

Learning Motor Skills: From Algorithms to Robot Experiments

Erlernen Motorischer Fähigkeiten: Von Algorithmen zu Roboter-Experimenten

Zur Erlangung des akademischen Grades Doktor-Ingenieur (Dr.-Ing.)

vorgelegte Dissertation von Dipl.-Ing. Jens Kober aus Künzelsau

März 2012 — Darmstadt — D 17



TECHNISCHE
UNIVERSITÄT
DARMSTADT



Learning Motor Skills:
From Algorithms to Robot Experiments
Erlernen Motorischer Fähigkeiten: Von Algorithmen zu Roboter-Experimenten

Vorgelegte Dissertation von Dipl.-Ing. Jens Kober aus Künzelsau

1. Gutachten: Prof. Dr. Jan Peters
2. Gutachten: Prof. Dr. Stefan Schaal

Tag der Einreichung:

Darmstadt — D 17

Bitte zitieren Sie dieses Dokument als:

URN: urn:nbn:de:tuda-tuprints-12345

URL: <http://tuprints.ulb.tu-darmstadt.de/12345>

Dieses Dokument wird bereitgestellt von tuprints,

E-Publishing-Service der TU Darmstadt

<http://tuprints.ulb.tu-darmstadt.de>

tuprints@ulb.tu-darmstadt.de



Die Veröffentlichung steht unter folgender Creative Commons Lizenz:

Namensnennung – Keine kommerzielle Nutzung – Keine Bearbeitung 2.0 Deutschland

<http://creativecommons.org/licenses/by-nc-nd/2.0/de/>

Erklärung zur Dissertation

Hiermit versichere ich, die vorliegende Dissertation ohne Hilfe Dritter nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die aus Quellen entnommen wurden, sind als solche kenntlich gemacht. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Darmstadt, den 13. März 2012

(J. Kober)



Abstract

Ever since the word “robot” was introduced to the English language by Karel Čapek’s play “Rossum’s Universal Robots” in 1921, robots have been expected to become part of our daily lives. In recent years, robots such as autonomous vacuum cleaners, lawn mowers, and window cleaners, as well as a huge number of toys have been made commercially available. However, a lot of additional research is required to turn robots into versatile household helpers and companions. One of the many challenges is that robots are still very specialized and cannot easily adapt to changing environments and requirements. Since the 1960s, scientists attempt to provide robots with more autonomy, adaptability, and intelligence. Research in this field is still very active but has shifted focus from reasoning based methods towards statistical machine learning. Both navigation (i.e., moving in unknown or changing environments) and motor control (i.e., coordinating movements to perform skilled actions) are important sub-tasks.

In this thesis, we will discuss approaches that allow robots to learn motor skills. We mainly consider tasks that need to take into account the dynamic behavior of the robot and its environment, where a kinematic movement plan is not sufficient. The presented tasks correspond to sports and games but the presented techniques will also be applicable to more mundane household tasks. Motor skills can often be represented by motor primitives. Such motor primitives encode elemental motions which can be generalized, sequenced, and combined to achieve more complex tasks. For example, a forehand and a backhand could be seen as two different motor primitives of playing table tennis. We show how motor primitives can be employed to learn motor skills on three different levels. First, we discuss how a single motor skill, represented by a motor primitive, can be learned using reinforcement learning. Second, we show how such learned motor primitives can be generalized to new situations. Finally, we present first steps towards using motor primitives in a hierarchical setting and how several motor primitives can be combined to achieve more complex tasks.

To date, there have been a number of successful applications of learning motor primitives employing imitation learning. However, many interesting motor learning problems are high-dimensional reinforcement learning problems which are often beyond the reach of current reinforcement learning methods. We review research on reinforcement learning applied to robotics and point out key challenges and important strategies to render reinforcement learning tractable. Based on these insights, we introduce novel learning approaches both for single and generalized motor skills.

For learning single motor skills, we study parametrized policy search methods and introduce a framework of reward-weighted imitation that allows us to derive both policy gradient methods and expectation-maximization (EM) inspired algorithms. We introduce a novel EM-inspired algorithm for policy learning that is particularly well-suited for motor primitives. We show that the proposed method out-performs several well-known parametrized policy search methods on an empirical benchmark both in simulation and on a real robot. We apply it in the context of motor learning and show that it can learn a complex ball-in-a-cup task on a real Barrett WAM.

In order to avoid re-learning the complete movement, such single motor skills need to be generalized to new situations. In this thesis, we propose a method that learns to generalize parametrized motor plans, obtained by imitation or reinforcement learning, by adapting a small set of global parameters. We employ reinforcement learning to learn the required parameters to deal with the current situation. Therefore, we introduce an appropriate kernel-based reinforcement learning algorithm. To show its feasibility, we evaluate this algorithm on a toy example and compare it to several previous approaches. Subsequently, we apply the approach to two robot tasks, i.e., the generalization of throwing movements in darts and of hitting movements in table tennis on several different real robots, i.e., a Barrett WAM, the JST-ICORP/SARCOS CBi and a Kuka KR 6.

We present first steps towards learning motor skills jointly with a higher level strategy and evaluate the approach with a target throwing task on a BioRob. Finally, we explore how several motor primitives,

representing sub-tasks, can be combined and prioritized to achieve a more complex task. This learning framework is validated with a ball-bouncing task on a Barrett WAM.

This thesis contributes to the state of the art in reinforcement learning applied to robotics both in terms of novel algorithms and applications. We have introduced the *Policy learning by Weighting Exploration with the Returns* algorithm for learning single motor skills and the *Cost-regularized Kernel Regression* to generalize motor skills to new situations. The applications explore highly dynamic tasks and exhibit a very efficient learning process. All proposed approaches have been extensively validated with benchmarks tasks, in simulation, and on real robots.

Zusammenfassung

Schon seit 1921 mit dem Theaterstück „Rossum’s Universal Robots“ von Karel Čapek das Wort „Roboter“ in die deutsche Sprache eingeführt wurde, besteht die Erwartung, dass Roboter Teil unseres täglichen Lebens werden. Seit ein paar Jahren sind sowohl Roboter wie autonome Staubsauger, Rasenmäher und Fensterreiniger als auch eine große Anzahl an Spielzeugrobotern im Handel erhältlich. Allerdings ist noch viel Forschung nötig, bis Roboter als universelle Haushalts-Helfer und Gefährten einsetzbar sind. Eine der größten Herausforderungen ist, dass Roboter immer noch sehr spezialisiert sind und sich nicht ohne Weiteres an sich ändernde Umgebungen und Anforderungen anpassen können. Seit den 1960ern versuchen Wissenschaftler, Roboter mit mehr Autonomie, Anpassungsfähigkeit und Intelligenz auszustatten. Die Forschung auf diesem Gebiet ist sehr aktiv, hat sich allerdings von regel-basierten Systemen hin zu statistischem maschinellen Lernen verlagert. Sowohl Navigation (d.h. sich in unbekannten oder sich ändernden Umgebungen zu bewegen) als auch Motorsteuerung (d.h. das Koordinieren von Bewegungen, um komplexe Aktionen auszuführen) sind hierbei wichtige Teilaufgaben.

In dieser Doktorarbeit diskutieren wir Ansätze, die es Robotern ermöglichen, motorische Fähigkeiten zu erlernen. Wir betrachten in erster Linie Aufgaben, bei denen das dynamische Verhalten des Roboters und seiner Umgebung berücksichtigt werden muss und wo ein kinematischer Bewegungsplan nicht ausreichend ist. Die vorgestellten Anwendungen kommen aus dem Sport- und Spiel-Bereich, aber die vorgestellten Techniken können auch bei alltäglichen Aufgaben im Haushalt Anwendung finden. Motorische Fähigkeiten können oft durch Motor-Primitive dargestellt werden. Solche Motor-Primitive kodieren elementare Bewegungen, die verallgemeinert, aneinandergereiht und kombiniert werden können, um komplexere Aufgaben zu erfüllen. Zum Beispiel könnte ein Vorhand- und Rückhand-Spiel als zwei verschiedene Motor-Primitive für Tischtennis angesehen werden. Wir zeigen, wie Motor-Primitive verwendet werden können, um motorische Fähigkeiten auf drei verschiedenen Ebenen zu erlernen. Zuerst zeigen wir, wie eine einzelne motorische Fertigkeit, die durch eine Motor-Primitive dargestellt wird, mittels Reinforcement-Learning (bestärkendes Lernen) gelernt werden kann. Zweitens zeigen wir, wie solche erlernten Motor-Primitiven verallgemeinert werden können, um auf neue Situationen zu reagieren. Schließlich präsentieren wir erste Schritte, wie Motor-Primitive in einer hierarchischen Struktur gelernt werden können und wie sich mehrere Motor-Primitive kombinieren lassen, um komplexere Aufgaben zu erfüllen.

Es gab schon eine Reihe von erfolgreichen Anwendungen des Erlernens von Motor-Primitiven durch überwachtes Lernen. Allerdings sind viele interessante motorische Lernprobleme hochdimensionale Reinforcement-Learning-Probleme, die oft außerhalb der Anwendbarkeit der aktuellen Reinforcement-Learning-Methoden liegen. Wir besprechen Publikationen, die sich mit Reinforcement-Learning, angewandt auf die Robotik, beschäftigen und zeigen sowohl zentrale Herausforderungen als auch Möglichkeiten, Reinforcement-Learning trotzdem anwenden zu können, auf. Basierend auf diesen Erkenntnissen stellen wir neuartige Lernansätze für einzelne motorische Fähigkeiten vor, auch um diese zu generalisieren.

Für das Erlernen einzelner motorischer Fähigkeiten untersuchen wir parametrisierte Policy-Search-Methoden und führen ein Framework der erfolgsgewichteten Nachahmung ein, das es uns ermöglicht, sowohl bekannte Policy-Search-Gradientenverfahren als auch neue Erwartungswert-Maximierung-inspirierte-Algorithmen herzuleiten. Wir stellen einen neuartigen EM-inspirierten Policy-Search-Algorithmus vor, der insbesondere für Motor-Primitive gut geeignet ist. Wir zeigen, dass das vorgeschlagene Verfahren mehrere bekannte parametrisierte Policy-Search-Methoden auf empirischen Benchmarks, in der Simulation und auf einem realen Roboter übertrifft. Wir wenden den Algorithmus für motorisches Lernen an und zeigen, dass eine komplexe Becherspiel-Aufgabe auf einem echten Barrett WAM gelernt werden kann.

Um zu vermeiden, dass die Bewegung komplett neu gelernt werden muss, brauchen wir eine Möglichkeit, motorische Fähigkeiten an neue Situationen anzupassen. In dieser Doktorarbeit schlagen wir eine Methode vor, die lernt, Motor-Primitive (die durch Nachahmung oder Reinforcement-Learning erlernt wurden)

durch das Anpassen einer kleinen Anzahl von globalen Parametern zu verallgemeinern. Wir verwenden Reinforcement-Learning, um die erforderlichen Parameter zu erlernen, die es ermöglichen, mit der aktuellen Situation umzugehen. Hierfür stellen wir einen kernel-basierten Reinforcement-Learning-Algorithmus vor. Um die Anwendbarkeit zu zeigen, testen wir diesen Algorithmus mit einem einfachen Beispiel und vergleichen ihn mit einigen bekannten Ansätzen. Anschließend wenden wir den Ansatz auf zwei Roboter-Aufgaben an, nämlich das Verallgemeinern von Dart-Würfen und Tischtennis-Schlägen mit verschiedenen realen Robotern: mit einem Barrett WAM, dem JST-ICORP/SARCOS CBI und einem KUKA KR 6 Roboter.

Wir präsentieren erste Schritte zum gleichzeitigen Erlernen motorischer Fähigkeiten und einer übergeordneten Strategie. Diesen Ansatz zeigen wir mit Hilfe einer Zielwurf-Aufgabe auf einem BioRob. Abschließend untersuchen wir, wie mehrere Motor-Primitive, die Teilaufgaben repräsentieren, mit unterschiedlicher Gewichtung kombiniert werden können, um eine komplexere Aufgabe zu erfüllen. Das Framework validieren wir, indem wir eine Tischtennisball-Jonglier-Aufgabe auf einem Barrett WAM erlernen.

Diese Doktorarbeit leistet einen Beitrag zum Stand der Technik in Reinforcement-Learning, angewandt auf die Robotik, sowohl durch neuartige Algorithmen als auch durch neuartige Anwendungen. Wir haben den „Policy learning by Weighting Exploration with the Returns“-Algorithmus zum Erlernen einzelner motorischer Fähigkeiten und die „Cost-regularized Kernel Regression“ zum Verallgemeinern motorischer Fähigkeiten auf neue Situationen vorgestellt. Die Anwendungen zeigen hochdynamische Aufgaben und weisen einen sehr effizienten Lernprozess aus. Alle vorgeschlagenen Ansätze wurden ausgiebig mit Benchmark-Aufgaben, in der Simulation und auf realen Robotern validiert.

Acknowledgments

I would like to thank:

- Jan Peters for his patience, support, and guidance
- Bernhard Schölkopf and Stefan Schaal for their support and for providing an ideal research environment
- my labmates Abdeslam Boularias, Oliver Kroemer, Timm Meyer, Katharina Mülling, Duy Nguyen-Tuong, and Zhikun Wang for their help, uncountable discussions, and valuable feedback
- Erhan Oztop and Mitsuo Kawato for hosting me at ATR
- Michael Mistry and Jessica Hodgins for hosting me at Disney Research Pittsburgh
- Andreas Wilhelm and Konrad Wöhlhaf for hosting me at University of Applied Sciences Ravensburg-Weingarten
- my co-authors Drew Bagnell, Silvia Chiappa, Christoph Lampert, and Betty Mohler
- all colleagues and visitors at ABGS from whom I had the opportunity to learn a lot about machine learning beyond the scope of this thesis
- all colleagues at IAS for valuable discussions
- my family and Moritz Schöpfer for their encouragement

without all these people this thesis would not have been possible.



Contents

Abstract	iii
Zusammenfassung	v
Acknowledgments	vii
1 Introduction	1
1.1 Motivation	1
1.2 Contributions	2
1.2.1 Algorithms	2
1.2.2 Applications	3
1.3 Outline	4
2 Reinforcement Learning in Robotics: a Survey	7
2.1 Introduction	7
2.2 Challenges in Robot Reinforcement Learning	8
2.2.1 Curse of Dimensionality	8
2.2.2 Curse of Real-World Samples	10
2.2.3 Curse of Under-Modeling	11
2.2.4 Curse of Goal Specification	12
2.3 Foundations of Robot Reinforcement Learning	12
2.3.1 Value Function Approaches	13
2.3.2 Policy Search	14
2.4 Tractability through Representation	16
2.4.1 Smart State-Action Discretization	16
2.4.2 Function Approximation	18
2.4.3 Pre-structured Policies	18
2.5 Tractability through Prior Knowledge	19
2.5.1 Prior Knowledge through Demonstrations	20
2.5.2 Prior Knowledge through Task Structuring	20
2.5.3 Directing Exploration with Prior Knowledge	21
2.6 Tractability through Models	21
2.6.1 Role of Models	22
2.6.2 Mental Rehearsal	22
2.6.3 Direct Transfer from Simulated to Real Robots	23
2.7 A Case Study: Ball-in-a-Cup	24
2.7.1 Experimental Setting: Task and Reward	24
2.7.2 Appropriate Policy Representation	25
2.7.3 Generating a Teacher’s Demonstration	26
2.7.4 Reinforcement Learning by Policy Search	26
2.7.5 Use of Simulations in Robot Reinforcement Learning	27
2.7.6 Alternative Approach with Value Function Methods	27
2.8 Discussion	27
2.9 Next Challenges	28
2.10 Conclusion	28

3	Movement Templates for Learning of Hitting and Batting	31
3.1	Introduction	31
3.2	Movement Templates for Learning to Strike	32
3.2.1	Discrete Movement Primitives	33
3.2.2	Adapting the Motor Primitives for Striking Movements	34
3.2.3	Safer Dynamics for Generalization	36
3.2.4	Imitation Learning	37
3.3	Robot Evaluation	38
3.3.1	Generalizing Forehands on Static Targets	38
3.3.2	Playing against a Ball Launcher	39
3.4	Conclusion	40
4	Policy Search for Motor Primitives in Robotics	41
4.1	Introduction	41
4.2	Policy Search for Parametrized Motor Primitives	42
4.2.1	Problem Statement & Notation	42
4.2.2	Episodic Policy Learning	43
4.2.3	Policy learning by Weighting Exploration with the Returns (PoWER)	47
4.3	Benchmark Evaluation and Application in Robotics	49
4.3.1	Dynamical Systems Motor Primitives as Stochastic Policies	50
4.3.2	Benchmark Comparison I: Basic Motor Learning Examples	50
4.3.3	Benchmark Comparison II: Mountain-Car	53
4.3.4	Benchmark Comparison III: Tetherball Target Hitting	54
4.3.5	Benchmark Comparison IV: Underactuated Swing-Up	55
4.3.6	Benchmark Comparison V: Casting	57
4.3.7	Ball-in-a-Cup on a Barrett WAM	58
4.4	Discussion & Conclusion	61
4.4.1	Discussion: Robotics as Benchmark for Reinforcement Learning?	61
4.4.2	Discussion on Simulation to Robot Transfer Scenarios	62
4.4.3	Conclusion	65
4.A	Derivations	65
4.A.1	REINFORCE	65
4.A.2	Episodic Reward Weighted Regression (eRWR)	66
4.A.3	EM Policy learning by Weighting Exploration with the Returns (PoWER)	66
5	Reinforcement Learning to Adjust Parametrized Motor Primitives to New Situations	69
5.1	Introduction	69
5.2	Meta-Parameter Learning for Motor Primitives	71
5.2.1	Motor Primitives with Meta-Parameters	71
5.2.2	Problem Statement: Meta-Parameter Self-Improvement	73
5.2.3	A Task-Appropriate Reinforcement Learning Algorithm	74
5.2.4	Meta-Parameter Learning by Reinforcement Learning	75
5.3	Evaluations and Experiments	77
5.3.1	Benchmark Comparison: Toy Cannon Shots	77
5.3.2	Robot Dart-Throwing Games	80
5.3.3	Robot Table Tennis	82
5.3.4	Active Learning of Ball Throwing	84
5.4	Conclusion & Future Work	91

6	Learning Prioritized Control of Motor Primitives	93
6.1	Introduction	93
6.1.1	Proposed Approach	94
6.2	Learning the Prioritized Control Law	95
6.2.1	Single Primitive Control Law	96
6.2.2	Prioritized Primitives Control Law	97
6.3	Evaluation: Ball-Bouncing	98
6.3.1	Task Description	99
6.3.2	Bouncing Strategy	99
6.3.3	Learning Results	100
6.4	Conclusion	101
7	Conclusion	103
7.1	Summary	103
7.1.1	Algorithms	103
7.1.2	Applications	104
7.2	Open Problems	104
7.3	Publications	107
	Bibliography	111
	List of Figures	127
	List of Algorithms	133
	List of Tables	135
	Abbreviations	137
	Curriculum Vitæ	139



1 Introduction

Since Issac Asimov started to write short stories about robots in the 1940s, the idea of robots as household helpers, companions and soldiers has shaped the popular view of robotics. Science fiction movies depict robots both as friends and enemies of the human race, but in both cases their capabilities far exceed the capabilities of current real robots. Simon [1965], one of the artificial intelligence (AI) research pioneers, claimed that “machines will be capable, within twenty years, of doing any work a man can do.” This over-optimistic promise led to more conservative predictions. Nowadays robots slowly start to penetrate our daily lives in the form of toys and household helpers, like autonomous vacuum cleaners, lawn mowers, and window cleaners. Most other robotic helpers are still confined to research labs and industrial settings. Many tasks of our daily lives can only be performed very slowly by a robot which often has very limited generalization capabilities. Hence, all these systems are still disconnected from the expectation raised by literature and movies as well as from the dreams of AI researchers.

Especially in Japan, the need of robotic household companions has been recognized due to the aging population. One of the main challenges remains the need to adapt to changing environments in a co-inhabited household (e.g., furniture being moved, changing lighting conditions) and the need to adapt to individual requirements and expectations of the human owner. Most current products either feature a “one size fits all” approach that often is not optimal (e.g., vacuum robots that are not aware of their surrounding but use an approach for obstacle treatment, that guarantees coverage of the whole floor [BotJunkie, 2012]) or an approach that requires a setup step either in software (e.g. providing a floor map) or in hardware (e.g., by placing beacons). As an alternative one could imagine a self-learning system. In this thesis, we will not treat navigation problems but rather focus on learning motor skills [Wulf, 2007]. We are mainly interested in motor skills that need to take into account the dynamics of the robot and its environment. For these motor skills, a kinematic plan of the movement will not be sufficient to perform the task successfully. A motor skill can often be represented by a motor primitive, i.e., a representation of a single movement that is adapted to varying situations (e.g., a forehand or a backhand in table tennis that is adapted to the ball position and velocity). We focus on learning how to perform a motor primitive optimally and how to generalize it to new situations. The presented tasks correspond to games and sports, which are activities that a user might enjoy with a robot companion, but the presented techniques could also be applied to more mundane household tasks.

1.1 Motivation

Machine-learning research has resulted in a huge number of algorithms. Unfortunately most standard approaches are not directly applicable to robotics, mainly due to the inherent high dimensionality. Hence there is a need to develop approaches specifically tailored for the needs of robot learning. In this thesis we focus on reinforcement learning of motor primitives.

Research in reinforcement learning started in the 1950s (e.g., [Minsky, 1954, Farley and Clark, 1954, Bellman, 1957]) but has been mainly focused on theoretical problems. The algorithms are often evaluated on synthetic benchmark problems involving discrete states and actions. Probably the best known real-world applications of reinforcement learning are games, like backgammon [Tesauro, 1994] or Go [Chan et al., 1996], but also robotic applications can be found from the 1990s on (e.g., [Mahadevan and Connell, 1992, Gullapalli et al., 1994]). In contrast to many problems studied in the reinforcement learning literature, robotic problems have inherently continuous states and actions. Furthermore, experiments in robotics often deal with expensive and potentially fragile hardware and often require human supervision and intervention. These differences result in the need of adapting existing reinforcement learning approaches or developing tailored ones.

Policy search, also known as policy learning, has become an accepted alternative of value function-based reinforcement learning [Strens and Moore, 2001, Kwee et al., 2001, Peshkin, 2001, Bagnell et al., 2004, El-Fakdi et al., 2006, Taylor et al., 2007]. Especially for learning motor skills in robotics, searching directly for the policy instead of solving the dual problem (i.e., finding the value function) has been shown to be promising, see Section 2.3 for a more detailed discussion. Additionally, incorporating prior knowledge in the form of the policy structure, an initial policy, or a model of the system can drastically reduce the learning time. In this thesis we will focus on policy search methods that employ a pre-structured policy and an initial policy.

1.2 Contributions

This thesis contributes to the state of the art in reinforcement learning applied to robotics both in terms of novel algorithms (Section 1.2.1) and applications (Section 1.2.2).

1.2.1 Algorithms

Reinforcement learning applied to robotics faces a number of challenges as discussed in Section 2.2. The dimensionality is inherently high and continuous. Acquiring real world samples is expensive as they require time, human supervision and potentially expensive and fragile hardware. Using models can alleviate these problems but poses different challenges. Every learning algorithm is limited by its goal specifications. Unfortunately specifying a cost or a reward function is not straightforward. Even the cost function of a simple human reaching movement is not completely understood yet [Bays and Wolpert, 2007]. Inverse reinforcement learning is a promising alternative to specifying the reward function manually and has been explored for the Ball-in-a-Cup task (Section 4.3.7) in [Boularias et al., 2011].

Ideally an algorithm applied to robotics should be safe, i.e., avoid damaging the robot, and it should be fast, both in terms of convergence and computation time. Having a sample-efficient algorithm, with very few open parameters, and the ability to incorporate prior knowledge all contribute significantly to fast convergence. In this thesis we propose a framework of reward-weighted imitation that fits these requirements. In Section 7.1.1, we will review how the proposed algorithms meet these requirements.

Policy learning by Weighting Exploration with the Returns (PoWER)

In Chapter 4, we introduce the Policy learning by Weighting Exploration with the Returns (PoWER) algorithm. PoWER is an expectation-maximization inspired policy search algorithm that is based on structured exploration in the parameter space. In Chapter 4, we derive a framework of reward-weighted imitation. Based on [Dayan and Hinton, 1997], we consider the return of an episode as an improper probability distribution. We maximize a lower bound of the logarithm of the expected return. Depending on the strategy of optimizing this lower bound and the exploration strategy, the framework yields several well known policy search algorithms: episodic REINFORCE [Williams, 1992], the policy gradient theorem [Sutton et al., 1999], episodic natural actor critic [Peters and Schaal, 2008c], as well as a generalization of the reward-weighted regression [Peters and Schaal, 2008b]. Our novel algorithm, PoWER, is based on an expectation-maximization inspired optimization and a structured, state-dependent exploration. Our approach has already given rise to follow-up work in other contexts, for example, [Vlassis et al., 2009, Kormushev et al., 2010]. Theodorou et al. [2010] have shown that an algorithm very similar to PoWER can also be derived from a completely different perspective, that is, the path integral approach.

Cost-regularized Kernel Regression (CrKR)

In Chapter 5, we introduce the algorithm Cost-regularized Kernel Regression (CrKR). CrKR is a non-parametric policy search approach that is particularly suited for learning meta-parameters, i.e., a limited set of parameters that influence the movement globally. In this setting, designing good parametrized policies can be challenging, a non-parametric policy offers more flexibility. We derive CrKR based on the reward-weighted regression [Peters and Schaal, 2008b, Chapter 4]. The resulting algorithm is related to Gaussian process regression and similarly yields a predictive variance, which is employed to guide the exploration in on-policy reinforcement learning. This approach is used to learn a mapping from situation to meta-parameters while the parameters of the motor primitive can still be acquired through traditional approaches.

1.2.2 Applications

In this thesis we show a large number of benchmark task and evaluate the approaches on robotic tasks both with simulated and real robots. The employed robots include a Barrett WAM, a BioRob, the JST-ICORP/SARCOS CBi and a Kuka KR 6. This thesis studies single motor skills and how these can be generalized to new situations. The presented work has contributed to the state-of-the-art of reinforcement learning applied to robotics by exploring highly dynamic tasks and exhibiting a very efficient learning process.

Single Motor Skills

Because of the curse of dimensionality, we cannot start with an arbitrary solution to learn a motor skill. Instead, we mimic the way children learn and first present an example movement for imitation learning, which is recorded using, e.g., kinesthetic teach-in. Subsequently, our reinforcement algorithm learns how to perform the skill reliably. After only realistically few episodes, the task can be regularly fulfilled and the robot shows very good average performance. We demonstrate this approach with a number of different policy representations including dynamical systems motor primitives [Ijspeert et al., 2002b] and other parametric representations. We benchmark PoWER against a variety of policy search approaches on both synthetic and robotic tasks. Finally, we demonstrate that single movement skills like the Underactuated Swing-Up and Ball-in-a-Cup can be learned on a Barrett WAM efficiently.

Generalized Motor Skills

In order to increase the efficiency of a learning process it is often advantageous to generalize a motor skill to new situations instead of re-learning it from scratch. This kind of learning often requires a non-intuitive mapping from situation to actions. To demonstrate the proposed system in a complex scenario, we have chosen the *Around the Clock* dart throwing game, table tennis, and ball throwing implemented both on simulated and real robots. In these scenarios we show that our approach performs well in a wide variety of settings, i.e. on four different real robots (namely a Barrett WAM, a BioRob, the JST-ICORP/SARCOS CBi and a Kuka KR 6), with different cost functions (both with and without secondary objectives), and with different policies in conjunction with their associated meta-parameters. The ball throwing task presents first steps towards a hierarchical reinforcement learning system.

Many robotic tasks can be decomposed into sub-tasks. However, often these sub-tasks cannot be achieved simultaneously and a dominance structure needs to be established. We evaluate initial steps towards a control law based on a set of prioritized motor primitives with a ball-bouncing task on a Barrett WAM.




Figure 1.1: This figure illustrates the structure of this thesis. Chapter 2 reviews the field of reinforcement learning applied to robotics, Chapter 3 describes the employed motor skill representation. Chapter 4 provides a framework for policy search approaches. Chapter 5 describes how movements can be generalized to new situations. Chapter 6 outlines how several motor skills can be performed simultaneously to achieve a more complex task. Chapter 1 gives an introduction and Chapter 7 provides a summary and an outlook on future work.

1.3 Outline

The chapters in this thesis can largely be read independently but partially build upon results of the preceding chapters. Figure 1.1 illustrates the outline of this thesis. Chapter 3 describes the employed motor primitive representation that is used as policy parametrization for the evaluations throughout the thesis. The theoretical contributions are applicable to a wider variety of policies. Chapter 4 describes how single movement skills can be learned using policy search and employing a parametrized policy that is linear in parameters. Chapter 5 explains how single movements can be generalized to new situations using a nonparametric policy representation. The initial movement is obtained using results from Chapters 3 and 4.

Chapter 2 : provides a survey on reinforcement learning applied to robotics. The survey provides an overview of techniques and applications. The particular focus lies on challenges specific to reinforcement learning in robotics and approaches to render the learning problem tractable nevertheless. This chapter is based on [Kober et al., 2012a].

Chapter 3 : describes how the [Ijspeert et al., 2002b] representation for motor skills can be generalized for hitting and batting movements. This chapter is based on [Kober et al., 2010a].

Chapter 4: discusses a framework of reward-weighted imitation that allows us to re-derive several well-know policy gradient approaches and to derive novel EM-inspired policy search approaches. The resulting algorithm, PoWER is evaluated in a number of benchmark and robotic tasks. This chapter is based on [Kober and Peters, 2011a] and a preliminary version of some of the work in this chapter was shown in [Kober et al., 2008, Kober and Peters, 2008, 2009, 2010].

Chapter 5: discusses how behaviors acquired by imitation and reinforcement learning as described in Chapters 3 and 4 can be generalized to novel situations using the Cost-regularized Kernel Regression. We also discuss first steps towards a hierarchical framework. The approach is evaluated on several robotic platforms. This chapter is based on [Kober et al., 2012b] and a preliminary version of some of the work in this chapter was shown in [Kober et al., 2010b, 2011, Kober and Peters, 2011b].

Chapter 6: discusses first steps towards performing sub-tasks simultaneously to achieve a task. We discuss a novel learning approach that allows to learn a prioritized control law built on a set of sub-tasks represented by motor primitives. The approach is evaluated with a ball bouncing task on a Barrett WAM. This chapter is based on [Kober and Peters, 2012].

Chapter 7: concludes the thesis with a summary and gives an overview of open problems and an outlook on future work.



2 Reinforcement Learning in Robotics: a Survey

Robotics offers a tremendously interesting platform for the application and impact of reinforcement learning; conversely, the challenges of robotics provide both inspiration and validation for reinforcement learning. The relationship between disciplines has enough promise to be likened to that as between physics and mathematics. Nevertheless, only a fraction of the scientists working on reinforcement learning are sufficiently tied to robotics to oversee most problems encountered in this context. Here we attempt to highlight both important challenges in robot reinforcement learning and important success by surveying work that has successfully applied reinforcement learning to behavior generation for real robots. We discuss how each work tamed the complexity of the domain and study the role of algorithms, representations and prior knowledge in achieving success. As a result, a particular focus of our chapter lies on the choice between model-based and model-free as well as between value function-based and policy search methods.

2.1 Introduction

Robotics has a near infinite amount of interesting learning problems, a large percentage of which can be phrased as reinforcement learning problems. See Figure 2.1 for an illustration of the wide variety of robots that have learned tasks using reinforcement learning. However, robotics as a domain differs significantly from well-defined typical reinforcement learning benchmark problems, which usually have discrete states and actions. In contrast, many real-world problems in robotics are best represented with high-dimensional, continuous states and actions. Every single trial run is costly and, as a result, such applications force us to focus on problems that do not arise that frequently in classical reinforcement learning benchmark examples. In this chapter, we highlight the challenges faced in robot reinforcement learning and bring many of the inherent problems of this domain to the reader's attention.

Robotics is characterized by inherently continuous state and action spaces which can be high dimensional. In order to learn in a reasonable time frame, suitable approximations need to be introduced. Real-world measurements are intrinsically noisy and often not all states can be observed. This inability to measure gives rise to uncertainty in the state of the system. Experience on the real system is costly and often hard to reproduce. However, it usually cannot be replaced by learning in simulations alone. In analytical or learned models of the system even small modeling errors accumulate to substantially different dynamic behavior, at least for highly dynamic tasks. Hence, the algorithms need to be robust with respect to under-modeling. Another challenge faced in robot reinforcement learning is the generation of appropriate reward functions. Good rewards that lead the systems quickly to success are needed to cope with the cost of real-world experience but are a substantial manual contribution.

Obviously, not every reinforcement learning method is equally suitable for the robotics domain. In fact, many of the methods that scale to more interesting tasks are model-based [Atkeson et al., 1997, Abbeel et al., 2007] and often employ policy search rather than value function-based approaches [Gullapalli et al., 1994, Miyamoto et al., 1996, Kohl and Stone, 2004, Tedrake et al., 2005, Peters and Schaal, 2008b,c, Kober and Peters, 2008]. This stands in contrast to much of mainstream reinforcement [Kaelbling et al., 1996, Sutton and Barto, 1998]. We attempt to give a fairly complete overview on real robot reinforcement learning citing most original papers while distinguishing mainly on a methodological level.

As none of the presented methods extends to robotics with ease, we discuss how robot reinforcement learning can be made tractable. We present several approaches to this problem such as choosing an appropriate representation for your value function or policy, incorporating prior knowledge, and transfer from simulations.

In this chapter, we survey real robot reinforcement learning and highlight how these approaches were able to handle the challenges posed by this setting. We focus mainly on results that were obtained on

physical robots with tasks going beyond typical reinforcement learning benchmarks. The goal of this chapter is twofold. This chapter surveys a wide variety of tasks where reinforcement learning has been successfully applied to robotics. Nearly any problem that can be phrased as an optimization problem, could be attempted to be solved using reinforcement learning. On the one hand, we hope that this chapter can provide indications for the robotics community which type of problems can be tackled by reinforcement learning and provide pointers to approaches that are promising. On the other hand, this chapter can point out novel real-world test beds and related interesting open questions for the reinforcement learning community. The challenges in applying reinforcement learning in robotics are discussed in Section 2.2.

Standard reinforcement learning methods suffer from the discussed challenges. As already pointed out in the reinforcement learning review paper by Kaelbling et al. [1996] “we must give up tabula rasa learning techniques and begin to incorporate bias that will give leverage to the learning process”. Hence, we concisely present reinforcement learning techniques in the context of robotics in Section 2.3. Different approaches of making reinforcement learning tractable are treated in Sections 2.4 to 2.6. Finally in Section 2.7, the example of ball-in-a-cup is employed to highlight which of the various approaches discussed in the chapter have been particularly helpful to make such a complex task tractable. In Section 2.8 we summarize the specific problems and benefits of reinforcement learning in robotics, in Section 2.9 we give an outlook on interesting problems, and in Section 2.10, we give a conclusion.


2.2 Challenges in Robot Reinforcement Learning

Reinforcement learning is generically a hard problem and many of its challenges apply particularly in the robotics setting. As the states and actions of most robots are inherently continuous we have to deal with the resolution. Additionally the dimensionality can be high and we face the ‘Curse of Dimensionality’ [Bellman, 1957] as discussed in Section 2.2.1. As we deal with complex physical systems, samples can be expensive due to the long execution time of complete tasks, required manual interventions as well as maintenance and repair. Due to the real-world measurements we have to cope with uncertainty of the system. A robot requires that the algorithm runs in real-time and that it is capable of dealing with delays in the sensing and execution which are inherent in physical systems (see Section 2.2.2). Obviously, a simulation could alleviate many problems but the approaches need to be robust with respect to model errors as discussed in Section 2.2.3. An often underestimated problem is the goal specification, which is achieved by designing a good reward function. As noted in Section 2.2.4 this choice can make the difference between feasibility and an unreasonable amount of exploration.


2.2.1 Curse of Dimensionality

When Bellman [1957] explored optimal control in higher dimensions, he faced an exponential explosion of states and actions for which he coined the term ‘Curse of Dimensionality’. Robotic systems often have to deal with such high dimensional states and actions, e.g., due to the many degrees of freedom of modern anthropomorphic robots. For example, in a ball-paddling task as shown in Figure 2.2, a proper representation of a robot’s state would consist of its joint angles and velocities for each of its seven degrees of freedom as well as Cartesian position and velocity of the ball. The robot’s actions would be the generated motor commands which often are torques or accelerations. In this example, we have $2 \times (7 + 3) = 20$ state dimensions and 7-dimensional continuous actions. Obviously, other tasks may require even more dimensions, e.g., human-like actuation often follows the antagonistic principle [Yamaguchi and Takanishi, 1997] which additionally allows to control stiffness. Obviously such dimensionality is a major challenge for both the robotics and the reinforcement learning communities.

In robotics, tasks with such problems are often made more accessible to the robot engineer by shifting some of the complexity to a lower layer of functionality. In the ball paddling example, we can simplify by controlling the robot in racket space (which is lower-dimensional as the racket is orientation-invariant




(a) OBELIX robot



(b) Zebra Zero robot



(c) An autonomous blimp



(d) Sarcos humanoid DB

Figure 2.1: This figure illustrates robots to which reinforcement learning has been applied. The robots cover the whole range of wheeled mobile robots, robotic arms, autonomous vehicles, to humanoid robots. (a) The OBELIX robot is a wheeled mobile robot that learned to push boxes [Mahadevan and Connell, 1992] with a value function-based approach (Picture reprint with permission of Sridhar Mahadevan). (b) The Zebra Zero robot is a robot arm that learned a peg-in-hole insertion task [Gullapalli et al., 1994] with a model-free policy gradient approach (Picture reprint with permission of Rod Grupen). (c) The control of such autonomous blimps (Picture reprint with permission of Axel Rottmann) was learned with both a value function based approach [Rottmann et al., 2007] and model-based policy search [Ko et al., 2007]. (d) The Sarcos humanoid DB learned a pole-balancing task [Schaal, 1996] using forward models (Picture reprint with permission of Stefan Schaal).




Figure 2.2: This Figure illustrates the state space of a robot reinforcement learning task.

around the string’s mounting point) with an operational space control law [Nakanishi et al., 2008]. Many commercial robot systems also encapsulate some of the state and action components in an embedded control system (e.g., trajectory fragments are frequently used as actions for industrial robots); however, this form of a state dimensionality reduction severely limits the dynamic capabilities of the robot according to our experience [Schaal et al., 2002, Peters et al., 2010b].

The reinforcement learning community has a long history of dealing with dimensionality using computational abstractions. It offers a larger set of applicable tools ranging from adaptive discretizations [Buşoniu et al., 2010] over function approximation approaches [Sutton and Barto, 1998] to macro actions or options [Barto and Mahadevan, 2003]. Macro actions allow decomposing a task in elementary components and quite naturally translate to robotics. For example, a macro action “move one meter to the left” could be achieved by a lower level controller that takes care of accelerating, moving, and stopping while ensuring the precision. Using a limited set of manually generated macro actions, standard reinforcement learning approaches can be made tractable for navigational tasks for mobile robots. However, the automatic generation of such sets of macro actions is the key issue in order to enable such approaches. We will discuss approaches that have been successful in robot reinforcement learning in Section 2.4.

2.2.2 Curse of Real-World Samples

Robots inherently interact with the real-world and, hence, robot reinforcement learning suffers from most of the resulting real-world problems. For example, robot hardware is usually expensive, suffers from wear and tear, and requires careful maintenance. Repairing a robot system is a non-negligible effort associated with cost, physical labor and long waiting periods. Hence, to apply reinforcement learning in robotics, safe exploration becomes a key issue of the learning process [Schneider, 1996, Bagnell, 2004]; a problem often neglected in the general reinforcement learning community.

However, several more aspects of the real-world make robotics a challenging domain. As the dynamics of a robot can change due to many external factors ranging from temperature to wear, the learning process may never fully converge, i.e., it needs a ‘tracking solution’ [Sutton et al., 2007]. Frequently, the environment settings during an earlier learning period cannot be reproduced and the external factors are not clear, e.g., how did the light conditions affect the performance of the vision system and, as a result,

the task's performance. This problem makes comparisons of algorithms particularly hard. Furthermore, the approaches often have to deal with uncertainty due to inherent measurement noise and the inability to observe all states directly with sensors.

Most real robot learning tasks require some form of human supervision, e.g., putting the pole back on the robot's end-effector during pole balancing, see Figure 2.1d, after a failure. Even when an automatic reset exists (e.g., by having a smart contraption that resets the pole), learning speed becomes essential as a task on a real robot cannot be sped up. The whole episode needs to be complete as it is often not possible to start from arbitrary states.

For such reasons, real-world samples are expensive in terms of time, labor and, potentially, finances. In robotic reinforcement learning, it is often considered to be more important to limit the real-world interaction time instead of limiting the memory consumption or computational complexity. Thus, sample efficient algorithms that are able to learn from a small number of trials are essential. In Sections 2.6.2 and 2.6.3 we will discuss several approaches that allow reducing the amount of required real-world interactions.

As the robot is a physical system there are strict constraints on the interaction between the learning algorithm and the robot setup. Usually the robot needs to get commands at fixed frequency and for dynamic tasks the movement cannot be paused. Thus, the agent has to select actions in real-time. It is often not possible to pause to think, learn or plan between each action but rather the learning algorithm has to deal with a fixed amount of time. Thus, not only are samples expensive to obtain, but also often only a very limited number of samples can be used, if the runtime of the algorithms depends on the number of samples. These constraints are less severe in an episodic setting where the time intensive part of the learning can be postponed to the period between episodes.

On physical systems there are always delays in sensing and actuation. The state of the setup represented by the sensors slightly lags behind the real state due to processing and communication delays. More critically there are also communication delays in the actuation as well as delays due to the fact that a physical system cannot instantly change its movement. For example, a car will not stop instantly when you step on the break. Due to these delays, actions do not have instantaneous effects but are observable only several time steps later. In contrast, in most general reinforcement learning algorithms, the actions are assumed to take effect instantaneously.

2.2.3 Curse of Under-Modeling

One way to offset the cost of real-world interaction would be accurate models that are being used as simulators. In an ideal setting, such an approach would render it possible to learn the behavior in simulation and subsequently transfer it to the real robot. Unfortunately, creating a sufficiently accurate model of the robot and the environment is challenging. As small model errors may accumulate, we can frequently see a fast divergence of the simulated robot from the real-world system. When a policy is trained using an imprecise forward model as simulator, the behavior will not transfer without significant modifications as experienced by Atkeson [1994] when learning the underactuated swing-up. Only in a limited number of experiments, the authors have achieved such a direct transfer, see Section 2.6.3 for examples. If the task is inherently stable it is safer to assume that approaches that can be applied in simulation may work similarly in the real-world [Kober and Peters, 2010].

In such stable scenarios, transferring policies poses a low risk of damaging the robot. Nevertheless, tasks can often be learned better in the real-world than in simulation due to complex interactions between mechanical objects. For unstable tasks transferring policies poses a high risk. As we will see later, models are best used to test the algorithms in simulations but can also be used to check the proximity to theoretically optimal solutions, to calculate approximate gradients, or to perform 'mental rehearsal'.

2.2.4 Curse of Goal Specification

In reinforcement learning, the goal of the task is implicitly specified by the reward. Defining a good reward function in robot reinforcement learning is hence often a daunting task. Giving rewards only upon task achievement, e.g., did a table tennis robot win the match, will result in a simple binary reward. However, the robot would receive such a reward so rarely that it is unlikely to ever succeed in the lifetime of a real-world system. Hence, instead of using only simpler binary rewards, we frequently need to include additional knowledge into such scalar rewards to guide the learning process to a reasonable solution. The trade-off between different factors may be essential as hitting a table tennis ball very fast will result in a high score but is likely to destroy the robot. Similarly, changes in actions can be penalized to avoid high frequency control. Good reinforcement learning algorithms often exploit the reward function in unexpected ways, especially if the reinforcement learning is done locally and not globally. For example, if the distance between the ball and the desired highest point is part of the reward in ball paddling (see Figure 2.2), many locally optimal solutions would attempt to simply move the racket upwards and keep the ball on it. Reward shaping is employed to have a notion of closeness to the desired behavior instead of relying on a reward that only encodes success or failure. Furthermore, robotic reinforcement learning is often interested in local optimization. In conjunction with local approximation many techniques from the optimal control literature can be employed.

Often the desired behavior can be represented in a high dimensional state-space which is unfeasible for learning. In such cases the reward artfully specified in the features of the lower dimensional space in which the learning algorithm operates. There is a trade-off between the complexity of the reward function and the complexity of the learning problem. For example, crusher [Ratliff et al., 2006a], a outdoor robot, reasons about the world on a long time horizon scale as if it is a holonomic robot living on a fine grid of costs. The human operator is interested in a combination of time and risk, which cannot be represented straightforwardly in this state-space. Nevertheless, by carefully mapping features to cost, a remarkably human like behavior that seems to respect time and risk priorities can be achieved.

Inverse optimal control, also known as inverse reinforcement learning, is a promising alternative to specifying the reward function manually. Instead, it assumes that a reward function can be reconstructed from a set of expert demonstrations. This reward function does not necessarily correspond to the true reward function, but provides guarantees on the resulting performance. Inverse optimal control was initially studied in the control community [Kalman, 1964] and in the field of economics [Keeney and Raiffa, 1976]. The initial results were only applicable to linear quadratic regulator problems and required required closed form access to plant and controller, hence samples from human demonstrations could not be used. Russell [1998] brought the field to the attention of the machine learning community. Abbeel and Ng [2004] defined the constraint that no solution can be better than the demonstrated ones under the recovered reward function. Ratliff et al. [2006b] introduced a margin to deal with instabilities due to noise and Ziebart et al. [2008] rendered the idea robust and probabilistic, making inverse optimal control applicable to robotics. It has been recently successfully applied to outdoor robot navigation [Ratliff et al., 2006a, Silver et al., 2008], manipulation [Ratliff et al., 2007], and quadruped locomotion [Ratliff et al., 2006a, 2007, Kolter et al., 2007]. Please note that apprenticeship learning as defined by Pieter Abbeel rather corresponds to imitation learning with subsequent refinement by reinforcement learning as discussed in Section 2.5.1.

2.3 Foundations of Robot Reinforcement Learning

Real-world domains such as robotics are affected more strongly by the basic approach choices than synthetic benchmark tasks. Hence, we introduce reinforcement learning in this chapter with a particular point of view. The goal of reinforcement learning is to find a policy $\pi(s, a)$ that gathers maximal rewards $R(s, a)$. However, in real-world domains the average reward is often more suitable than a discounted formulation due to its stability properties [Peters et al., 2004]. In order to incorporate exploration,

the policy is considered a conditional probability distribution $\pi(s, a) = f(a|s, \theta)$ with parameters θ . Reinforcement learning aims at finding the optimal policy π^* or equivalent policy parameters θ^* which maximize the average return $J(\pi) = \sum_{s,a} \mu^\pi(s) \pi(s, a) R(s, a)$ where μ^π is the stationary state distribution generated by policy π acting in the environment $T(s, a, s') = P(s'|s, a)$. Hence, we have an optimization problem of

$$\max_{\pi} J(\pi) = \sum_{s,a} \mu^\pi(s) \pi(s, a) R(s, a), \quad (2.1)$$

$$\text{s.t. } \mu^\pi(s') = \sum_{s,a} \mu^\pi(s) \pi(s, a) T(s, a, s'), \quad \forall s' \in S, \quad (2.2)$$

$$1 = \sum_{s,a} \mu^\pi(s) \pi(s, a). \quad (2.3)$$

Here, Equation (2.2) defines stationarity of the state distributions μ^π (i.e., it ensures that it converges) and Equation (2.3) ensures a proper state-action probability distribution. This optimization problem can be tackled in two substantially different ways [Bellman, 1967, 1971], i.e., we can search the optimal solution directly in the original, primal problem, and we can optimize in the dual formulation. Optimizing in primal formulation is known as *policy search* in reinforcement learning while searching in the dual is called a *value function-based approach*.

2.3.1 Value Function Approaches

Most of reinforcement has focused on solving the optimization problem in Equations (2.1-2.3) not directly but rather in its dual form. Using the Lagrangian multipliers $V(s')$ and \bar{R} , we can express the Lagrangian of the problem by

$$L = \sum_{s,a} \mu^\pi(s) \pi(s, a) \left[R(s, a) + \sum_{s'} V(s') T(s, a, s') - \bar{R} \right] - \sum_{s'} V(s') \mu^\pi(s') + \bar{R}.$$

Using straightforward insights and the stationarity condition $\sum_{s',a'} V(s') \mu^\pi(s') \pi(s', a') = \sum_{s,a} V(s) \mu^\pi(s) \pi(s, a)$, we can obtain the Karush-Kuhn-Tucker conditions [Kuhn and Tucker, 1950] by differentiating with respect to $\mu^\pi(s) \pi(s, a)$ which yields

$$\partial_{\mu^\pi \pi} L = R(s, a) + \sum_{s'} V(s') T(s, a, s') - \bar{R} - V(s) = 0.$$

As this implies that there are as many equations the number of states multiplied by the number of actions, it is clear that only one action a^* can be optimal. Thus, following the *Bellman Principle of Optimality* [Bellman, 1957] “An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.” we have

$$V(s) = \max_{a^*} \left[R(s, a^*) - \bar{R} + \sum_{s'} V(s') T(s, a^*, s') \right]. \quad (2.4)$$

When evaluating Equation 2.4, we realize that $V(s)$ corresponds to the sum of the reward difference from the average reward \bar{R} encountered after taking the optimal action a^* in state s . Note that this function is usually discovered by human insight [Sutton and Barto, 1998]. This principle of optimality has given birth to the field of optimal control [Kirk, 1970] and the solution above corresponds to the dynamic programming solution from the viewpoint of reinforcement learning.

VALUE FUNCTION APPROACHES

Approach	Employed by...
<i>Model-Based</i>	Abbeel et al. [2006, 2007], Atkeson and Schaal [1997], Atkeson [1998], Bagnell and Schneider [2001], Bagnell [2004], Bakker et al. [2006], Coates et al. [2009], Donnart and Meyer [1996], Hester et al. [2010], Kalmár et al. [1998], Ko et al. [2007], Kolter et al. [2008], Martínez-Marín and Duckett [2005], Michels et al. [2005], Morimoto and Doya [2001], Ng et al. [2004b,a], Pendrith [1999], Schaal and Atkeson [1994], Schaal [1996], Touzet [1997], Willgoss and Iqbal [1999]
<i>Model-Free</i>	Asada et al. [1996], Bakker et al. [2003], Benbrahim et al. [1992], Benbrahim and Franklin [1997], Birdwell and Livingston [2007], Bitzer et al. [2010], Conn and Peters II [2007], Duan et al. [2007, 2008], Fagg et al. [1998], Gaskett et al. [2000], Gräve et al. [2010], Hafner and Riedmiller [2007], Huang and Weng [2002], Ilg et al. [1999], Katz et al. [2008], Kimura et al. [2001], Kirchner [1997], Kroemer et al. [2009, 2010], Latzke et al. [2007], Lizotte et al. [2007], Mahadevan and Connell [1992], Mataric [1997], Nemec et al. [2009, 2010], Oßwald et al. [2010], Paletta et al. [2007], Platt et al. [2006], Riedmiller et al. [2009], Rottmann et al. [2007], Smart and Kaelbling [1998, 2002], Soni and Singh [2006], Tamosiunaite et al. [2011], Thrun [1995], Tokic et al. [2009], Uchibe et al. [1998], Wang et al. [2006]

Table 2.1: This table illustrates different value function based reinforcement learning methods employed for robotic tasks and associated publications.

Hence, we have a dual formulation of the original problem as condition for optimality. Many traditional reinforcement learning approaches are based on this equation, these are called the value function methods. Instead of directly learning a policy, they first approximate the Lagrangian multiplier $V(s)$, also called the value function, and use it to reconstruct the optimal policy. A wide variety of methods exist and can be split mainly in three classes: (i) dynamic programming-based optimal control approaches such as policy iteration or value iteration, (ii) rollout-based Monte Carlo methods and (iii) temporal difference methods such as $TD(\lambda)$, Q -learning and SARSA. However, such value function based approaches often do not translate well into high dimensional robotics as proper representations for the value function become intractable and even finding the optimal action can already be a hard problem. A particularly drastic problem is the error propagation in value functions where a small change in the policy may cause a large change in the value function which again causes a large change in the policy. While this may lead faster to good, possibly globally optimal solutions, such a learning process is considerably more dangerous when applied on real systems where it is likely to cause significant damage. An overview of publications using value function based methods is presented in Table 2.1. Here, model-based methods refers to all methods that employ a predetermined or a learned model.

2.3.2 Policy Search

It is straightforward to realize that the primal formulation has a lot of features relevant to robotics. It allows a natural integration of expert knowledge, e.g., through initializations of the policy. It allows domain-appropriate pre-structuring of the policy in an approximate form without changing the original problem. Optimal policies often have many fewer parameters than optimal value functions, e.g., in linear quadratic control, the value function has quadratically many parameters while the policy requires only linearly many parameters. Extensions to continuous state and action spaces follow straightforwardly. Local search in policy space can directly lead to good results as exhibited by early hill-climbing approaches

POLICY SEARCH

Approach	Employed by...
<i>Gradient</i>	Deisenroth and Rasmussen [2010], Endo et al. [2008], Geng et al. [2006], Guenter et al. [2007], Gullapalli et al. [1994], Hailu and Sommer [1998], Kohl and Stone [2004], Kolter and Ng [2009], Mitsunaga et al. [2005], Miyamoto et al. [1996], Peters and Schaal [2008a,c], Tamei and Shibata [2009], Tedrake [2004], Tedrake et al. [2005]
<i>Heuristic</i>	Erden and Leblebicioğlu [2008], Dorigo and Colombetti [1993], Mataric [1994], Svinin et al. [2001], Yasuda and Ohkura [2008], Youssef [2005]
<i>Sample</i>	Buchli et al. [2011], Kalakrishnan et al. [2011], Kober and Peters [2008], Kober et al. [2010a], Pastor et al. [2011], Peters and Schaal [2008b], Peters et al. [2010a], Stulp et al. [2011], Tamosiunaite et al. [2011]

Table 2.2: This table illustrates different policy search reinforcement learning methods employed for robotic tasks and associated publications.

[Kirk, 1970]. Additional constraints can be incorporated naturally. As a result, policy search appears more natural to robotics.

Nevertheless, policy search has been considered the harder problem for a long time as the optimal solution cannot directly be determined from Equations (2.1-2.3) while the solution of the dual problem directly arises from the problems' Karush-Kuhn-Tucker conditions [Kuhn and Tucker, 1950] and the *Bellman Principle of Optimality* [Bellman, 1957].

Notwithstanding, in robotics, policy search has recently become an important alternative to value function based methods due to the reasons described above as well as the convergence problems of approximate value function methods. Most policy search methods optimize locally around existing policies π_i by computing policy changes $\delta\pi_i$ that will increase the expected return and results in iterative updates in the form

$$\pi_{i+1} = \pi_i + \delta\pi_i.$$

The computation of the policy update is the key step here and a variety of updates have been proposed ranging from pairwise comparisons [Strens and Moore, 2001, Ng et al., 2004a] over gradient estimation using finite policy differences [Geng et al., 2006, Mitsunaga et al., 2005, Sato et al., 2002, Tedrake et al., 2005], cross entropy [Rubinstein and Kroese, 2004], and heuristic parallel search methods (such as genetic algorithms, see [Goldberg, 1989]) to approaches coming from optimal control such as differential dynamic programming (DDP) [Atkeson, 1998] and multiple shooting approaches [Betts, 2001] as well as core reinforcement learning methods.

In recent years, general reinforcement learning has yielded three kinds of policy search approaches that have translated particularly well into the domain of robotics: (i) policy gradients approaches based on likelihood-ratio estimation [Sutton et al., 1999], (ii) policy updates inspired by expectation maximization [Toussaint et al., 2010], and (iii) the path integral methods [Kappen, 2005]. Likelihood-ratio policy gradient methods rely on perturbing the motor command instead of comparing in policy space. Initial approaches such as REINFORCE [Williams, 1992] have been rather slow but recent natural policy gradient approaches [Peters and Schaal, 2008a,c] have allowed faster convergence which may be useful for robotics. When the reward is treated as an improper probability distribution [Dayan and Hinton, 1997], safe and fast methods can be derived that are inspired by expectation-maximization. Some of these approaches have proven successful in robotics, e.g., reward-weighted regression [Peters and Schaal, 2008b], PoWER [Kober and Peters, 2008], MCEM [Vlassis et al., 2009], cost-regularized kernel regression [Kober et al., 2010a], and PI^2 [Theodorou et al., 2010]. An overview of publications using policy search methods is presented in Table 2.2.

2.4 Tractability through Representation

Much of the success of reinforcement learning methods has been due to the smart use of approximate representations. In a domain that is so inherently beyond the reach of complete tabular representation, the need of such approximations is particularly pronounced. The different ways of making reinforcement learning methods tractable in robotics are tightly coupled to the underlying framework. Policy search methods require a choice of policy representation that limits the number of representable policies to enhance learning speed, see Section 2.4.3. A value function-based approach requires an accurate, robust but general function approximator that can capture the value function sufficiently precisely, see Section 2.4.2. Reducing the dimensionality of states or actions by smart state-action discretization is a representational simplification that may enhance both policy search and value function-based methods, see 2.4.1. An overview of publications using representations to render the learning problem tractable is presented in Table 2.3.

2.4.1 Smart State-Action Discretization

Having lower-dimensional states or actions eases most reinforcement learning problems significantly, particularly in the context of robotics. Here, we give a quick overview of different attempts to achieve this goal with smart discretization.

Hand Crafted: A variety of authors have manually tuned discretizations so that basic tasks can be learned on real robots. For low-dimensional tasks, such as balancing a ball on a beam [Benbrahim et al., 1992], we can generate discretizations straightforwardly while much more

human experience is needed for more complex tasks. Such tasks range from basic navigation with noisy sensors [Willgoss and Iqbal, 1999] over one degree of freedom ball-in-a-cup [Nemec et al., 2010], two degree of freedom crawling motions [Tokic et al., 2009], and learning object affordances [Paletta et al., 2007] up to gait patterns for four legged walking [Kimura et al., 2001].

Learned from Data: Instead of specifying the discretizations by hand, they can also be learned from data. For example, a rule based reinforcement learning approach automatically segmented the state space to learn a cooperative task with mobile robots [Yasuda and Ohkura, 2008]. In the related field of computer vision, Piater et al. [2011] propose an approach that adaptively and incrementally discretizes a perceptual space into discrete states.

Meta-Actions: Automatic construction of meta actions has fascinated reinforcement learning researchers and there are various example in the literature. For example, in [Asada et al., 1996], the state and action sets are constructed in a way that repeated action primitives leads to a change in state to overcome problems associated with the discretization. Q-learning and dynamic programming based approaches have been compared in a pick-n-place task [Kalmár et al., 1998] using modules. A task of transporting a ball with a dog robot [Soni and Singh, 2006] can be learned with semi-automatically discovered options. Using only the sub-goals of primitive motions, a pouring task can be learned by a humanoid robot [Nemec et al., 2009]. Various other examples range from foraging [Mataric, 1997] and cooperative tasks [Mataric, 1994] with multiple robots, to grasping with restricted search spaces [Platt et al., 2006] navigation of a mobile robot [Dorigo and Colombetti, 1993]. These approaches belong to hierarchical reinforcement learning approaches.

Relational Representation: In a relational representation the states, actions, and transitions are not represented individually but entities of the same, predefined type are grouped and their relations are considered. This representation has been employed to learn to navigate buildings with a real robot in a supervised setting [Cocora et al., 2006] and to manipulate articulated objects in simulation [Katz et al., 2008].

SMART STATE-ACTION DISCRETIZATION

Approach	Employed by...
<i>Hand crafted</i>	Benbrahim et al. [1992], Kimura et al. [2001], Nemec et al. [2010], Paletta et al. [2007], Tokic et al. [2009], Willgoss and Iqbal [1999]
<i>Learned</i>	Piater et al. [2011], Yasuda and Ohkura [2008]
<i>Meta-actions</i>	Asada et al. [1996], Dorigo and Colombetti [1993], Kalmár et al. [1998], Mataric [1994, 1997], Platt et al. [2006], Soni and Singh [2006], Nemec et al. [2009]
<i>Relational Representation</i>	Cocora et al. [2006], Katz et al. [2008]

FUNCTION APPROXIMATION

Approach	Employed by...
<i>Local Models</i>	Bentivegna et al. [2004b], Schaal [1996], Smart and Kaelbling [1998]
<i>Neural Networks</i>	Benbrahim and Franklin [1997], Duan et al. [2008], Gaskett et al. [2000], Hafner and Riedmiller [2003], Riedmiller et al. [2009], Thrun [1995]
<i>GPR</i>	Gräve et al. [2010], Kroemer et al. [2009, 2010], Lizotte et al. [2007], Rottmann et al. [2007]
<i>Neighbors</i>	Hester et al. [2010], Mahadevan and Connell [1992], Touzet [1997]

PRE-STRUCTURED POLICIES

Approach	Employed by...
<i>Motor Primitives</i>	Kohl and Stone [2004], Kober and Peters [2008], Peters and Schaal [2008a,c], Stulp et al. [2011], Tamosiunaite et al. [2011], Theodorou et al. [2010]
<i>Neural Networks</i>	Endo et al. [2008], Geng et al. [2006], Gullapalli et al. [1994], Hailu and Sommer [1998]
<i>Via Points</i>	Miyamoto et al. [1996]
<i>Linear Models</i>	Tamei and Shibata [2009]
<i>GMM & LLM</i>	Deisenroth and Rasmussen [2010], Guenter et al. [2007], Peters and Schaal [2008b]
<i>Controller</i>	Kolter and Ng [2009], Tedrake [2004], Tedrake et al. [2005], Vlassis et al. [2009]
<i>Non-parametric</i>	Kober et al. [2010a], Mitsunaga et al. [2005], Peters et al. [2010a]

Table 2.3: This table illustrates different methods of making robot reinforcement learning tractable by employing a suitable representation.

2.4.2 Function Approximation

Function approximation has always been the key component that allowed value function methods to scale into interesting domains. In robot reinforcement learning, the following function approximation schemes have been popular and successful.

Neural networks: Neural networks as function approximators for continuous states and actions have been used by various groups, e.g., multi-layer perceptrons were used to learn a wandering behavior and visual servoing [Gaskett et al., 2000]; fuzzy neural networks [Duan et al., 2008] as well as explanation-based neural networks [Thrun, 1995] have allowed learning basic navigation while CMAC neural networks have been used for biped locomotion [Benbrahim and Franklin, 1997].

A particularly impressive application of value function approximation has been the success of Brainstormers RoboCup soccer team that used multi-layer perceptrons to learn various sub-tasks [Hafner and Riedmiller, 2003, Riedmiller et al., 2009]. The resulting components contributed to winning the world cup several times in different leagues. Retaining all collected data and using domain knowledge to only consider the task relevant state space, allows to efficiently fill up the state space.

Generalize to Neighboring Cells: As neural networks are globally affected from local errors, much work has focused on simply generalizing from neighboring cells. One of the earliest papers in robot reinforcement learning [Mahadevan and Connell, 1992] introduced this idea by statistical clustering states to speed up a box pushing task with a mobile robot, see Figure 2.1a. This approach was also used for a navigation and obstacle avoidance task with a mobile robot [Touzet, 1997]. Similarly, decision trees have been used to generalize states and actions to unseen ones, e.g., to learn a penalty kick on a humanoid robot [Hester et al., 2010].

Local Models: Locally weighted regression is known to be a particularly efficient function approximator. Using it for value function approximation has allowed learning a navigation task with obstacle avoidance [Smart and Kaelbling, 1998], a pole balancing task Schaal [1996] as well as an air hockey task [Bentivegna et al., 2004b].

Gaussian Process Regression: Using GPs as function approximator for the value function has allowed learning of hovering with an autonomous blimp [Rottmann et al., 2007], see Figure 2.1c. Similarly, another paper shows that grasping can be learned using Gaussian Process Regression [Gräve et al., 2010]. Grasping locations can be learned by focusing on rewards, modeled by GPR, by trying candidates with predicted high rewards [Kroemer et al., 2009]. High reward uncertainty allows intelligent exploration in reward-based grasping [Kroemer et al., 2010]. Gait of robot dogs can be optimized by first learning the expected return function with a Gaussian process regression and subsequently searching for the optimal solutions [Lizotte et al., 2007].

2.4.3 Pre-structured Policies

To make the policy search approach tractable, the policy needs to be represented with an appropriate function approximation.

Motor Primitives: Motor primitives are a biologically-inspired concept and represent basic movements. For discrete movements the dynamical system motor primitives [Ijspeert et al., 2002b, Schaal et al., 2007] representation has been employed to learn a T-ball batting task [Peters and Schaal, 2008a,c], the underactuated swing-up and ball-in-a-cup [Kober and Peters, 2008], flipping a light switch [Buchli et al., 2011], a pouring water [Tamosiunaite et al., 2011], as well as playing pool and manipulating a box [Pastor et al., 2011]. For rhythmic behaviors half-elliptical locuses have been used as a representation of the gait pattern of a robot dog [Kohl and Stone, 2004].

Neural Networks: Instead of analytically describing rhythmic movements, neural networks can be used as oscillators to learn gaits of a two legged robot [Geng et al., 2006, Endo et al., 2008]. Also a peg-in-hole (see Figure 2.1b) and a ball-balancing task [Gullapalli et al., 1994] as well as a navigation task [Hailu and Sommer, 1998] have been learned with neural networks as policy function approximators.




Figure 2.3: Boston Dynamics LittleDog jumping [Kolter and Ng, 2009] (Picture reprint with permission of Zico Kolter).

Via Points: Optimizing the position and timing of via-points Miyamoto et al. [1996] learned a kendama task.

Linear Models: Tamei and Shibata [2009] used reinforcement learning to to adjust a model that maps from EMG signals to forces that in turn is used in a cooperative holding task.

Gaussian Mixture Models and Locally Linear Models: One of the most general function approximators is based on radial basis functions, also called Gaussian kernels. However, specifying the centers and widths of these is challenging. These locations and variances can also be estimated from data prior to the reinforcement learning process which has been used to generalize a reaching movement [Guenter et al., 2007] and to learn the cart-pole task [Deisenroth and Rasmussen, 2010]. Operational space control was learned by Peters and Schaal [2008b] using locally linear models.

Controller: Here, parameters of a local linear controller are learned. Applications include learning to walk in 20 minutes with a biped robot [Tedrake, 2004, Tedrake et al., 2005], to drive a radio-controlled (RC) car as well as a jumping behavior for a robot dog jump [Kolter and Ng, 2009], as illustrated in Figure 2.3, and to balance a two wheeled robot [Vlassis et al., 2009].

Non-parametric: Also in this context non-parametric representations can be used. The weights of different robot human interaction possibilities [Mitsunaga et al., 2005], the weights of different striking movements in a table tennis task [Peters et al., 2010a], and the parameters of meta-actions for dart and table tennis tasks [Kober et al., 2010a] can be optimized.

2.5 Tractability through Prior Knowledge

Prior knowledge can significantly help to guide the learning process. Prior knowledge can be included in the form of initial policies, initial models, or a predefined structure of the task. These approaches significantly reduce the search space and, thus, speed up the learning process. Providing a goal achieving initial policy allows a reinforcement learning method to quickly explore promising regions in the value functions or in policy space, see Section 2.5.1. Pre-structuring the task to break a complicated task down into several more tractable ones can be very successful, see Section 2.5.2. An overview of publications using prior knowledge to render the learning problem tractable is presented in Table 2.4.

2.5.1 Prior Knowledge through Demonstrations

Animals and humans frequently learn using a combination of imitation and trial and error. For example, when learning to play tennis, an instructor usually shows the student how to do a proper swing, e.g., a forehand or backhand. The student will subsequently imitate this behavior but still needs hours of practicing to successfully return balls to the opponent's court. Input from a teacher is not limited to initial instruction. The instructor can give additional demonstrations in a later learning stage [Latzke et al., 2007], these can also be used as differential feedback [Argall et al., 2008]. Global exploration is not necessary as the student can improve by locally optimizing his striking movements previously obtained by imitation. A similar approach can speed up robot reinforcement learning based on human demonstrations or initial hand coded policies. Pieter Abbeel coined the term ‘apprenticeship learning’ for this type of learning.¹ For a recent survey on imitation learning for robotics see [Argall et al., 2009].

Both value function based approaches and policy search methods work well for real system applications if they do not make big changes of the distribution over states. Policy search approaches implicitly maintain the state distribution by limiting the changes to the policy. On the other hand, for value function based approaches an unstable estimate of the value function can lead to drastic changes in the policy if the state distribution is not maintained. The learning problem becomes significantly more tractable if it is known beforehand which states matter, that is the state distribution the approach needs to optimize over. Such state distributions can be obtained both from demonstrations by a teacher and hand coded policies.

Demonstrations by a Teacher: Demonstrations can be obtained by remote controlling the robot, which was used to initialize a Q-table for a navigation task [Conn and Peters II, 2007]. If the robot is back-drivable, kinesthetic teach-in (i.e., by taking it by the hand and moving it) can be employed. This method has resulted in applications including T-ball batting [Peters and Schaal, 2008a,c], reaching tasks [Guenther et al., 2007, Bitzer et al., 2010], ball-in-a-cup [Kober and Peters, 2008], flipping a light switch [Buchli et al., 2011], playing pool and manipulating a box [Pastor et al., 2011], as well as opening a door and picking up objects [Kalakrishnan et al., 2011]. A marble maze task can be learned using demonstrations by a human player [Bentivegna et al., 2004a]. Motion-capture setups can be used alternatively, but the demonstrations are often not as informative due to the correspondence problem. Demonstrations obtained by motion capture have been used to learn ball-in-a-cup [Kober et al., 2008] and grasping [Gräve et al., 2010].

Hand Coded Policy: A pre-programmed policy can provide demonstrations instead of a human teacher. A vision-based mobile robot docking task can be learned faster with such a basic behavior than using Q-learning alone as demonstrated in [Martínez-Marín and Duckett, 2005]. As an alternative corrective actions when the robot deviates significantly from the desired behavior can be employed as prior knowledge. This approach has been applied to adapt walking patterns of a robot dog to new surfaces [Birdwell and Livingston, 2007] by Q-learning. Having hand-coded stable initial gaits can significantly help as shown on six-legged robot gait [Erden and Leblebicioğlu, 2008] as well as on a biped [Tadrake, 2004, Tadrake et al., 2005].

2.5.2 Prior Knowledge through Task Structuring

Often a task can be decomposed hierarchically into basic components or in a sequence of increasingly difficult tasks. In both cases the complexity of the learning task is significantly reduced.

Hierarchical Reinforcement Learning: Easier tasks

can be used as building blocks for a more complex behavior. For example, hierarchical Q-learning has been used to learn different behavioral levels for a six legged robot: moving single legs, locally moving the complete body, and globally moving the robot towards a goal [Kirchner, 1997]. A stand-up

¹ The term ‘apprenticeship learning’ is often misconstrued to ‘inverse reinforcement learning’ or ‘inverse optimal control’ but actually should be employed in a much broader meaning.

DEMONSTRATION	
Approach	Employed by...
<i>Teacher</i>	Bentivegna et al. [2004a], Bitzer et al. [2010], Conn and Peters II [2007], Gräve et al. [2010], Kober et al. [2008], Kober and Peters [2008], Latzke et al. [2007], Peters and Schaal [2008a,c]
<i>Policy</i>	Birdwell and Livingston [2007], Erden and Leblebicioğlu [2008], Martínez-Marín and Duckett [2005], Smart and Kaelbling [1998], Tedrake [2004], Tedrake et al. [2005], Wang et al. [2006]
TASK STRUCTURE	
Approach	Employed by...
<i>Hierarchical</i>	Donnart and Meyer [1996], Kirchner [1997], Morimoto and Doya [2001]
<i>Progressive Tasks</i>	Asada et al. [1996]
DIRECTED EXPLORATION	
	Employed by...
	Huang and Weng [2002], Kroemer et al. [2010], Pendrith [1999]

Table 2.4: This table illustrates different methods of making robot reinforcement learning tractable by incorporating prior knowledge.

behavior considered as a hierarchical reinforcement learning task has been learned using Q-learning in the upper-level and TD-learning in the lower level [Morimoto and Doya, 2001]. Navigation in a maze can be learned using an actor-critic architecture by tuning the influence of different control modules and learning these modules [Donnart and Meyer, 1996].

Progressive Tasks: Often complicated tasks are easier to learn if simpler tasks can already be performed. A sequence of increasingly difficult missions has been employed to learn a goal shooting task in [Asada et al., 1996] using Q-learning.

2.5.3 Directing Exploration with Prior Knowledge

A mobile robot learns to direct attention [Huang and Weng, 2002] by employing a modified Q-learning approach using novelty. Using ‘corrected truncated returns’ and taking into account the estimator variance, a six legged robot employed with stepping reflexes can learn to walk [Pendrith, 1999]. Using upper confidence bounds to direct exploration grasping can be learned efficiently [Kroemer et al., 2010]. Offline search can be used to enhance Q-learning during a grasping task [Wang et al., 2006].

2.6 Tractability through Models

Using a simulation instead of the real physical robot has major advantages such as safety and speed. A simulation can be used to eliminate obviously bad behaviors and often runs much faster than real time. Simulations are without doubt a helpful testbed for debugging algorithms. A popular approach is to combine simulations and real evaluations by only testing promising policies on the real system and using it to collect new data to refine the simulation (Section 2.6.2). Unfortunately, directly transferring policies learned in simulation to a real system can be challenging (Section 2.6.3). An overview of publications using simulations to render the learning problem tractable is presented in Table 2.5.



Figure 2.4: Autonomous inverted helicopter flight [Ng et al., 2004b](Picture reprint with permission of Andrew Ng).

2.6.1 Role of Models

Model-free algorithms try to directly learn the value function or the policy. Model-based approaches jointly learn a model of the system and the value function or the policy. Model-based methods can make the learning process a lot more sample efficient. However, depending on the type of model these may require a lot of memory. Model-based approaches rely on an approach that finds good policies in the model. These methods encounter the risk of exploiting model inaccuracies to decrease the cost. If the learning methods require predicting the future or using derivatives, the inaccuracies may accumulate quickly, and, thus, significantly amplify noise and errors [An et al., 1988]. These effects lead to value functions or policies that work well in the model but poorly on the real system. This issue is highly related to the transfer problem discussed in Section 2.2.3. A solution is to overestimate the noise, to introduce a controlled amount of inconsistency [Atkeson, 1998], to use a crude model to find a policy that compensates the derivative of the behavior in the model and on the real system [Abbeel et al., 2006], or to solely rely on the gradient of the model.

In Section 2.2.3, we discussed that policies learned in simulation often cannot be transferred to the real system. However, simulations are still a very useful tool. Most simulations run significantly faster than real time and many problems associated with expensive samples (Section 2.2.2) can be avoided. For these reasons simulations are usually used to debug, test and optimize algorithms. Learning in simulation often can be made significantly easier than on real robots. The noise can be controlled and all variables can be accessed. If the approach does not work in simulation it is often unlikely that it works on the real system. Many papers also use simulations to benchmark approaches as repeating the experiment frequently to observe the average behavior and to compare many algorithms is often not feasible on the real system.

2.6.2 Mental Rehearsal

The idea of combining learning in simulation and in the real environment has been popularized by the Dyna-architecture [Sutton, 1990] in reinforcement learning. Due to the obvious advantages in the robotics domain, it has been proposed in this context as well. Experience collected in the real-world can be used to learn a forward model [Åström and Wittenmark, 1989] from data. Such a forward model allows training in a simulated environment and the resulting policy is subsequently transferred to the real environment. This approach can also be iterated and may significantly reduce the needed interactions with the real-world. However, often the learning process can exploit the model errors which may lead to biased solutions and slow convergence.

Benefits of Noise: A complex real-world system such as a robot can never be simulated without any kind of simulation errors. Even small inaccuracies accumulate and lead to simulation bias. Reinforcement learning approaches will exploit such inaccuracies if they are beneficial for the reward. This effect is closely related to overfitting in supervised learning. The problem can be addressed by introducing stochastic models or distributions over models even if the system is very close to deterministic. Artificially adding a little noise will smoothen model errors and therefore avoid overfitting [Jakobi et al., 1995, Atkeson, 1998]. On the downside, some very precise policies may be eliminated. This technique can be beneficial in all approaches described below.

Re-using Random Numbers: When comparing results of different simulation runs, it is often hard to tell from a small number of samples whether a policy really worked better or whether it was an effect of the simulated stochasticity. Using a large number of samples to obtain proper estimates of the expectations become prohibitively expensive if a large number of such comparisons need to be performed (e.g., for gradient estimation within an algorithm). A common technique in the statistics and simulation community [Glynn, 1987] to address this problem is to re-use the series of random numbers in fixed models, hence, removing the noise contribution. Ng et al. [2004b,a] extended this approach for learned simulators. The resulting approach, PEGASUS, found various applications in the learning of artistic maneuvers for autonomous helicopters [Bagnell and Schneider, 2001, Bagnell, 2004, Ng et al., 2004b,a], as illustrated in Figure 2.4, it has been used to learn control parameters for a RC car [Michels et al., 2005] and an autonomous blimp [Ko et al., 2007].

Taking Turns: Such mental rehearsal has found a lot of applications in robot reinforcement learning. Parallel learning in simulation and directed exploration allows Q-learning to learn a navigation task from scratch in 20 minutes [Bakker et al., 2006]. Two robots taking turns in learning a simplified soccer task were also able to profit from mental rehearsal [Uchibe et al., 1998]. Atkeson et al. [1997] learned a billiard and a devil sticking task employing forward models. Nemec et al. [2010] used a value function learned in simulation to initialize the real robot learning.

Iterative Learning Control: A powerful idea is to use crude approximate models to determine gradients, e.g., for an update step. The resulting new policy is then evaluated in the real world and the model is updated. This approach is known as iterative learning control [Arimoto et al., 1984] and is related to feedback error learning [Kawato, 1990]. A similar preceding idea was employed to minimize trajectory tracking errors [An et al., 1988]. More recently, iterative learning control has been employed to learn robot control [Norrlöf, 2002, Bukkems et al., 2005], the steering of a RC car [Abbeel et al., 2006], a pick and place task [Freeman et al., 2010], as well as tying knots with a surgical robot at superhuman speeds [Berg et al., 2010].

Linear Quadratic Regulators: Instead of sampling from a forward model-based simulator, such learned models can also be directly used for computing optimal control policies. This has resulted in a variety of robot reinforcement learning applications ranging from pendulum swing-up tasks learned with DDP [Atkeson and Schaal, 1997, Atkeson, 1998], devil-sticking (a form of gyroscopic juggling) obtained with local LQR solutions [Schaal and Atkeson, 1994], trajectory following with space-indexed controllers trained with DDP for an autonomous RC car [Kolter et al., 2008], the cart-pole task [Deisenroth and Rasmussen, 2010], a block stacking task [Deisenroth et al., 2011], to aerobatic helicopter flight trained with DDP [Coates et al., 2009]. Solving an LQR problem with multiple probabilistic models and combining the resulting closed-loop control with open-loop control has resulted in autonomous sideways sliding into a parking spot [Kolter et al., 2010]. A promising new related approach are LQR-trees [Tedrake et al., 2010].

2.6.3 Direct Transfer from Simulated to Real Robots

Only few papers claim that a policy learned in simulation can directly be transferred to a real robot while maintaining its high level of performance. The few examples include maze navigation tasks [Bakker et al., 2003, Oßwald et al., 2010, Youssef, 2005] and obstacle avoidance [Fagg et al., 1998] for a mobile

Approach	Employed by...
<i>Re-using Random Numbers</i>	Bagnell and Schneider [2001], Bagnell [2004], Ko et al. [2007], Michels et al. [2005], Ng et al. [2004b,a]
<i>Taking Turns</i>	Atkeson et al. [1997], Bakker et al. [2006], Nemec et al. [2010], Uchibe et al. [1998]
<i>Iterative Learning Control</i>	Abbeel et al. [2006], An et al. [1988], Berg et al. [2010], Bukkems et al. [2005], Freeman et al. [2010], Norrlöf [2002]
<i>Linear Quadratic Regulators</i>	Atkeson and Schaal [1997], Atkeson [1998], Coates et al. [2009], Deisenroth and Rasmussen [2010], Deisenroth et al. [2011], Kolter et al. [2008, 2010], Schaal and Atkeson [1994], Tedrake et al. [2010]

DIRECT POLICY TRANSFER

	Employed by...
	Bakker et al. [2003], Duan et al. [2007], Fagg et al. [1998], Ilg et al. [1999], Oßwald et al. [2010], Svinin et al. [2001], Youssef [2005]

Table 2.5: This table illustrates different methods of making robot reinforcement learning tractable using models.

robot. Similar transfer was achieved in very basic robot soccer [Duan et al., 2007] and multi-legged robot locomotion [Ilg et al., 1999, Svinin et al., 2001].

2.7 A Case Study: Ball-in-a-Cup

Up to this point in this chapter, we have reviewed a large variety of problems and associated possible solutions of robot reinforcement learning. In this section, we will take an orthogonal approach and discuss one task in detail that has previously been studied. This task is called ball-in-a-cup and due to its complexity, it can serve as an example to highlight some of the various discussed challenges and methods. In Section 2.7.1, the experimental setting is described with a focus on the task and the reward. Section 2.7.2 discusses a type of pre-structured policies that has been particularly useful in robotics. Inclusion of prior knowledge is presented in Section 2.7.3. The advantages of the employed policy search algorithm are explained in Section 2.7.4. The use of simulations in this task is discussed in Section 2.7.5. Finally, an alternative reinforcement learning approach is explored in Section 2.7.6.

2.7.1 Experimental Setting: Task and Reward

The children’s motor game ball-in-a-cup, also known as balero and bilboquet, is challenging even for most adults. The toy consists of a small cup held in one hand (or, in this case, is attached to the end-effector of the robot) and a small ball is hanging on a string attached to the cup’s bottom (for the employed toy, the string is 40cm long). Initially, the ball is at rest, hanging down vertically. The player needs to move fast to induce motion in the ball through the string, toss it up and catch it with the cup. A possible movement is illustrated in Figure 2.5a. As the string frequently entangles in failures and the robot cannot unravel it, human intervention is required and an automatic reset is often not possible.

The state of the system can be described by joint angles and joint velocities of the robot as well as the the Cartesian coordinates and velocities of the ball. The actions are the joint space accelerations (which are translated into torques by an inverse dynamics controller). Thus, the reinforcement learning approach has to deal with twenty state and seven action dimensions. Discretizing the state-action space for a value function based approach is not likely to be feasible.

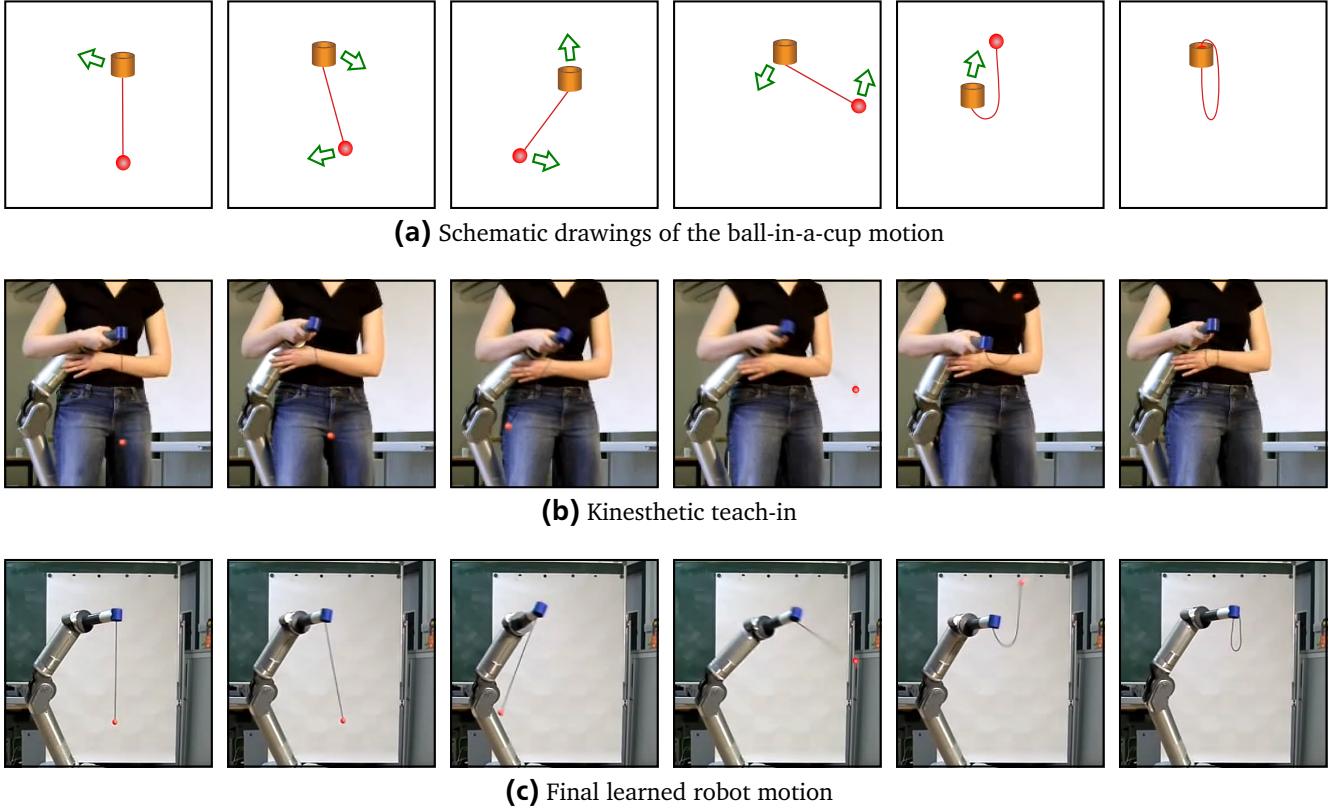


Figure 2.5: This figure shows schematic drawings of the ball-in-a-cup motion (a), the final learned robot motion (c), as well as a kinesthetic teach-in (b). The green arrows show the directions of the current movements in that frame. The human cup motion was taught to the robot by imitation learning with 31 parameters per joint for an approximately 3 seconds long movement. The robot manages to reproduce the imitated motion quite accurately, but the ball misses the cup by several centimeters. After approximately 75 iterations of the Policy learning by Weighting Exploration with the Returns (PoWER) algorithm the robot has improved its motion so that the ball regularly goes into the cup.

At the time t_c where the ball passes the rim of the cup with a downward direction, the reward is computed as $r(t_c) = \exp(-\alpha(x_c - x_b)^2 - \alpha(y_c - y_b)^2)$ while $r(t) = 0$ for all $t \neq t_c$. Here, the cup position is denoted by $[x_c, y_c, z_c] \in \mathbb{R}^3$, the ball position $[x_b, y_b, z_b] \in \mathbb{R}^3$ and the scaling parameter $\alpha = 100$. Initially, Kober and Peters [2010] used a reward function based solely on the minimal distance between the ball and the cup. However, the algorithm has exploited rewards resulting from hitting the cup with the ball from below or from the side as such behaviors are easier to achieve and yield comparatively high rewards. To avoid such local optima, it was essential to find a good reward function such as the initially described one.

The task is quite complex as the reward is not only affected by the cup's movements but foremost by the ball's movements. As the ball's movements are very sensitive to small perturbations, the initial conditions or small arm movement changes, will drastically affect the outcome. Creating an accurate simulation is hard due to the nonlinear, unobservable dynamics of the string and its non-negligible weight.

2.7.2 Appropriate Policy Representation

The policy is represented by dynamical system motor primitives [Ijspeert et al., 2002b, Schaal et al., 2007]. The global movement is encoded as a point attractor linear dynamical system. The details of the movement are generated by a transformation function that allows learning complex behaviors. This

transformation function is modeled using locally linear function approximation. This combination of the global attractor behavior and local transformation allows a very parsimonious representation of the policy. This policy is linear in parameters $a = \theta\mu(s)$ and, thus, it is straightforward to include prior knowledge from a demonstration using supervised learning by locally weighted regression.

2.7.3 Generating a Teacher's Demonstration

Due to the complexity of the task, ball-in-a-cup is a hard motor task even for children who usually only succeed after observing another person presenting a demonstration, and require additional trial-and-error-based learning. Mimicking how children learn ball-in-a-cup, the motor primitives are first initialized by imitation and, subsequently, improved by reinforcement learning.

A demonstration for imitation was obtained by recording the motions of a human player performing kinesthetic teach-in as shown in Figure 2.5b. Kinesthetic teach-in means ‘taking the robot by the hand’, performing the task by moving the robot while it is in gravity-compensation mode and recording the joint angles, velocities and accelerations. It requires a back-drivable robot system similar to a human arm. From the imitation, the number of needed policy parameters can be determined by cross-validation. As the robot fails to catch the ball with the cup, additional reinforcement learning is needed for self-improvement.

2.7.4 Reinforcement Learning by Policy Search

When learning motor primitives, a deterministic mean policy $\bar{a} = \theta\mu(s) = f(z)$ is learned. This policy is linear in parameters θ and augmented by additive exploration $\epsilon(s, t)$ to make model-free reinforcement learning possible. As a result, the explorative policy can be given in the form $a = \theta\mu(s, t) + \epsilon(\mu(s, t))$. Policy search approaches often focus on state-independent, white Gaussian exploration, i.e., $\epsilon(\mu(s, t)) \sim \mathcal{N}(0, \Sigma)$ which has resulted into applications such as T-Ball batting [Peters and Schaal, 2008a] and constrained movement [Guenther et al., 2007]. However, such unstructured exploration at every step has several disadvantages, i.e., (i) it causes a large variance which grows with the number of time-steps, (ii) it perturbs actions too frequently, thus, ‘washing’ out their effects and (iii) can damage the system executing the trajectory.

Alternatively, one could generate a form of structured, state-dependent exploration [Rückstieß et al., 2008] $\epsilon(\mu(s, t)) = \varepsilon_t\mu(s, t)$ with $[\varepsilon_t]_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$, where σ_{ij}^2 are meta-parameters of the exploration that can also be optimized. This argument results into the policy $a \sim \pi(a_t|s_t, t) = \mathcal{N}(a|\mu(s, t), \hat{\Sigma}(s, t))$.

Kober and Peters [2008] have derived a framework of reward weighted imitation. Based on [Dayan and Hinton, 1997] they consider the return of an episode as an improper probability distribution. A lower bound of the logarithm of the expected return is maximized. Depending on the strategy of optimizing this lower bound and the exploration strategy, the framework yields several well known policy search algorithms: episodic REINFORCE [Williams, 1992], the policy gradient theorem [Sutton et al., 1999], episodic natural actor critic [Peters and Schaal, 2008c], a generalization of the reward-weighted regression [Peters and Schaal, 2008b] as well as the novel Policy learning by Weighting Exploration with the Returns (PoWER) algorithm. PoWER is an expectation-maximization inspired algorithm that employs state-dependent exploration. The update rule is given by

$$\theta' = \theta + \frac{E_\tau \left\{ \sum_{t=1}^T \varepsilon_t Q^\pi(s_t, a_t, t) \right\}}{E_\tau \left\{ \sum_{t=1}^T Q^\pi(s_t, a_t, t) \right\}}.$$

To reduce the number of trials in this on-policy scenario, the trials are reused through importance sampling [Sutton and Barto, 1998]. To avoid the fragility sometimes resulting from importance sampling in reinforcement learning, samples with very small importance weights are discarded. This algorithm performs basically a local search around the policy learned from prior knowledge.

The policy converges to the maximum after 100 episodes and the robot regularly succeeds at bringing the ball into the cup after approximately 75 episodes.

Using a value function based approach would require an unrealistic amount of samples to get a good estimate of the value function. Greedily searching for an optimal motor command in such a high-dimensional action space is probably as hard as finding a locally optimal policy.

2.7.5 Use of Simulations in Robot Reinforcement Learning

The robot is simulated by rigid body dynamics with parameters estimated from data. The toy is simulated as a pendulum with an elastic string that switches to a ballistic point mass when the ball is closer to the cup than the string is long. The spring, damper and restitution constants were tuned to match recorded data. Even though this simulation matches recorded data very well, policies that get the ball in the cup in simulation usually miss the cup by several centimeters on the real system and vice-versa. However, this simulation was very helpful to develop and tune the algorithm as it runs faster in simulation than real-time and does not require human supervision or intervention.

2.7.6 Alternative Approach with Value Function Methods

Nemec et al. [2010] used a different reinforcement learning approach to achieve the ball-in-a-cup task with a Mitsubishi PA10 robot. They decomposed the task in two sub-tasks, the swing-up phase and the catching phase. In the swing-up phase the ball is moved above the cup. In the catching phase the ball is caught with the cup using an analytic prediction of the ball trajectory based on the movement of a flying point mass. The catching behavior is fixed, only the swing-up behavior is learned. The paper proposes to use SARSA to learn the swing-up movement. The states consist of the cup positions and velocities as well as the angular positions and velocities of the ball. The actions are the accelerations of the cup in a single Cartesian direction. Tractability is achieved by discretizing both the states (324 values) and the actions (5 values) and initialization by simulation. The behavior was first learned in simulation requiring 220 to 300 episodes. The state-action value function learned in simulation was used to initialize the learning on the real robot. The robot required an additional 40 to 90 episodes to adapt the behavior learned in simulation to the real environment.

2.8 Discussion

Incorporating prior knowledge is one of the main tools to make robotic reinforcement learning tractable. However, it is often hard to tell in advance how much prior knowledge is required to make the reinforcement learning algorithms succeed in a reasonable number of episodes. If local optimization is employed, an initial policy needs to be defined. Obtaining these from human demonstrations can be virtually impossible if the morphology of the robot is different from a human's and more than a small number of degrees of freedom need to be controlled. Programming a policy by hand is usually too time consuming to be a desirable alternative.

The action space in robotics often is inherently continuous and multi-dimensional. Care needs to be taken when using function approximation. Model based approaches can be used to reduce required real-world interactions. Often methods based on approximate models and using local optimization work well. On the one hand, due to under-modeling and resulting model errors, policies learned only in simulation usually cannot be transferred directly to the robot. On the other hand, real-world samples are expensive in most cases while calculations are comparatively cheap.

The state of the robot can only be obtained indirectly by sensors or is not observable at all. Often a lot of effort needs to be put into pre-processing the sensor information to obtain the states for the reinforcement learning algorithm. Often the environment and the robot itself change dynamically, for

example, vision systems depend on the lightning condition and the robot dynamics change with wear and the temperature of the grease. Thus, constant adaptation might be a requirement.

Reinforcement learning is not yet applicable to robotics out of the box. For each task, appropriate methods need to be carefully selected. The user has to decide when enough prior knowledge is given and learning can take over. All methods require hand-tuning as appropriate representations, reward functions, and prior knowledge need to be chosen. Inverse reinforcement learning could be used as an alternative to designing the reward function.

In the field of inverse optimal control, robotic problems clearly have driven theoretical research. Policy search and robustness concerns also has been pushed by robotic research. Nevertheless, more theoretical research is required to come up with algorithms that can handle the specific challenges of robotic reinforcement learning. These methods need to address multi-dimensional continuous action spaces, partial observability, model learning, and continuous adaptation to changes in the hardware and environment.

Robotic reinforcement learning has a number of advantages over algorithmic reinforcement learning. Prior knowledge of the behavior of the system as well as the task is often available and can be incorporated in the learning approach. The use of models permits to limit the number of real-world trials as the learning can be partially done based on the model. Using crude models to approximate gradients has yielded impressive results.

Policy search approaches implicitly maintain a state distribution. Also value function based methods benefit from such fixed state distributions. For some tasks, value function based approaches can be made tractable in a straightforward manner by using macro actions. Often the robotic learning approaches are closer to optimal control problems than the ones typically studied in algorithmic reinforcement learning. All approaches benefit from reward shaping, that is by using rewards that convey a notation of closeness and are not only based on success or failure. Value function heuristics can be successfully employed.

2.9 Next Challenges

When a human is learning a new task he relies on the skills he has learned previously. For example, learning to throw darts is significantly easier when the skill of throwing balls has been learned previously compared to learning from scratch. Also for reinforcement learning applied to robotics it will be crucial to have the possibility to transfer learned skills to other tasks and potentially to robots of a different type. For complex tasks learning cannot be achieved globally, hence it is essential to reuse other locally learned information.

One of the key differences to synthetic reinforcement learning benchmarks is that robotic reinforcement learning has to handle noisy and incomplete perceptions. Most of the surveyed work relies on a very simple perception system or pre-processed data. In contrast humans learn to optimally perceive in order to fulfill a task. For example, when catching a ball, people learn to cancel the optic acceleration in order to facilitate to intercept its path [McLeod et al., 2006]. Integrating computer vision and robotic reinforcement learning will allow many applications beyond the reach of current techniques. Learning active perception jointly with the robot's movement and semantic perception are open problems.

2.10 Conclusion

In this chapter, we have surveyed robot reinforcement learning in order to introduce general reinforcement learning audiences to the state of the art in this domain. We have pointed out the inherent challenges such as the high-dimensional continuous state and action space, the high cost associated with trials, the problems associated with transferring policies learned in simulation to real robots as well as the need for appropriate reward functions. A discussion of how different robot reinforcement learning approaches are affected by the domain has been given. We have surveyed different authors' approaches to render

robot reinforcement learning tractable through improved representation, inclusion of prior knowledge and usage of simulation. To highlight aspects that we found particularly important, we give a case study on how a robot can learn a complex task such as ball-in-a-cup.



3 Movement Templates for Learning of Hitting and Batting

Hitting and batting tasks, such as tennis forehands, ping-pong strokes, or baseball batting, depend on predictions where the ball can be intercepted and how it can properly be returned to the opponent. These predictions get more accurate over time, hence the behaviors need to be continuously modified. As a result, movement templates with a learned global shape need to be adapted during the execution so that the racket reaches a target position and velocity that will return the ball over to the other side of the net or court. It requires altering learned movements to hit a varying target with the necessary velocity at a specific instant in time. Such a task cannot be incorporated straightforwardly in most movement representations suitable for learning. For example, the standard formulation of the dynamical system based motor primitives (introduced by Ijspeert et al. [2002a]) does not satisfy this property despite their flexibility which has allowed learning tasks ranging from locomotion to kendama. In order to fulfill this requirement, we reformulate the Ijspeert framework to incorporate the possibility of specifying a desired hitting point and a desired hitting velocity while maintaining all advantages of the original formulation. We show that the proposed movement template formulation works well in two scenarios, i.e., for hitting a ball on a string with a table tennis racket at a specified velocity and for returning balls launched by a ball gun successfully over the net using forehand movements. All experiments were carried out on a Barrett WAM using a four camera vision system.

3.1 Introduction

Learning new skills can frequently be helped significantly by choosing a movement template representation that facilitates the process of acquiring and refining the desired behavior. For example, the work on dynamical systems-based motor primitives [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] has allowed speeding up both imitation and reinforcement learning while, at the same time, making them more reliable. Resulting successes have shown that it is possible to rapidly learn motor primitives for complex behaviors such as tennis swings [Ijspeert et al., 2002a] with only a final target, constrained reaching [Gams and Ude, 2009], drumming [Pongas et al., 2005], biped locomotion [Schaal et al., 2003, Nakanishi et al., 2004] and even in tasks with potential industrial application [Urbanek et al., 2004]. Although some of the presented examples, e.g., the tennis swing [Ijspeert et al., 2002a] or the T-ball batting [Peters and Schaal, 2006], are striking movements, these standard motor primitives cannot properly encode a hitting movement. Previous work needed to make simplistic assumptions such as having a static goal [Peters and Schaal, 2006], a learned complex goal function [Peters et al., 2010b] or a stationary goal that could only be lightly touched at the movement's end [Ijspeert et al., 2002a].

Most racket sports require that we hit a non-stationary target at various positions and with various velocities during the execution of a complete striking movement. For example, in table tennis, a typical movement consists of swinging back from a rest postures, hitting the ball at a desired position with a desired orientation and velocity, continuing the swing a bit further and finally returning to the rest posture. See Figure 3.1 for an illustration. Sports sciences literature [Ramanantsoa and Durey, 1994, Mülling and Peters, 2009] indicates that most striking movements are composed of similar phases that only appear to be modified by location and velocity at the interception point for the ball [Schmidt and Wrisberg, 2000, Bootsma and van Wieringen, 1990, Hubbard and Seng, 1954, Tyldesley and Whiting, 1975]. These findings indicate that similar motor primitives are being used that are invariant under these external influences similar to the Ijspeert motor primitives [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] being invariant under the modification of the final position, movement amplitude and duration. However, the standard formulation by Ijspeert et al. cannot be used properly in this context as there is no possibility to directly incorporate either a via-point or a target velocity (if the duration cannot be adapted as, e.g., for an approaching ball). Hence, a reformulation is needed that can deal with these requirements.




Figure 3.1: This figure illustrates the different phases of a table tennis stroke. The blue box on the left represents a ping-pong ball launcher, the table is shown in green and the different states of the robot are superposed. A typical racket trajectory is indicated by the dark gray solid line while the orange dashed line represents the ball's trajectory. The robot goes through four stages: it swings back (①→②), it strikes the ball at a virtual hitting point with a goal-oriented velocity at posture ③, it follows the strike through (③→④) and finally returns to the rest posture ①. While the complete arm movement can be modeled with the Ijspeert approach, the reformulation in this paper is required to be able to properly strike the ball.

In this paper, we augment the Ijspeert approach [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] of using dynamical systems as motor primitives in such a way that it includes the possibility to set arbitrary velocities at the hitting point without changing the overall shape of the motion or introducing delays that will prevent a proper strike. This modification allows the generalization of learned striking movements, such as hitting and batting, from demonstrated examples. In Section 3.2, we present the reformulation of the motor primitives as well as the intuition behind the adapted approach. We apply the presented method in Section 3.3 where we show two successful examples. First, we test the striking movements primitive in a static scenario of hitting a hanging ball with a table tennis racket and show that the movement generalizes well to new ball locations. After this proof of concept, we take the same movement template representation in order to learn an adaptable forehand for table tennis. The later is tested in the setup indicated by Figure 3.1 where we use a seven degrees of freedom Barrett WAM in order to return balls launched by a ball cannon.

3.2 Movement Templates for Learning to Strike

Ijspeert et al. [2002a] suggested to use a dynamical system approach in order to represent both discrete point-to-point movements as well as rhythmic motion with motor primitives. This framework ensures the stability of the movement¹, allows the representation of arbitrarily shaped movements through the primitive's policy parameters, and these parameters can be estimated straightforwardly by locally weighted regression. In the discrete case, these primitives can be modified through their meta-parameters in order to adapt to the final goal position, the movement amplitude or the duration of the movement. The resulting movement can start from arbitrary positions and velocities and go to arbitrary final positions while maintaining the overall shape of the trajectory. In Section 3.2.1, we review the most current version of this approach based on [Schaal et al., 2007].

However, as outlined in Section 3.1, this formulation of the motor primitives cannot be used straightforwardly in racket sports as incorporating a desired virtual hitting point [Ramanantsoa and Durey, 1994, Mülling and Peters, 2009] (consisting of a desired target position and velocity) cannot be achieved

¹ Note that the dynamical systems motor primitives ensure the stability of the movement generation but cannot guarantee the stability of the movement execution [Ijspeert et al., 2002a, Schaal et al., 2007].

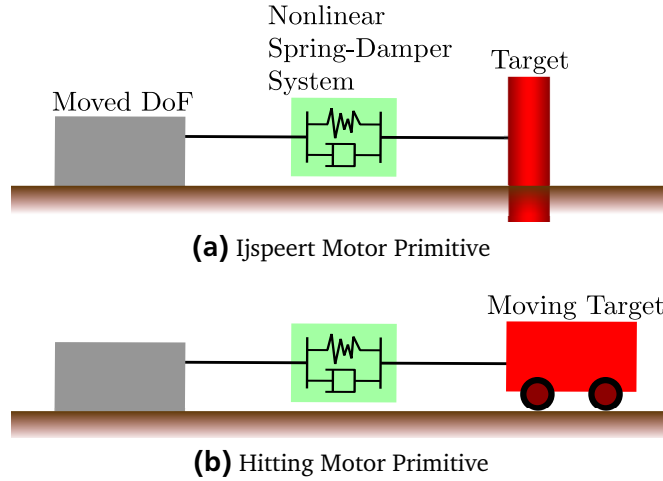


Figure 3.2: In this figure, we convey the intuition of the presented reactive templates for learning striking movements. The Ijspeert formulation can be seen as a nonlinear spring damper system that pulls a degree of freedom to a stationary goal while exhibiting a learned movement shape. The presented approach allows hitting a target with a specified velocity without replanning if the target is adapted and, in contrast to the Ijspeert formulation, can be seen as a degree of freedom pulled towards a moving goal.

straightforwardly. For example, in table tennis, fast forehand movements need to hit a ball at a pre-specified speed, hitting time and a continuously adapted location. In the original Ijspeert formulation, the goal needs to be determined at the start of the movement and at approximately zero velocity as in the experiments in [Ijspeert et al., 2002a]; a via-point target can only be hit properly by modifying either the policy parameters [Peters and Schaal, 2006] or, indirectly, by modifying the goal parameters [Peters et al., 2010b]. Hence, such changes of the target can only be achieved by drastically changing the shape of the trajectory and duration.

As an alternative, we propose a modified version of Ijspeert’s original framework that overcomes this limitation and is particularly well-suited for striking movements. This modification allows setting both a hitting point and a striking velocity while maintaining the desired duration and the learned shape of the movement. Online adaptation of these meta-parameters is possible and, hence, it is well-suited for learning racket sports such as table tennis as discussed in Chapter 5. The basic intuition behind the modified version is similar to the one of Ijspeert’s primitives, i.e., both assume that the controlled degree of freedom is connected to a specific spring damper system; however, the approach presented here allows overcoming previous limitations by assuming a connected moving target, see Figure 3.2. The resulting approach is explained in Section 3.2.2.

A further drawback of the Ijspeert motor primitives [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] is that, when generalizing to new targets, they tend to produce large accelerations early in the movement. Such an acceleration peak may not be well-suited for fast movements and can lead to execution problems due to physical limitations; Figure 3.4 illustrates this drawback. In Section 3.2.3, we propose a modification that alleviates this shortcoming.

3.2.1 Discrete Movement Primitives

While the original formulation in [Ijspeert et al., 2002a, Schaal et al., 2003] for discrete dynamical systems motor primitives used a second-order system to represent the phase z of the movement, this formulation has proven to be unnecessarily complicated in practice. Since then, it has been simplified and, in [Schaal et al., 2007], it was shown that a single first order system suffices

$$\dot{z} = -\tau \alpha_z z. \quad (3.1)$$

This canonical system has the time constant $\tau = 1/T$ where T is the duration of the motor primitive, a parameter α_z which is chosen such that $z \approx 0$ at T to ensure that the influence of the transformation function, shown in Equation (3.3), vanishes. Subsequently, the internal state \mathbf{x} of a second system is chosen such that positions \mathbf{q} of all degrees of freedom are given by $\mathbf{q} = \mathbf{x}_1$, the velocities $\dot{\mathbf{q}}$ by $\dot{\mathbf{q}} = \tau \mathbf{x}_2 = \dot{\mathbf{x}}_1$ and the accelerations $\ddot{\mathbf{q}}$ by $\ddot{\mathbf{q}} = \tau \dot{\mathbf{x}}_2$. Under these assumptions, the learned dynamics of Ijspeert motor primitives can be expressed in the following form

$$\begin{aligned}\dot{\mathbf{x}}_2 &= \tau \alpha_x (\beta_x (\mathbf{g} - \mathbf{x}_1) - \mathbf{x}_2) + \tau \mathbf{A} \mathbf{f}(z), \\ \dot{\mathbf{x}}_1 &= \tau \mathbf{x}_2.\end{aligned}\tag{3.2}$$

This set of differential equations has the same time constant τ as the canonical system, parameters α_x , β_x set such that the system is critically damped, a goal parameter \mathbf{g} , a transformation function \mathbf{f} and an amplitude matrix $\mathbf{A} = \text{diag}(a_1, a_2, \dots, a_n)$, with the amplitude modifier $\mathbf{a} = [a_1, a_2, \dots, a_n]$. In [Schaal et al., 2007], they use $\mathbf{a} = \mathbf{g} - \mathbf{x}_1^0$ with the initial position \mathbf{x}_1^0 , which ensures linear scaling. Alternative choices are possibly better suited for specific tasks, see e.g., [Park et al., 2008]. The transformation function $\mathbf{f}(z)$ alters the output of the first system, in Equation (3.1), so that the second system, in Equation (3.2), can represent complex nonlinear patterns and it is given by

$$\mathbf{f}(z) = \sum_{i=1}^N \psi_i(z) \mathbf{w}_i z.\tag{3.3}$$

Here, \mathbf{w}_i contains the i^{th} adjustable parameter of all degrees of freedom, N is the number of parameters per degree of freedom, and $\psi_i(z)$ are the corresponding weighting functions [Schaal et al., 2007]. Normalized Gaussian kernels are used as weighting functions given by

$$\psi_i(z) = \frac{\exp(-h_i(z - c_i)^2)}{\sum_{j=1}^N \exp(-h_j(z - c_j)^2)}.\tag{3.4}$$

These weighting functions localize the interaction in phase space using the centers c_i and widths h_i . Note that the degrees of freedom (DoF) are usually all modeled as independent in Equation (3.2). All DoFs are synchronous as the dynamical systems for all DoFs start at the same time, have the same duration, and the shape of the movement is generated using the transformation $\mathbf{f}(z)$ in Equation (3.3). This transformation function is learned as a function of the shared canonical system in Equation (3.1).

One of the biggest advantages of this motor primitive framework [Ijspeert et al., 2002b, Schaal et al., 2007] is that the second system in Equation (3.2), is linear in the shape parameters $\boldsymbol{\theta}$. Therefore, these parameters can be obtained efficiently, and the resulting framework is well-suited for imitation (Section 3.2.4) and reinforcement learning (Chapter 4). Additional feedback terms can be added as shown in [Schaal et al., 2007, Kober et al., 2008, Park et al., 2008].

3.2.2 Adapting the Motor Primitives for Striking Movements

The regular formulation [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] which was reviewed in Section 3.2.1, allows to change the initial position \mathbf{x}_1^0 and goal position \mathbf{g} (which corresponds to the target, i.e., the position at the end of the movement at time T) of the motor primitive while maintaining the overall shape of the movement determined by the parameters \mathbf{w}_i . For disturbed initial conditions, the attractor dynamics that pull the motor primitive to the trained behavior and it is guaranteed to finally reach to the goal position \mathbf{g} , see [Ijspeert et al., 2002a]. However, the formulation above only considers the case of a final goal with a favored velocity of $\dot{\mathbf{x}}_1(T) = 0$ at the goal \mathbf{g} and final time T . However, using the transformation function $\mathbf{f}(z)$ in Equation (3.3), it can be forced to arbitrary final velocities by changing the




Figure 3.3: Target velocity adaptation is essential for striking movements. This figure illustrates how different versions of the dynamical system based motor primitives are affected by a change of the target velocity. Here, an artificial training example (i.e., $q = 2t^2 + \cos(4t\pi) - 1$) is generated. After learning, all motor primitive formulations manage to reproduce the movements accurately from the training example for the same target velocity and cannot be distinguished. When the target velocity is tripled, this picture changes drastically. For Ijspeert's original model the amplitude modifier \mathbf{a} had to be increased to yield the desired velocity. The increased amplitude of the trajectory is clearly visible for the positions and even more drastic for the velocities and accelerations. The reformulations presented in this paper, stay closer to the movement shape and amplitude. Particularly the velocities and accelerations exhibit that the new approach allows much better generalizing of the learned behavior. This figure furthermore demonstrates how a large initial step in acceleration appears for Ijspeert's original model (and the reformulation for hitting) even if a transformation function is used to partially suppress it for the training example.

shape parameters of the movement. As the last basis function in the transformation function $\mathbf{f}(\mathbf{z})$ decays almost to zero at time T the active parameters, the last parameter \mathbf{w}_N needs to be over-proportionally large. If the motor primitive is trained with $\dot{\mathbf{x}}_1(T) = 0$ it simply rests at $\mathbf{x}_1 = \mathbf{g}$ if it runs for longer than T . However, large \mathbf{w}_N often cause overshooting in \mathbf{x}_1 and the trajectory is subsequently pulled back to the final position \mathbf{g} only using the linear attractor dynamics in Equation (3.2) which may not be suitable for a given task. The goal velocity $\dot{\mathbf{x}}_1(T)$ can only be changed either by scaling the duration of the movement T or with the amplitude modifier \mathbf{a} ; however a mapping of \mathbf{g} and $\dot{\mathbf{x}}_1(T)$ to \mathbf{a} has to be established first. The main downsides of these approaches, respectively, are that either the total duration is changed (which makes the interception of a table tennis ball hard) or that \mathbf{a} modifies the whole motion including shape and amplitude (which causes undesired movements and often requires overly strenuous movements in table tennis). These effects are illustrated in Figure 3.3. Note, if the goal is constantly adapted as in table tennis (where the ball trajectory is not certain until the ball has bounced on the table the last time), these effects will produce significantly stronger undesired effects and, possibly, unstable behavior.

As an alternative for striking movements, we propose a modification of the dynamical system based motor primitives that allows us to directly specify the desired $\dot{\mathbf{x}}_1(T)$ while maintaining the duration of the movement and having the possibility to change the amplitude of the motion independently. For doing so, we introduce a moving goal and include the desired final velocity in Equation (3.2). We use a linearly

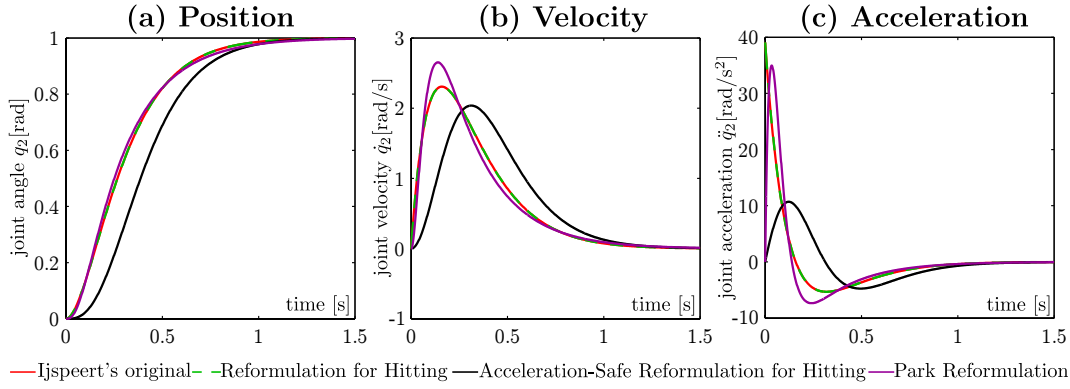


Figure 3.4: An important aspect of the Ijspeert framework is that such primitives are guaranteed to be stable and, hence, safe for learning. A problem of the regular formulation highly unevenly distributed acceleration with a jump at the beginning of the movement of its unaltered dynamics. These high accelerations affect the movement when the behavior is either generalized to new goals or when during trial-and-error learning where the initial parameters are small. Some of these problem have previously been noted by Park et al. [2008], and are particularly bad in the context of fast striking movements. Here, we compare the different formulations with respect to their acceleration in the unaltered dynamics case (i.e., $\mathbf{w} = \mathbf{0}$). For a better comparison, we set the goal velocity to zero ($\dot{\mathbf{g}} = \mathbf{0}$). The Ijspeert formulation clearly shows the problem with the large acceleration, as does the reformulation for hitting (with a hitting speed of $\dot{\mathbf{g}} = \mathbf{0}$ both are identical). While the Park modification starts without the jump in acceleration, it requires almost as high accelerations shortly afterwards. The acceleration-safe reformulation for hitting also starts out without a step in acceleration and does not require huge accelerations.

moving goal but other choices may be better suited for different tasks. This reformulation results in the following equations for the learned dynamics

$$\dot{\mathbf{x}}_2 = \tau \alpha_g \left(\beta_g (\mathbf{g}_m - \mathbf{x}_1) + \frac{(\dot{\mathbf{g}} - \dot{\mathbf{x}}_1)}{\tau} \right) + \tau \mathbf{A} \mathbf{f}, \quad (3.5)$$

$$\dot{\mathbf{x}}_1 = \tau \mathbf{x}_2,$$

$$\mathbf{g}_m = \mathbf{g}_m^0 - \dot{\mathbf{g}} \frac{\ln(z)}{\tau \alpha_h}, \quad (3.6)$$

where $\dot{\mathbf{g}}$ is the desired final velocity, \mathbf{g}_m is the moving goal and the initial position of the moving goal $\mathbf{g}_m^0 = \mathbf{g} - \tau \dot{\mathbf{g}}$ ensures that $\mathbf{g}_m(T) = \mathbf{g}$. The term $-\ln(z) / (\tau \alpha_h)$ is proportional to the time if the canonical system in Equation (3.1) runs unaltered; however, adaptation of z allows the straightforward adaptation of the hitting time. If $\dot{\mathbf{g}} = \mathbf{0}$, this formulation is exactly the same as the original formulation. The imitation learning approach mentioned in Section 3.2.1 can be adapted straightforwardly to this formulation. Figure 3.3 illustrates how the different approaches behave when forced to achieve a specified desired final velocity.

3.2.3 Safer Dynamics for Generalization

Generalizing fast movements such as a forehand in table tennis can become highly dangerous if the primitive requires exceedingly high accelerations or has large jumps in the acceleration (e.g., the fastest table tennis moves that we have executed on our WAM had a peak velocity of $7m/s$ and $10g$ maximum acceleration). Hence, the initial jump in acceleration often observed during the execution of the Ijspeert primitives may lead to desired accelerations that a physical robot cannot properly execute, and may even

cause damage to the robot system. In the following, we will discuss several sources of these acceleration jumps and how to overcome them. If the dynamics are not altered by the transformation function, i.e., $\mathbf{w} = \mathbf{0}$, the highest acceleration during the original Ijspeert motor primitive occurs at the very first time-step and then decays rapidly. If the motor primitives are properly initialized by imitation learning, the transformation function will cancel this initial acceleration, and, thus, this usually does not pose a problem in the absence of generalization. However, when changing the amplitude \mathbf{a} of the motion (e.g., in order to achieve a specific goal velocity) the transformation function will over- or undercompensate for this initial acceleration jump. The adaptation proposed in Section 3.2.2 does not require a change in amplitude, but suffers from a related shortcoming, i.e., changing the goal velocity also changes the initial position of the goal, thus results in a similar jump in acceleration that needs to be compensated. Using the motor primitives with an initial velocity that differs from the one used during imitation learning has the same effect. Figures 3.3 and 3.4 illustrate these initial steps in acceleration for various motor primitive formulations. As an alternative, we propose to gradually activate the attractor dynamics of the motor primitive (e.g., by reweighting them using the output of the canonical system). When combined, these two adaptations result in

$$\dot{\mathbf{x}}_2 = (1 - z) \tau \alpha_g \left(\beta_g (\mathbf{g}_m - \mathbf{x}_1) + \frac{(\dot{\mathbf{g}} - \dot{\mathbf{x}}_1)}{\tau} \right) + \tau \mathbf{A} \mathbf{f}. \quad (3.7)$$

Surprisingly, after this modification the unaltered dynamics (i.e., where $\mathbf{w} = \mathbf{0}$ and, hence, $\tau \mathbf{A} \mathbf{f}(z) = \mathbf{0}$) result in trajectories that roughly resemble a minimum jerk movements and, hence, look very similar to human movements. Exactly as for the Ijspeert formulation, we can arbitrarily shape the behavior by learning the weights of the transformation function. Note that [Park et al., 2008] also introduced a similar modification canceling the initial acceleration caused by the offset between initial and goal position; however, their approach cannot deal with a deviating initial velocity.

The proposed acceleration jump compensation also yields smoother movements during the adaptation of the hitting point as well as smoother transitions if motor primitives are sequenced. The later becomes particularly important when the preceding motor primitive has a significantly different velocity than during training (by imitation learning) or if it is terminated early due to external events. All presented modifications are compatible with the imitation learning approach discussed in Section 3.2.1 and the adaptation is straightforward. Figures 3.3 and 3.4 show how the presented modifications overcome the problems with the initial jumps in acceleration.

3.2.4 Imitation Learning

In the following chapters we use imitation learning from a single example to generate a sensible initial policy. This step can be performed efficiently in the context of dynamical systems motor primitives in both the original and adapted forms, as the transformation function Equation (3.3) is linear in its parameters. As a result, we can choose the weighted squared error (WSE)

$$\text{WSE}_n = \sum_{t=1}^T \psi_t^n \left(f_t^{\text{ref}} - z_t \theta^n \right)^2 \quad (3.8)$$

as cost function and minimize it for all parameter vectors θ^n with $n \in \{1, 2, \dots, N\}$. Here, the corresponding weighting functions are denoted by ψ_t^n and the basis function by z_t . The reference or target signal f_t^{ref} is the desired transformation function and t indicates the time-step of the sample. The error in Equation (3.8) can be rewritten in matrix form as

$$\text{WSE}_n = \left(\mathbf{f}^{\text{ref}} - \mathbf{Z} \theta^n \right)^T \mathbf{\Psi} \left(\mathbf{f}^{\text{ref}} - \mathbf{Z} \theta^n \right)$$

with \mathbf{f}^{ref} containing the values of f_t^{ref} for all time-steps t , $\Psi = \text{diag}(\psi_1^n, \dots, \psi_t^n, \dots, \psi_T^n)$, and $[\mathbf{Z}]_t = z_t$. As a result, we have a standard locally-weighted linear regression problem that is straightforward to solve and yields the unbiased parameter estimator

$$\theta^n = (\mathbf{Z}^T \Psi \mathbf{Z})^{-1} \mathbf{Z}^T \Psi \mathbf{f}^{\text{ref}}.$$

This approach was originally suggested for imitation learning by Ijspeert et al. [2002a]. Estimating the parameters of the dynamical system is slightly more difficult; the duration of the movement is extracted using motion detection and the time-constant is set accordingly.

3.3 Robot Evaluation

In this section, we evaluate the presented reactive templates for representing, learning and executing forehands in the setting of table tennis. For doing so, we evaluate our representation for striking movements first on hitting a hanging ball in Section 3.3.1 and, subsequently, in the task of returning a ball served by a ball launcher presented in Section 3.3.2.

When hitting a ping-pong ball that is hanging from the ceiling, the task consists of hitting the ball with an appropriate desired Cartesian velocity and orientation of the paddle. Hitting a ping-pong ball shot by a ball launcher requires predicting the ball's future positions and velocities in order to choose an interception point. The latter is only sufficiently accurate after the ball has hit the table for the last time. This short reaction time underlines that the movement templates can be adapted during the trajectory under strict time limitations when there is no recovery from a bad generalization, long replanning or inaccurate movements.

3.3.1 Generalizing Forehands on Static Targets

As a first experiment, we evaluated how well this new formulation of hitting primitives generalizes forehand movements learned from imitation as shown in Figure 3.6 (a). First, we collected arm, racket and ball trajectories for imitation learning using the 7 DoF Barrett WAM robot as an haptic input device for kinesthetic teach-in where all inertial forces and gravity were compensated. In the second step, we employ this data to automatically extract the duration of the striking movement, the duration of the

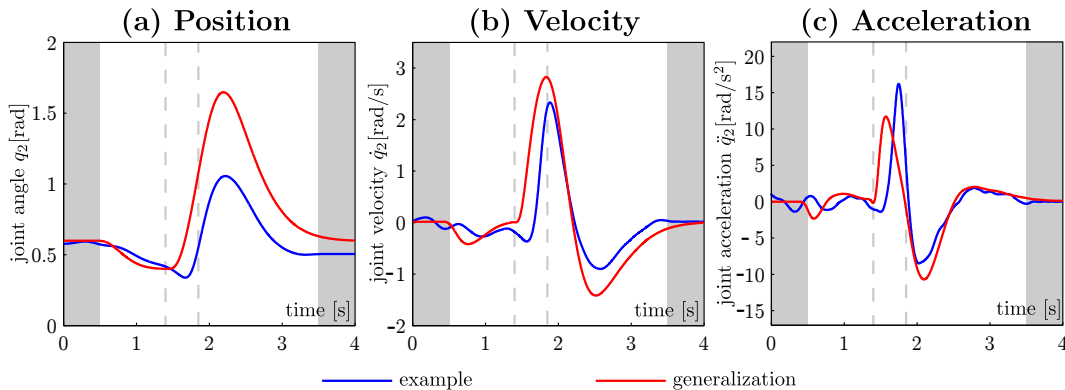





Figure 3.5: This figure demonstrates the generalization of an imitated behavior to a different target that is 15cm away from the original target. Note that this trajectory is for a static target, hence the slow motion. The depicted degree of freedom (DoF) is shoulder adduction-abduction (i.e., the second DoF). The solid gray bars indicate the time before and after the main movement, the gray dashed lines indicate the phase borders also depicted in Figure 3.1 and the target is hit at the second border.



(a) Demonstration by a Human Instructor



(b) Example: Reproduction for Hitting a Stationary Ball



(c) Application: Returning Balls launched by a Ball Gun

Figure 3.6: This figure presents a hitting sequence from the demonstration, a generalization on the robot with a ball attached by a string as well as a generalization hitting a ball shot by a ping-pong ball launcher. The demonstration and the flying ball generalization are captured by a $25Hz$ video camera, the generalization with the attached ball is captured with $200Hz$ through our vision system. From left to right the stills represent: rest posture, swing-back posture, hitting point, swing-through and rest posture. The postures (①-④) are the same as in Figure 3.2.

individual phases as well as the Cartesian target velocity and orientation of the racket when hitting the ball. We employ a model (as shown in Section 3.2) that has phases for swinging back, hitting and going to a rest posture. Both the phase for swing-back and return-to-home phases will go into intermediary still phases while the hitting phase goes through a target point with a pre-specified target velocity. All phases can only be safely executed due to the “safer dynamics” which we introduced in Section 3.2.3.

In this experiment, the ball is a stationary target and detected by a stereo camera setup. Subsequently, the supervisory level proposed in [Mülling and Peters, 2009] determines the hitting point and the striking velocity in configuration space. The motor primitives are adjusted accordingly and executed on the robot in joint-space using an inverse dynamics control law. The robot successfully hits the ball at different positions within a diameter of approximately $1.2m$ if kinematically feasible. The adaptation for striking movements achieves the desired velocities and the safer dynamics allow generalization to a much larger area while successfully removing the possibly large accelerations at the transitions between motor primitives. See Figure 3.5 for a comparison of the training example and the generalized motion for one degree of freedom and Figure 3.6 (b) for a few frames from a hit of a static ball.

3.3.2 Playing against a Ball Launcher

This evaluation adds an additional layer of complexity as the hitting point and the hitting time has to be estimated from the trajectory of the ball and continuously adapted as the hitting point cannot be reliably determined until the ball has bounced off the table for the last time. In this setting, the ball is tracked by two overlapping high speed stereo vision setups with $200Hz$ cameras. In order to obtain better estimates




Figure 3.7: Generalization to various targets (five different forehands at posture ③) are shown approximately when hitting the ball.

of the current position and to calculate the velocities, the raw 3D positions are filtered by a specialized Kalman filter [Kalman, 1960] that takes contacts of the ball with the table and the racket into account [Mülling and Peters, 2009]. When used as a Kalman predictor, we can again determine the target point for the primitive with a pre-specified target velocity with the method described in [Mülling and Peters, 2009]. The results obtained for the still ball generalize well from the static ball to the one launched by a ball launcher at $3m/s$ which are returned at speeds up to $8m/s$. A sequence of frames from the attached video is shown in Figure 3.6. The plane of possible virtual hitting points again has a diameter of roughly $1m$ as shown in Figure 3.7. The modified motor primitives generated movements with the desired hitting position and velocity. The robot hit the ball in the air in approx. 95% of the trials. However, due to a simplistic ball model and execution inaccuracies the ball was often not properly returned on the table. Please see the videos accompanying this chapter <http://www.robot-learning.de/Research/HittingMPs>.

Note that our results differ significantly from previous approaches as we use a framework that allows us to learn striking movements from human demonstrations unlike previous work in batting [Senoo et al., 2006] and table tennis [Andersson, 1988]. Unlike baseball which only requires four degrees of freedom (as, e.g., in [Senoo et al., 2006] who used a 4 DoF WAM arm in a manually coded high speed setting), and previous work in table tennis (which had only low-inertia, was overpowered and had mostly prismatic joints [Andersson, 1988, Fässler et al., 1990, Matsushima et al., 2005]), we use a full seven degrees of freedom revolutionary joint robot and, thus, have to deal with larger inertia as the wrist adds roughly $2.5kg$ weight at the elbow. Hence, it was essential to train trajectories by imitation learning that distribute the torques well over the redundant joints as the human teacher was suffering from the same constraints.

3.4 Conclusion

In this paper, we rethink previous work on dynamic systems motor primitive [Ijspeert et al., 2002a, Schaal et al., 2003, 2007] in order to obtain movement templates that can be used reactively in batting and hitting sports. This reformulation allows to change the target velocity of the movement while maintaining the overall duration and shape. Furthermore, we present a modification that overcomes the problem of an initial acceleration step which is particularly important for safe generalization of learned movements. Our adaptations retain the advantages of the original formulation and perform well in practice. We evaluate this novel motor primitive formulation first in hitting a stationary table tennis ball and, subsequently, in returning ball served by a ping pong ball launcher. In both cases, the novel motor primitives manage to generalize well while maintaining the features of the demonstration. This new formulation of the motor primitives can hopefully be used together with *meta-parameter learning* (Chapter 5) in a *mixture of motor primitives* [Mülling et al., 2010] in order to create a complete framework for learning tasks like table tennis autonomously.

4 Policy Search for Motor Primitives in Robotics

Many motor skills in humanoid robotics can be learned using parametrized motor primitives. While successful applications to date have been achieved with imitation learning, most of the interesting motor learning problems are high-dimensional reinforcement learning problems. These problems are often beyond the reach of current reinforcement learning methods. In this chapter, we study parametrized policy search methods and apply these to benchmark problems of motor primitive learning in robotics. We show that many well-known parametrized policy search methods can be derived from a general, common framework. This framework yields both policy gradient methods and expectation-maximization (EM) inspired algorithms. We introduce a novel EM-inspired algorithm for policy learning that is particularly well-suited for dynamical system motor primitives. We compare this algorithm, both in simulation and on a real robot, to several well-known parametrized policy search methods such as episodic REINFORCE, ‘Vanilla’ Policy Gradients with optimal baselines, episodic Natural Actor Critic, and episodic Reward-Weighted Regression. We show that the proposed method out-performs them on an empirical benchmark of learning dynamical system motor primitives both in simulation and on a real robot. We apply it in the context of motor learning and show that it can learn a complex Ball-in-a-Cup task on a real Barrett WAM robot arm.

4.1 Introduction

To date, most robots are still taught by a skilled human operator either via direct programming or a teach-in. Learning approaches for automatic task acquisition and refinement would be a key step for making robots progress towards autonomous behavior. Although imitation learning can make this task more straightforward, it will always be limited by the observed demonstrations. For many motor learning tasks, skill transfer by imitation learning is prohibitively hard given that the human teacher is not capable of conveying sufficient task knowledge in the demonstration. In such cases, reinforcement learning is often an alternative to a teacher’s presentation, or a means of improving upon it. In the high-dimensional domain of anthropomorphic robotics with its continuous states and actions, reinforcement learning suffers particularly from the curse of dimensionality. However, by using a task-appropriate policy representation and encoding prior knowledge into the system by imitation learning, local reinforcement learning approaches are capable of dealing with the problems of this domain. Policy search (also known as policy learning) is particularly well-suited in this context, as it allows the usage of domain-appropriate pre-structured policies [Toussaint and Goerick, 2007], the straightforward integration of a teacher’s presentation [Guenter et al., 2007, Peters and Schaal, 2006] as well as fast online learning [Bagnell et al., 2004, Ng and Jordan, 2000, Hoffman et al., 2007]. Recently, policy search has become an accepted alternative of value-function-based reinforcement learning [Bagnell et al., 2004, Strens and Moore, 2001, Kwee et al., 2001, Peshkin, 2001, El-Fakdi et al., 2006, Taylor et al., 2007] due to many of these advantages.

In this chapter, we will introduce a policy search framework for episodic reinforcement learning and show how it relates to policy gradient methods [Williams, 1992, Sutton et al., 1999, Lawrence et al., 2003, Tedrake et al., 2004, Peters and Schaal, 2006] as well as expectation-maximization (EM) inspired algorithms [Dayan and Hinton, 1997, Peters and Schaal, 2007]. This framework allows us to re-derive or to generalize well-known approaches such as episodic REINFORCE [Williams, 1992], the policy gradient theorem [Sutton et al., 1999, Peters and Schaal, 2006], the episodic Natural Actor Critic [Peters et al., 2003, 2005], and an episodic generalization of the Reward-Weighted Regression [Peters and Schaal, 2007]. We derive a new algorithm called Policy Learning by Weighting Exploration with the Returns (PoWER), which is particularly well-suited for the learning of trial-based tasks in motor control.

We evaluate the algorithms derived from this framework to determine how they can be used for refining parametrized policies in robot skill learning. To address this problem, we follow a methodology suitable for robotics where the policy is first initialized by imitation learning and, subsequently, the policy search algorithm is used for self-improvement. As a result, we need a suitable representation in order to apply this approach in anthropomorphic robot systems. In imitation learning, a particular kind of motor control policy has been very successful, which is known as dynamical system motor primitives [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007]. In this approach, dynamical systems are used to encode a control policy suitable for motor tasks. The representation is linear in the parameters; hence, it can be learned straightforwardly from demonstrations. Such dynamical system motor primitives can represent both point-to-point and rhythmic behaviors. We focus on the point-to-point variant which is suitable for representing single-stroke, episodic behaviors. As a result, they are particularly well-suited for episodic policy search.

We show that all presented algorithms work sufficiently well when employed in the context of learning dynamical system motor primitives in different benchmark and application settings. We compare these methods on the two benchmark problems from [Peters and Schaal, 2006] for dynamical system motor primitives learning, the Underactuated Swing-Up [Atkeson, 1994] robotic benchmark problem, and the Casting task. Using entirely different parametrizations, we evaluate policy search methods on the mountain-car benchmark [Sutton and Barto, 1998] and the Tetherball Target Hitting task. On the mountain-car benchmark, we additionally compare to a value function based approach. The method with the best performance, PoWER, is evaluated on the complex task of Ball-in-a-Cup [Sumners, 1997]. Both the Underactuated Swing-Up as well as Ball-in-a-Cup are achieved on a real Barrett WAM robot arm. Please also refer to the videos at <http://www.robot-learning.de/Research/ReinforcementLearning>. For all real robot experiments, the presented movement is learned by imitation from a kinesthetic demonstration, and the Barrett WAM robot arm subsequently improves its behavior by reinforcement learning.

4.2 Policy Search for Parametrized Motor Primitives

Our goal is to find reinforcement learning techniques that can be applied in robotics in the context of learning high-dimensional motor control tasks. We first introduce the required notation for the derivation of the reinforcement learning framework in Section 4.2.1. We discuss the problem in the general setting of reinforcement learning using a generalization of the approach in [Dayan and Hinton, 1997, Attias, 2003, Peters and Schaal, 2007]. We extend the existing approach to episodic reinforcement learning for continuous states, in a manner suitable for robotics.

We derive a new expectation-maximization (EM) inspired algorithm [Dempster et al., 1977] called Policy Learning by Weighting Exploration with the Returns (PoWER) in Section 4.2.3 and show how the general framework is related to policy gradient methods and the Reward-Weighted Regression method in Section 4.2.2.

4.2.1 Problem Statement & Notation

In this chapter, we treat motor primitive learning problems in the framework of reinforcement learning [Sutton and Barto, 1998] with a strong focus on the episodic case. At time t , there is an actor in a state \mathbf{s}_t that chooses an action \mathbf{a}_t according to a stochastic policy $\pi(\mathbf{a}_t|\mathbf{s}_t, t)$. Such a policy is a probability distribution over actions given the current state and time. The stochastic formulation allows a natural incorporation of exploration, and the optimal time-invariant policy has been shown to be stochastic in the case of hidden state variables [Sutton et al., 1999, Jaakkola et al., 1993]. Upon the completion of the action, the actor transfers to a state \mathbf{s}_{t+1} and receives a reward r_t . As we are interested in learning complex motor tasks consisting of a single stroke [Schaal et al., 2007], we focus on finite horizons of length T with episodic restarts and learn the optimal parametrized, stochastic policy for such episodic reinforcement learning problems [Sutton and Barto, 1998]. We assume an explorative parametrized

policy π with parameters $\theta \in \mathbb{R}^n$. In Section 4.3.1, we discuss how the dynamical system motor primitives [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007] can be employed in this setting. In this section, we will keep most derivations sufficiently general such that they are transferable to various other parametrized policies that are linear in the parameters.

The general goal in reinforcement learning is to optimize the *expected return* of the policy π with parameters θ defined by

$$J(\theta) = \int_{\mathbb{T}} p_{\theta}(\tau) R(\tau) d\tau,$$

where \mathbb{T} is the set of all possible paths. A rollout $\tau = [\mathbf{s}_{1:T+1}, \mathbf{a}_{1:T}]$, also called path, episode or trial, denotes a series of states $\mathbf{s}_{1:T+1} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_{T+1}]$ and actions $\mathbf{a}_{1:T} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_T]$. The probability of rollout τ is denoted by $p_{\theta}(\tau)$, while $R(\tau)$ refers to its aggregated return. Using the standard Markov assumption and additive accumulated rewards, we can write

$$\begin{aligned} p_{\theta}(\tau) &= p(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) \pi(\mathbf{a}_t | \mathbf{s}_t, t), \\ R(\tau) &= T^{-1} \sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, t), \end{aligned} \quad (4.1)$$

where $p(\mathbf{s}_1)$ denotes the initial state distribution, $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ the next state distribution conditioned on the last state and action, and $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, t)$ denotes the immediate reward.

While episodic Reinforcement Learning (RL) problems with finite horizons are common in both human Wulf [2007] and robot motor control problems, few methods exist in the RL literature. Examples are episodic REINFORCE [Williams, 1992], the episodic Natural Actor Critic eNAC [Peters et al., 2003, 2005] and model-based methods using differential dynamic programming [Atkeson, 1994].

4.2.2 Episodic Policy Learning

In this section, we discuss episodic reinforcement learning in policy space, which we will refer to as Episodic Policy Learning. We first discuss the lower bound on the expected return as suggested in [Dayan and Hinton, 1997] for guaranteeing that policy update steps are improvements. In [Dayan and Hinton, 1997, Peters and Schaal, 2007] only the immediate reward case is discussed; we extend this framework to episodic reinforcement learning. Subsequently, we derive a general update rule, which yields the policy gradient theorem [Sutton et al., 1999], a generalization of the reward-weighted regression [Peters and Schaal, 2007], as well as the novel Policy learning by Weighting Exploration with the Returns (PoWER) algorithm.

Bounds on Policy Improvements

Unlike in reinforcement learning, other branches of machine learning have focused on maximizing lower bounds on the cost functions, which often results in expectation-maximization (EM) algorithms [McLachan and Krishnan, 1997]. The reasons for this preference also apply in policy learning: if the lower bound also becomes an equality for the sampling policy, we can guarantee that the policy will be improved by maximizing the lower bound. Results from supervised learning can be transferred with ease. First, we generalize the scenario suggested by Dayan and Hinton [1997] to the episodic case. Here, we generate rollouts τ using the current policy with parameters θ , which we then weight with the returns $R(\tau)$, and subsequently match it with a new policy parametrized by θ' . This matching of the success-weighted path distribution is equivalent to minimizing the Kullback-Leibler divergence $D(p_{\theta}(\tau)R(\tau) || p_{\theta'}(\tau))$ between the new path distribution $p_{\theta'}(\tau)$ and the reward-weighted previous one $p_{\theta}(\tau)R(\tau)$. The Kullback-Leibler divergence is considered a natural distance measure between probability distributions [Bagnell and Schneider, 2003, van der Maaten et al., 2009]. As shown in [Dayan and Hinton, 1997, Peters and Schaal,

2007], such a derivation results in a lower bound on the expected return using Jensen's inequality and the concavity of the logarithm. Thus, we obtain

$$\begin{aligned}\log J(\theta') &= \log \int_{\mathbb{T}} p_{\theta'}(\tau) R(\tau) d\tau = \log \int_{\mathbb{T}} \frac{p_{\theta}(\tau)}{p_{\theta}(\tau)} p_{\theta'}(\tau) R(\tau) d\tau, \\ &\geq \int_{\mathbb{T}} p_{\theta}(\tau) R(\tau) \log \frac{p_{\theta'}(\tau)}{p_{\theta}(\tau)} d\tau + \text{const},\end{aligned}$$

which is proportional to

$$-D(p_{\theta}(\tau) R(\tau) \| p_{\theta'}(\tau)) = L_{\theta}(\theta'),$$

where

$$D(p(\tau) \| q(\tau)) = \int p(\tau) \log \frac{p(\tau)}{q(\tau)} d\tau$$

denotes the Kullback-Leibler divergence, and the constant is needed for tightness of the bound. Note that $p_{\theta}(\tau) R(\tau)$ is an improper probability distribution as pointed out by Dayan and Hinton [1997]. The policy improvement step is equivalent to maximizing the lower bound on the expected return $L_{\theta}(\theta')$, and we will now show how it relates to previous policy learning methods.

Resulting Policy Updates

In this section, we will discuss three different policy updates, which are directly derived from the results of Section 4.2.2. First, we show that policy gradients [Williams, 1992, Sutton et al., 1999, Lawrence et al., 2003, Tedrake et al., 2004, Peters and Schaal, 2006] can be derived from the lower bound $L_{\theta}(\theta')$, which is straightforward from a supervised learning perspective [Binder et al., 1997]. Subsequently, we show that natural policy gradients [Bagnell and Schneider, 2003, Peters and Schaal, 2006] can be seen as an additional constraint regularizing the change in the path distribution resulting from a policy update when improving the policy incrementally. Finally, we will show how expectation-maximization (EM) algorithms for policy learning can be generated.

Policy Gradients.

When differentiating the function $L_{\theta}(\theta')$ that defines the lower bound on the expected return, we directly obtain

$$\partial_{\theta'} L_{\theta}(\theta') = \int_{\mathbb{T}} p_{\theta}(\tau) R(\tau) \partial_{\theta'} \log p_{\theta'}(\tau) d\tau = E \left\{ \left(\sum_{t=1}^T \partial_{\theta'} \log \pi(\mathbf{a}_t | \mathbf{s}_t, t) \right) R(\tau) \right\}, \quad (4.2)$$

where

$$\partial_{\theta'} \log p_{\theta'}(\tau) = \sum_{t=1}^T \partial_{\theta'} \log \pi(\mathbf{a}_t | \mathbf{s}_t, t)$$

denotes the log-derivative of the path distribution. As this log-derivative depends only on the policy we can estimate a gradient from rollouts, without having a model, by simply replacing the expectation by a sum. When θ' is close to θ , we have the policy gradient estimator, which is widely known as episodic REINFORCE [Williams, 1992]

$$\lim_{\theta' \rightarrow \theta} \partial_{\theta'} L_{\theta}(\theta') = \partial_{\theta} J(\theta).$$

See Algorithm 4.1 for an example implementation of this algorithm and Appendix 4.A.1 for the detailed steps of the derivation. A MATLAB implementation of this algorithm is available at <http://www.robot-learning.de/Member/JensKober>.

Algorithm 4.1 ‘Vanilla’ Policy Gradients (VPG)

Input: initial policy parameters θ_0

repeat

Sample: Perform $h = \{1, \dots, H\}$ rollouts using $\mathbf{a} = \theta^\top \phi(\mathbf{s}, t) + \epsilon_t$ with $[\epsilon_t^n] \sim \mathcal{N}(0, (\sigma^{h,n})^2)$ as stochastic policy and collect all $(t, \mathbf{s}_t^h, \mathbf{a}_t^h, \mathbf{s}_{t+1}^h, \epsilon_t^h, r_{t+1}^h)$ for $t = \{1, 2, \dots, T+1\}$.

Compute: Return $R^h = \sum_{t=1}^{T+1} r_t^h$, eligibility

$$\psi^{h,n} = \frac{\partial \log p(\tau^h)}{\partial \theta^n} = \sum_{t=1}^T \frac{\partial \log \pi(a_t^h | s_t^h, t)}{\partial \theta^n} = \sum_{t=1}^T \frac{\epsilon_t^{h,n}}{(\sigma_h^n)^2} \phi^n(s_t^{h,n}, t)$$

and baseline

$$b^n = \frac{\sum_{h=1}^H (\psi^{h,n})^2 R^h}{\sum_{h=1}^H (\psi^{h,n})^2}$$

for each parameter $n = \{1, \dots, N\}$ from rollouts.

Compute Gradient:

$$g_{\text{VP}}^n = E \left\{ \frac{\partial \log p(\tau^h)}{\partial \theta^n} (R(\tau^h) - b^n) \right\} = \frac{1}{H} \sum_{h=1}^H \psi^{h,n} (R^h - b^n).$$

Update policy using

$$\theta_{k+1} = \theta_k + \alpha g_{\text{VP}}.$$

until Convergence $\theta_{k+1} \approx \theta_k$.

Algorithm 4.2 episodic Natural Actor Critic (eNAC)

Input: initial policy parameters θ_0

repeat

Sample: Perform $h = \{1, \dots, H\}$ rollouts using $\mathbf{a} = \theta^T \phi(\mathbf{s}, t) + \epsilon_t$ with $[\epsilon_t^n] \sim \mathcal{N}(0, (\sigma^{h,n})^2)$ as stochastic policy and collect all $(t, \mathbf{s}_t^h, \mathbf{a}_t^h, \mathbf{s}_{t+1}^h, \epsilon_t^h, r_{t+1}^h)$ for $t = \{1, 2, \dots, T+1\}$.

Compute: Return $R^h = \sum_{t=1}^{T+1} r_t^h$ and eligibility $\psi^{h,n} = \sum_{t=1}^T (\sigma_h^n)^{-2} \epsilon_t^{h,n} \phi^n(s_t^{h,n}, t)$ for each parameter $n = \{1, \dots, N\}$ from rollouts.

Compute Gradient:

$$[\mathbf{g}_{\text{eNAC}}, R_{\text{ref}}]^T = (\Psi^T \Psi)^{-1} \Psi^T \mathbf{R}$$

with $\mathbf{R} = [R^1, \dots, R^H]^T$ and $\Psi = \begin{bmatrix} \psi^1, \dots, \psi^H \\ 1, \dots, 1 \end{bmatrix}^T$ where $\psi^h = [\psi^{h,1}, \dots, \psi^{h,N}]^T$.

Update policy using

$$\theta_{k+1} = \theta_k + \alpha \mathbf{g}_{\text{eNAC}}.$$

until Convergence $\theta_{k+1} \approx \theta_k$.

A reward, which precedes an action in a rollout, can neither be caused by the action nor cause an action in the same rollout. Thus, when inserting Equation (4.1) into Equation (4.2), all cross-products between r_t and $\partial_{\theta'} \log \pi(\mathbf{a}_{t+\delta t} | \mathbf{s}_{t+\delta t}, t + \delta t)$ for $\delta t > 0$ become zero in expectation [Peters and Schaal, 2006]. Therefore, we can omit these terms and rewrite the estimator as

$$\partial_{\theta'} L_{\theta}(\theta') = E \left\{ \sum_{t=1}^T \partial_{\theta'} \log \pi(\mathbf{a}_t | \mathbf{s}_t, t) Q^{\pi}(\mathbf{s}, \mathbf{a}, t) \right\}, \quad (4.3)$$

where

$$Q^{\pi}(\mathbf{s}, \mathbf{a}, t) = E \left\{ \sum_{\tilde{t}=t}^T r(\mathbf{s}_{\tilde{t}}, \mathbf{a}_{\tilde{t}}, \mathbf{s}_{\tilde{t}+1}, \tilde{t}) | \mathbf{s}_t = \mathbf{s}, \mathbf{a}_t = \mathbf{a} \right\}$$

is called the state-action value function [Sutton and Barto, 1998]. Equation (4.3) is equivalent to the policy gradient theorem [Sutton et al., 1999] for $\theta' \rightarrow \theta$ in the infinite horizon case, where the dependence on time t can be dropped.

The derivation results in the episodic Natural Actor Critic as discussed in [Peters et al., 2003, 2005] when adding an additional cost in Equation (4.2) to penalize large steps away from the observed path distribution. Such a regularization can be achieved by restricting the amount of change in the path distribution and subsequently, determining the steepest descent for a fixed step away from the observed trajectories. Change in probability distributions is naturally measured using the Kullback-Leibler divergence, thus after adding the additional constraint of

$$D(p_{\theta}(\tau) || p_{\theta'}(\tau)) \approx 0.5 (\theta' - \theta)^T \mathbf{F}(\theta) (\theta' - \theta) = \delta$$

using a second-order expansion as an approximation where $\mathbf{F}(\theta)$ denotes the Fisher information matrix [Bagnell and Schneider, 2003, Peters et al., 2003, 2005]. See Algorithm 4.2 for an example implementation of the episodic Natural Actor Critic. A MATLAB implementation of this algorithm is available at <http://www.robot-learning.de/Member/JensKober>.

Policy Search via Expectation Maximization.

One major drawback of gradient-based approaches is the learning rate, which is an open parameter that can be hard to tune in control problems but is essential for good performance. Expectation-Maximization

algorithms are well-known to avoid this problem in supervised learning while even yielding faster convergence [McLachan and Krishnan, 1997]. Previously, similar ideas have been explored in immediate reinforcement learning [Dayan and Hinton, 1997, Peters and Schaal, 2007]. In general, an EM-algorithm chooses the next policy parameters θ_{n+1} such that

$$\theta_{n+1} = \operatorname{argmax}_{\theta'} L_{\theta}(\theta').$$

In the case where $\pi(\mathbf{a}_t|\mathbf{s}_t, t)$ belongs to the exponential family, the next policy can be determined analytically by setting Equation (4.2) or Equation (4.3) to zero

$$E \left\{ \sum_{t=1}^T \partial_{\theta'} \log \pi(\mathbf{a}_t|\mathbf{s}_t, t) Q(\mathbf{s}, \mathbf{a}, t) \right\} = 0, \quad (4.4)$$

and solving for θ' . Depending on the choice of stochastic policy, we will obtain different solutions and different learning algorithms. It allows the extension of the reward-weighted regression to longer horizons as well as the introduction of the Policy learning by Weighting Exploration with the Returns (PoWER) algorithm.

4.2.3 Policy learning by Weighting Exploration with the Returns (PoWER)

In most learning control problems, we attempt to have a deterministic mean policy $\bar{\mathbf{a}} = \theta^T \boldsymbol{\phi}(\mathbf{s}, t)$ with parameters θ and basis functions $\boldsymbol{\phi}$. In Section 4.3.1, we will introduce a particular type of basis function well-suited for robotics. These basis functions derive from the motor primitive formulation. Given such a deterministic mean policy $\bar{\mathbf{a}} = \theta^T \boldsymbol{\phi}(\mathbf{s}, t)$, we generate a stochastic policy using additive exploration $\epsilon(\mathbf{s}, t)$ in order to make model-free reinforcement learning possible. We have a policy $\pi(\mathbf{a}_t|\mathbf{s}_t, t)$ that can be brought into the form

$$\mathbf{a} = \theta^T \boldsymbol{\phi}(\mathbf{s}, t) + \epsilon(\boldsymbol{\phi}(\mathbf{s}, t)).$$

Previous work in this setting [Williams, 1992, Guenter et al., 2007, Peters and Schaal, 2006, 2007], with the notable exception of [Rückstieß et al., 2008], has focused on state-independent, white Gaussian exploration, namely $\epsilon(\boldsymbol{\phi}(\mathbf{s}, t)) \sim \mathcal{N}(\epsilon|\mathbf{0}, \Sigma)$. It is straightforward to obtain the Reward-Weighted Regression for episodic RL by solving Equation (4.4) for θ' , which naturally yields a weighted regression method with the state-action values $Q^\pi(\mathbf{s}, \mathbf{a}, t)$ as weights. See Algorithm 4.3 for an exemplary implementation and Appendix 4.A.2 for the derivation. An optimized MATLAB implementation of this algorithm is available at <http://www.robot-learning.de/Member/JensKober>. This form of exploration has resulted in various applications in robotics such as T-Ball batting [Peters and Schaal, 2006], Peg-In-Hole [Gullapalli et al., 1994], constrained reaching movements [Guenter et al., 2007] and operational space control [Peters and Schaal, 2007].

However, such unstructured exploration at every step has several disadvantages: (i) it causes a large variance in parameter updates that grows with the number of time-steps [Rückstieß et al., 2008, Peters and Schaal, 2006], (ii) it perturbs actions too frequently as the system acts as a low pass filter, and the perturbations average out, thus their effect is washed out, and (iii) it can damage the system executing the trajectory. As the action is perturbed in every time-step the outcome of a trial can change drastically. This effect accumulates with the number of trials and the exploration is not equal over the progress of the trial. This behavior leads to a large variance in parameter updates. Random exploration in every time-step leads to jumps in the actions. A physical robot can not execute instantaneous changes in actions as either the controller needs time to react or the motor and the links of the robot have inertia that forces the robot to continue the motion induced by the previous actions. Globally speaking, the system acts as a low pass filter. If the robot tries to follow the desired high frequency action changes, a lot of strain is placed on the mechanics of the robot and can lead to oscillations. Furthermore, the accumulating effect of the exploration can lead the robot far from previously seen states, which is potentially dangerous.

Algorithm 4.3 episodic Reward Weighted Regression (eRWR)

Input: initial policy parameters θ_0

repeat

Sample: Perform $h = \{1, \dots, H\}$ rollouts using $\mathbf{a} = \theta^T \phi(\mathbf{s}, t) + \epsilon_t$ with $[\epsilon_t^n] \sim \mathcal{N}(0, (\sigma^{h,n})^2)$ as stochastic policy and collect all $(t, \mathbf{s}_t^h, \mathbf{a}_t^h, \mathbf{s}_{t+1}^h, \epsilon_t^h, r_{t+1}^h)$ for $t = \{1, 2, \dots, T+1\}$.

Compute: State-action value function $Q_t^{\pi,h} = \sum_{\tilde{t}=t}^T r_{\tilde{t}}^h$ from rollouts.

Update policy using

$$\theta_{k+1}^n = \left((\Phi^n)^T \mathbf{Q}^\pi \Phi^n \right)^{-1} (\Phi^n)^T \mathbf{Q}^\pi \mathbf{A}^n$$

with basis functions

$$\Phi^n = \left[\phi_1^{1,n}, \dots, \phi_T^{1,n}, \phi_1^{2,n}, \dots, \phi_1^{H,n}, \dots, \phi_T^{H,n} \right]^T,$$

where $\phi_t^{h,n}$ is the value of the basis function of rollout h and parameter n at time t , actions

$$\mathbf{A}^n = \left[a_1^{1,n}, \dots, a_T^{1,n}, a_1^{2,n}, \dots, a_1^{H,n}, \dots, a_T^{H,n} \right]^T,$$

and returns

$$\mathbf{Q}^\pi = \text{diag} \left(Q_1^{\pi,1}, \dots, Q_T^{\pi,1}, Q_1^{\pi,2}, \dots, Q_1^{\pi,H}, \dots, Q_T^{\pi,H} \right)$$

until Convergence $\theta_{k+1} \approx \theta_k$.

As a result, all methods relying on this state-independent exploration have proven too fragile for learning tasks such as the Ball-in-a-Cup (see Section 4.3.7) on a real robot system. Alternatively, as introduced in [Rückstieß et al., 2008], one could generate a form of structured, state-dependent exploration. We use

$$\epsilon(\phi(\mathbf{s}, t)) = \epsilon_t^T \phi(\mathbf{s}, t)$$

with $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \hat{\Sigma})$, where $\hat{\Sigma}$ is a hyper-parameter of the exploration that can be optimized in a similar manner (see Appendix 4.A.3). This argument results in the policy

$$\mathbf{a} \sim \pi(\mathbf{a}_t | \mathbf{s}_t, t) = \mathcal{N}(\mathbf{a} | \theta^T \phi(\mathbf{s}, t), \phi(\mathbf{s}, t)^T \hat{\Sigma} \phi(\mathbf{s}, t)).$$

Inserting the resulting policy into Equation (4.4), we obtain the optimality condition update and can derive the update rule

$$\theta' = \theta + E \left\{ \sum_{t=1}^T \mathbf{W}(\mathbf{s}, t) Q^\pi(\mathbf{s}, \mathbf{a}, t) \right\}^{-1} E \left\{ \sum_{t=1}^T \mathbf{W}(\mathbf{s}, t) \epsilon_t Q^\pi(\mathbf{s}, \mathbf{a}, t) \right\}$$

with $\mathbf{W}(\mathbf{s}, t) = \phi(\mathbf{s}, t) \phi(\mathbf{s}, t)^T (\phi(\mathbf{s}, t)^T \hat{\Sigma} \phi(\mathbf{s}, t))^{-1}$.

In order to reduce the number of rollouts in this on-policy scenario, we reuse the rollouts through importance sampling as described, in the context of reinforcement learning, in [Andrieu et al., 2003, Sutton and Barto, 1998]. The expectations $E\{\cdot\}$ are replaced by the importance sampler denoted by $\langle \cdot \rangle_{w(\tau)}$. To avoid the fragility sometimes resulting from importance sampling in reinforcement learning, samples with very small importance weights are discarded. This step is necessary as a lot of rollouts with a low return accumulate mass and can bias the update. A simple heuristic that works well in practice is to discard all but the j best rollouts, where j is chosen in the same order of magnitude as the number of parameters N . The derivation is shown in Appendix 4.A.3 and the resulting algorithm in Algorithm 4.4. Note that for our motor primitives, some simplifications of \mathbf{W} are possible. These and other simplifications are shown in Appendix 4.A.3. A MATLAB implementation of this algorithm in

Algorithm 4.4 EM Policy learning by Weighting Exploration with the Returns (PoWER)

Input: initial policy parameters θ_0

repeat

Sample: Perform rollout(s) using $\mathbf{a} = (\theta + \epsilon_t)^T \phi(\mathbf{s}, t)$ with $\epsilon_t^T \phi(\mathbf{s}, t) \sim \mathcal{N}(0, \phi(\mathbf{s}, t)^T \hat{\Sigma} \phi(\mathbf{s}, t))$ as stochastic policy and collect all $(t, \mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, \epsilon_t, r_{t+1})$ for $t = \{1, 2, \dots, T + 1\}$.

Estimate: Use unbiased estimate

$$\hat{Q}^\pi(\mathbf{s}, \mathbf{a}, t) = \sum_{\tilde{t}=t}^T r(\mathbf{s}_{\tilde{t}}, \mathbf{a}_{\tilde{t}}, \mathbf{s}_{\tilde{t}+1}, \tilde{t}).$$

Reweight: Compute importance weights and reweight rollouts, discard low-importance rollouts.

Update policy using

$$\theta_{k+1} = \theta_k + \left\langle \sum_{t=1}^T \mathbf{W}(\mathbf{s}, t) Q^\pi(\mathbf{s}, \mathbf{a}, t) \right\rangle_{w(\tau)}^{-1} \left\langle \sum_{t=1}^T \mathbf{W}(\mathbf{s}, t) \epsilon_t Q^\pi(\mathbf{s}, \mathbf{a}, t) \right\rangle_{w(\tau)}$$

with $\mathbf{W}(\mathbf{s}, t) = \phi(\mathbf{s}, t) \phi(\mathbf{s}, t)^T (\phi(\mathbf{s}, t)^T \hat{\Sigma} \phi(\mathbf{s}, t))^{-1}$.

until Convergence $\theta_{k+1} \approx \theta_k$.

several variants is available at <http://www.robot-learning.de/Member/JensKober>. As we will see in Section 4.3, this PoWER method significantly outperforms all other described methods.

PoWER is very robust with respect to reward functions. The key constraint is that it has to be an improper probability distribution which means that the rewards have to be positive. It can be beneficial for learning speed if the reward function sums up to one as a proper probability distribution.

Like most learning algorithms, PoWER achieves convergence faster for lower numbers of parameters. However, as it is an EM-inspired approach, it suffers significantly less from this problem than gradient based approaches. Including more prior knowledge, either in the parametrization or the initial policy, leads to faster convergence. As discussed above, changing exploration at every time-step has a number of disadvantages. Fixing the exploration for the whole episode (if each basis function is only active for a short time) or using a slowly varying exploration (for example based on random walks) can increase the performance. All algorithms introduced in this chapter optimize locally and can get stuck in local optima. An initial policy should be chosen to avoid local optima on the progress towards the desired final solution.

4.3 Benchmark Evaluation and Application in Robotics

In this section, we demonstrate the effectiveness of the algorithms presented in Section 4.2.3 in the context of motor primitive learning for robotics. We will first give a quick overview of how the motor primitives [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007] work and how learning algorithms can be used to adapt them. Subsequently, we will discuss how we can turn the parametrized motor primitives [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007] into explorative, stochastic policies [Rückstieß et al., 2008]. We show that the novel PoWER algorithm outperforms many previous well-known methods, particularly ‘Vanilla’ Policy Gradients [Williams, 1992, Sutton et al., 1999, Lawrence et al., 2003, Peters and Schaal, 2006], Finite Difference Gradients [Sehnke et al., 2010, Peters and Schaal, 2006], the episodic Natural Actor Critic [Peters et al., 2003, 2005], and the generalized Reward-Weighted Regression [Peters and Schaal, 2007] on the two simulated benchmark problems suggested by Peters and Schaal [2006] and the Underactuated Swing-Up [Atkeson, 1994]. We compare policy search based algorithms to a value function based one on the mountain-car benchmark. Additionally, we evaluate policy search methods on the multidimensional robotic tasks Tetherball Target Hitting and Casting. As a significantly more

	<i>Open Parameters</i>	<i>DoF</i>	<i>Rollouts</i>	<i>Policy</i>	<i>Platform</i>	<i>Algorithms</i>
4.3.2	10 (shape)	1	4400	MP	simulation	FDG, VPG, eNAC, eRWR, PoWER
4.3.3	2 (switching)	1	80	bang-bang	simulation	FDG, PoWER, kNN-TD(λ)
4.3.4	6 (positions)	1	200	rhythmic	simulation	FDG, PoWER
4.3.5	10 (goal & shape)	1	200/100	MP	simu/robot	FDG, VPG, eNAC, eRWR, PoWER
4.3.6	10 (shape)	2	200	MP	simulation	eNAC, PoWER
4.3.7	217 (shape)	7	100	MP	robot	PoWER

Table 4.1: Overview of the Experiments: 4.3.2 Basic Motor Learning, 4.3.3 Mountain-Car, 4.3.4 Tetherball Target Hitting, 4.3.5 Underactuated Swing-Up, 4.3.6 Casting, and 4.3.7 Ball-in-a-Cup

complex motor learning task, we will show how the robot can learn a high-speed Ball-in-a-Cup movement [Sumners, 1997] with motor primitives for all seven degrees of freedom of our Barrett WAM robot arm. An overview of the experiments is presented in Table 4.1.

4.3.1 Dynamical Systems Motor Primitives as Stochastic Policies

In the analytically tractable cases, episodic Reinforcement Learning (RL) problems have been studied deeply in the optimal control community. In this field it is well-known that for a finite horizon problem, the optimal solution is non-stationary [Kirk, 1970] and, in general, cannot be represented by a time-independent policy. The motor primitives based on dynamical systems [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007] represent a particular type of time-variant policy that has an internal phase, which corresponds to a clock with additional flexibility (for example, for incorporating coupling effects, perceptual influences, etc.). Thus, they can represent optimal solutions for finite horizons. We embed this internal clock or movement phase into our state and from an optimal control perspective have ensured that the optimal solution can be represented. See Chapter 3 for a more detailed discussion.

One of the biggest advantages of this motor primitive framework [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007] is that the second system, in Equation (3.2), is linear in the policy parameters θ and is therefore well-suited for both imitation learning as well as for the presented reinforcement learning algorithms. For example, if we would have to learn only a motor primitive for a single degree of freedom q_i , then we could use a motor primitive in the form $\ddot{q}_i = \phi(\mathbf{s})^T \theta$ where $\mathbf{s} = [q_i, \dot{q}_i, z]$ is the state and where time is implicitly embedded in z . We use the output of $\ddot{q}_i = \phi(\mathbf{s})^T \theta = \bar{a}$ as the policy mean. The perturbed accelerations $\ddot{q}_i = a = \bar{a} + \varepsilon$ are given to the system.

In Sections 4.3.5 and 4.3.7, we use imitation learning from a single example to generate a sensible initial policy. This step can be performed efficiently in the context of dynamical systems motor primitives as the policy is linear in its parameters, see Section 3.2.4.

4.3.2 Benchmark Comparison I: Basic Motor Learning Examples

As a benchmark comparison, we follow a previously studied scenario in order to evaluate, which method is best-suited for our problem class. We perform our evaluations on exactly the same benchmark problems as in [Peters and Schaal, 2006] and use two tasks commonly studied in motor control literature for which the analytic solutions are known. The first task is a reaching task, wherein a goal has to be reached at a certain time, while the used motor commands have to be minimized. The second task is a reaching task of the same style with an additional via-point. The task is illustrated in Figure 4.1. This comparison mainly shows the suitability of our algorithm (Algorithm 4.4) and that it outperforms previous

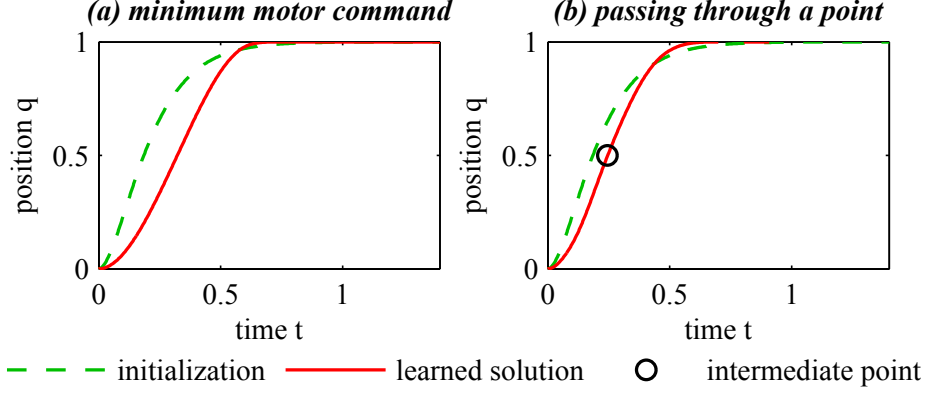


Figure 4.1: This figure shows the initial and the final trajectories for the two motor control tasks. Both start at 0 and have to go to 1 with minimal accelerations. From $T/2 = 0.75$ on the trajectory has to be as close to 1 as possible. For the passing through task the trajectory additionally has to pass through $p_M = 0.5$ at time $M = 7/40T$ indicated by the circle.

methods such as Finite Difference Gradient (FDG) methods [Sehnke et al., 2010, Peters and Schaal, 2006], see Algorithm 4.5, ‘Vanilla’ Policy Gradients (VPG) with optimal baselines [Williams, 1992, Sutton et al., 1999, Lawrence et al., 2003, Peters and Schaal, 2006], see Algorithm 4.1, the episodic Natural Actor Critic (eNAC) [Peters et al., 2003, 2005], see Algorithm 4.2, and the new episodic version of the Reward-Weighted Regression (eRWR) algorithm [Peters and Schaal, 2007], see Algorithm 4.3. MATLAB implementations of all algorithms are available at <http://www.robot-learning.de/Member/JensKober>. For all algorithms except PoWER, we used batches to update the policy. A sliding-window based approach is also possible. For VPG, eNAC, and eRWR a batch size of $H = 2N$ and for FDG a batch size of $H = N + 1$ are typical. For PoWER, we employed an importance sampling based approach, although a batch based update is also possible.

We consider two standard tasks taken from [Peters and Schaal, 2006], but we use the newer form of the motor primitives from [Schaal et al., 2007]. The first task is to achieve a goal with a minimum-squared movement acceleration and a given movement duration, that gives a return of

$$R(\tau) = - \sum_{t=0}^{T/2} c_1 \ddot{q}_t^2 - \sum_{t=T/2+1}^T c_2 \left((q_t - g)^2 + \dot{q}_t^2 \right)$$

for optimization, where $T = 1.5$, $c_1 = 1/100$ is the weight of the transient rewards for the movement duration $T/2$, while $c_2 = 1000$ is the importance of the final reward, extended over the time interval $[T/2 + 1, T]$ which insures that the goal state $g = 1.0$ is reached and maintained properly. The initial state of the motor primitive is always zero in this evaluation.

The second task involves passing through an intermediate point during the trajectory, while minimizing the squared accelerations, that is, we have a similar return with an additional punishment term for missing the intermediate point p_M at time M given by

$$R(\tau) = - \sum_{t=0}^{T/2} \tilde{c}_1 \ddot{q}_t^2 - \sum_{t=T/2+1}^T \tilde{c}_2 \left((q_t - g)^2 + \dot{q}_t^2 \right) - \tilde{c}_3 (q_M - p_M)^2$$

where $\tilde{c}_1 = 1/10000$, $\tilde{c}_2 = 200$, $\tilde{c}_3 = 20000$. The goal is given by $g = 1.0$, the intermediate point a value of $p_M = 0.5$ at time $M = 7/40T$, and the initial state was zero. This return yields a smooth movement, which passes through the intermediate point before reaching the goal, even though the optimal solution is not identical to the analytic solution with hard constraints.

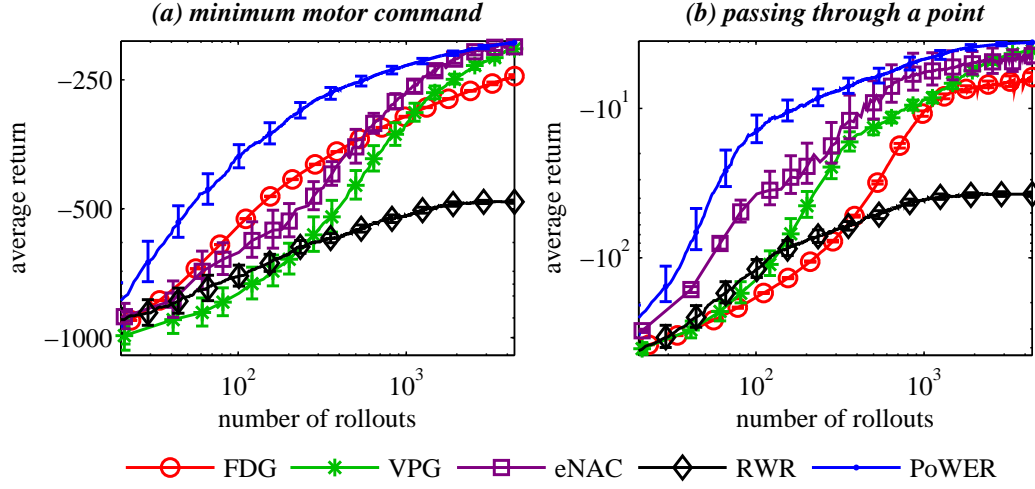


Figure 4.2: This figure shows the mean performance of all compared methods in two benchmark tasks averaged over twenty learning runs with the error bars indicating the standard deviation. Policy learning by Weighting Exploration with the Returns (PoWER) clearly outperforms Finite Difference Gradients (FDG), ‘Vanilla’ Policy Gradients (VPG), the episodic Natural Actor Critic (eNAC), and the adapted Reward-Weighted Regression (eRWR) for both tasks. Note that this plot has logarithmic scales on both axes, thus a unit difference corresponds to an order of magnitude. The omission of the first twenty rollouts was necessary to cope with the log-log presentation.

Algorithm 4.5 Finite Difference Gradients (FDG)

Input: initial policy parameters θ_0

repeat

 Generate policy variations: $\Delta\theta^h \sim \mathcal{U}_{[-\Delta\theta_{\min}, \Delta\theta_{\max}]}$ for $h = \{1, \dots, H\}$ rollouts.

 Sample: Perform $h = \{1, \dots, H\}$ rollouts using $\mathbf{a} = (\theta + \Delta\theta^h)^T \phi(\mathbf{s}, t)$ as policy and collect all $(t, \mathbf{s}_t^h, \mathbf{a}_t^h, \mathbf{s}_{t+1}^h, \mathbf{e}_t^h, r_{t+1}^h)$ for $t = \{1, 2, \dots, T+1\}$.

 Compute: Return $R^h(\theta + \Delta\theta^h) = \sum_{t=1}^{T+1} r_t^h$ from rollouts.

 Compute Gradient:


$$[\mathbf{g}_{\text{FD}}, R_{\text{ref}}]^T = (\Delta\Theta^T \Delta\Theta)^{-1} \Delta\Theta^T \mathbf{R}$$

with $\Delta\Theta = \begin{bmatrix} \Delta\theta^1, \dots, \Delta\theta^H \\ 1, \dots, 1 \end{bmatrix}^T$ and $\mathbf{R} = [R^1, \dots, R^H]^T$.


 Update policy using

$$\theta_{k+1} = \theta_k + \alpha \mathbf{g}_{\text{FD}}.$$

until Convergence $\theta_{k+1} \approx \theta_k$.



(a) The task consists of driving the underpowered car to the target on the mountain indicated by the yellow star.



(b) This figure shows the mean accumulated returns of the methods compared on the mountain car benchmark. The results are averaged over fifty learning runs with error bars indicating the standard deviation. Policy learning by Weighting Exploration with the Returns (PoWER) and Finite Difference Gradients (FDG) clearly outperform $kNN-TD(\lambda)$. All methods converge to the optimal solution.

Figure 4.3: This figure shows an illustration of the mountain-car task and the mean accumulated returns of the compared methods.


All open parameters were manually optimized for each algorithm in order to maximize the performance while not destabilizing the convergence of the learning process. When applied in the episodic scenario, Policy learning by Weighting Exploration with the Returns (PoWER) clearly outperformed the episodic Natural Actor Critic (eNAC), ‘Vanilla’ Policy Gradient (VPG), Finite Difference Gradient (FDG), and the episodic Reward-Weighted Regression (eRWR) for both tasks. The episodic Reward-Weighted Regression (eRWR) is outperformed by all other algorithms suggesting that this algorithm does not generalize well from the immediate reward case. While FDG gets stuck on a plateau, both eNAC and VPG converge to the same good final solution. PoWER finds the a slightly better solution while converging noticeably faster. The results are presented in Figure 4.2.

4.3.3 Benchmark Comparison II: Mountain-Car


As a typical reinforcement learning benchmark we chose the mountain-car task [Sutton and Barto, 1998] as it can be treated with episodic reinforcement learning. In this problem we have a car placed in a valley, and it is supposed to go on the top of the mountain in front of it, but does not have the necessary capabilities of acceleration to do so directly. Thus, the car has to first drive up the mountain on the opposite side of the valley to gain sufficient energy. The dynamics are given in [Sutton and Barto, 1998] as

$$\begin{aligned}\dot{x}_{t+1} &= \dot{x}_t + 0.001a_t - 0.0025 \cos(3x_t), \\ x_{t+1} &= x_t + \dot{x}_{t+1},\end{aligned}$$

with position $-1.2 \leq x_{t+1} \leq 0.5$ and velocity $-0.07 \leq \dot{x}_{t+1} \leq 0.07$. If the goal $x_{t+1} \geq 0.5$ is reached the episode is terminated. If the left bound is reached the velocity is reset to zero. The initial condition of the car is $x_0 = -0.5$ and $\dot{x}_0 = 0$. The reward is $r_t = -1$ for all time-steps until the car reaches the goal. We employed an undiscounted return. The set of actions a_t is slightly different to the setup proposed by Sutton and Barto [1998]. We only have two actions, the full throttle forward ($a_t = +1$) and the full



(a) The task consists of striking the yellow ball hanging on an elastic string such that it hits the green target.



(b) The returns are averaged over 20 learning runs with error bars indicating the standard deviation. Policy learning by Weighting Exploration with the Returns (PoWER) clearly outperforms Finite Difference Gradients (FDG).

Figure 4.4: This figure shows an illustration of the Tetherball Target Hitting task and the mean returns of the compared methods.

throttle reverse ($a_t = -1$). From a classical optimal control point of view, it is straightforward to see that a bang-bang controller can solve this problem. As an initial policy we chose a policy that accelerates forward until the car cannot climb the mountain further, accelerates reverse until the car cannot climb the opposite slope further, and finally accelerates forward until the car reaches the goal. This policy reaches the goal but is not optimal as the car can still accumulate enough energy if it reverses the direction slightly earlier. As a parametrization for the policy search approaches we chose to encode the switching points of the acceleration. The two parameters of the policy indicate at which timestep t the acceleration is reversed. For this kind of policy only algorithms that perturb the parameters are applicable and we compare a Finite Difference Gradient approach to PoWER. This parametrized policy is entirely different to motor primitives. Additionally we included a comparison to a value function based method. The Q-function was initialized with our initial policy. As the k NN-TD(λ) algorithm [Martín H. et al., 2009] won the Reinforcement Learning Competitions in 2008 and 2009, we selected it for this comparison. This comparison is contrived as our switching policy always starts in a similar initial state while the value function based policy can start in a wider range of states. Furthermore, the policy search approaches may be sped up by the initialization, while k NN-TD(λ) will learn the optimal policy without prior knowledge and does not benefit much from the initialization. However, the use of a global approach, such as k NN-TD(λ) requires a global search of the state space. Such a global search limits the scalability of these approaches. The more local approaches of policy search are less affected by these scalability problems. Figure 4.3b shows the performance of these algorithms. As k NN-TD(λ) initially explores the unseen parts of the Q-function, the policy search approaches converge faster. All methods find the optimal solution.

4.3.4 Benchmark Comparison III: Tetherball Target Hitting

In this task, a table tennis ball is hanging on an elastic string from the ceiling. The task consists of hitting the ball with a table tennis racket so that it hits a fixed target. The task is illustrated in Figure 4.4a. The return is based on the minimum distance between the ball and the target during one episode transformed by an exponential. The policy is parametrized as the position of the six lower degrees of freedom of the Barrett WAM. Only the first degree of freedom (shoulder rotation) is moved during an episode. The movement is represented by a rhythmic policy with a fixed amplitude and period. Due to the




Figure 4.5: This figure shows the time series of the Underactuated Swing-Up where only a single joint of the robot is moved with a torque limit ensured by limiting the maximal motor current of that joint. The resulting motion requires the robot to (ii) first move away from the target to limit the maximal required torque during the swing-up in (iii-v) and subsequent stabilization (vi).




Figure 4.6: This figure shows the performance of all compared methods for the swing-up in simulation and the mean performance averaged over 20 learning runs with the error bars indicating the standard deviation. PoWER outperforms the other algorithms from 50 rollouts on and finds a significantly better policy.

parametrization of the task only PoWER and Finite Difference Gradients are applicable. We observed reliable performance if the initial policy did not miss the target by more than approximately 50cm. In this experiment it took significantly less iterations to find a good initial policy than to tune the learning rate of Finite Difference Gradients, a problem from which PoWER did not suffer as it is an EM-like algorithm. Figure 4.4b illustrates the results. PoWER converges significantly faster.

4.3.5 Benchmark Comparison IV: Underactuated Swing-Up

As an additional simulated benchmark and for the real-robot evaluations, we employed the Underactuated Swing-Up [Atkeson, 1994]. Here, only a single degree of freedom is represented by the motor primitive as described in Section 4.3.1. The goal is to move a hanging heavy pendulum to an upright position and to stabilize it there. The objective is threefold: the pendulum has to be swung up in a minimal amount of time, has to be stabilized in the upright position, and achieve these goals with minimal motor torques. By limiting the motor current for this degree of freedom, we can ensure that the torque limits described in [Atkeson, 1994] are maintained and directly moving the joint to the right position is not possible. Under these torque limits, the robot needs to first move away from the target to reduce the maximal required torque during the swing-up, see Figure 4.5. This problem is similar to the mountain-car problem (Section 4.3.3). The standard mountain-car problem is designed to get the car to the top of the mountain in minimum time. It does not matter if it stops at this point or drives at a high speed as usage of the accelerator and brake is not punished. Adding the requirement of stabilizing the car at the top of the




Figure 4.7: This figure shows the improvement of the policy over rollouts. The snapshots from the video show the final positions. (0) Initial policy after imitation learning (without torque limit). (1) Initial policy after imitation learning (with active torque limit). (20) Policy after 20 rollouts, going further up. (30) Policy after 30 rollouts, going too far. (40) Policy after 40 rollouts, going only a bit too far. (65) Final policy after 65 rollouts.





Figure 4.8: This figure shows the performance of all compared methods for the swing-up on the real robot and the mean performance averaged over 3 learning runs with the error bars indicating the standard deviation. PoWER outperforms the other algorithms and finds a significantly better policy.


mountain makes the problem significantly harder. These additional constraints exist in the Underactuated Swing-Up task where it is required that the pendulum (the equivalent of the car) stops at the top to fulfill the task. The applied torque limits were the same as in [Atkeson, 1994] and so was the reward function, except that the complete return of the trajectory was transformed by an $\exp(\cdot)$ to ensure positivity. The reward function is given by

$$r(t) = -c_1 q(t)^2 + c_2 \log \cos \left(c_3 \frac{u(t)}{u_{\max}} \right),$$

where the constants are $c_1 = 5/\pi^2 \approx 0.507$, $c_2 = (2/\pi)^2 \approx 0.405$, and $c_3 = \pi/2 \approx 1.571$. Please note that π refers to the mathematics constant here, and not to the policy. The first term of the sum is punishing the distance to the desired upright position $q = 0$, and the second term is punishing the usage of motor torques u . A different trade-off can be achieved by changing the parameters or the structure of the reward function, as long as it remains an improper probability function. Again all open parameters of all algorithms were manually optimized. The motor primitive with nine shape parameters and one goal parameter was initialized by imitation learning from a kinesthetic teach-in. Kinesthetic teach-in means “taking the robot by the hand”, performing the task by moving the robot while it is in gravity-compensation mode, and recording the joint angles, velocities, and accelerations. This initial demonstration needs to include all the relevant features of the movement, e.g., it should first move away from the target and then towards the upright position. The performance of the algorithms is very robust, as long as the initial policy with active torque limits moves the pendulum approximately above the horizontal orientation.



(a) The task consists of placing the blue ball in the brown cup.



(b) The returns are averaged over 20 learning runs with error bars indicating the standard deviation. Policy learning by Weighting Exploration with the Returns (PoWER) clearly outperforms episodic Natural Actor Critic (eNAC).

Figure 4.9: This figure illustrates the Casting task and shows the mean returns of the compared methods.

As the batch size, and, thus the learning speed, of the gradient based approaches depend on the number of parameters (see Section 4.3.2), we tried to minimize the number of parameters. Using more parameters would allow us to control more details of the policy which could result in a better final policy, but would have significantly slowed down convergence. At least nine shape parameters were needed to ensure that the imitation can capture the movement away from the target, which is essential to accomplish the task. We compared all algorithms considered in Section 4.3.2 and could show that PoWER would again outperform the others. The convergence is a lot faster than in the basic motor learning examples (see Section 4.3.2), as we do not start from scratch but rather from an initial solution that allows significant improvements in each step for each algorithm. The results are presented in Figure 4.6. See Figure 4.7 and Figure 4.8 for the resulting real-robot performance.

4.3.6 Benchmark Comparison V: Casting

In this task a ball is attached to the endeffector of the Barrett WAM by a string. The task is to place the ball into a small cup in front of the robot. The task is illustrated in Figure 4.9a. The return is based on the sum of the minimum distance between the ball and the top, the center, and the bottom of the cup respectively during one episode. Using only a single distance, the return could be successfully optimized, but the final behavior often corresponded to a local maximum (for example hitting the cup from the side). The movement is in a plane and only one shoulder DoF and the elbow are moved. The policy is parametrized using motor primitives with five shape parameters per active degree of freedom. The policy is initialized with a movement that results in hitting the cup from the side in the upper quarter of the cup. If the ball hits the cup below the middle, approximately 300 rollouts were required for PoWER and we did not achieve reliable performance for the episodic Natural Actor Critic. We compare the two best performing algorithms from the basic motor learning examples (see Section 4.3.2) and the Underactuated Swing-Up (see Section 4.3.5), namely eNAC and PoWER. Figure 4.9b illustrates the results. PoWER again converges significantly faster.




Figure 4.10: This figure illustrates how the reward is calculated. The plane represents the level of the upper rim of the cup. For a successful rollout the ball has to be moved above the cup first and is then flying in a downward direction into the opening of the cup. The reward is calculated as the distance d of the center of the cup and the center of the ball on the plane at the moment the ball is passing the plane in a downward direction. If the ball is flying directly into the center of the cup, the distance is 0 and through the transformation $\exp(-d^2)$ yields the highest possible reward of 1. The further the ball passes the plane from the cup, the larger the distance and thus the smaller the resulting reward.

4.3.7 Ball-in-a-Cup on a Barrett WAM

The children’s motor skill game Ball-in-a-Cup [Sumners, 1997], also known as Balero, Bilboquet, and Kendama, is challenging even for adults. The toy has a small cup which is held in one hand (or, in our case, is attached to the end-effector of the robot) and the cup has a small ball hanging down on a string (the string has a length of 40cm in our setup). Initially, the ball is hanging down vertically in a rest position. The player needs to move fast in order to induce a motion in the ball through the string, toss it up, and catch it with the cup. A possible movement is illustrated in Figure 4.11 in the top row.

Note that learning of Ball-in-a-Cup and Kendama has previously been studied in robotics. We are going to contrast a few of the approaches here. While we learn directly in the joint space of the robot, Takenaka [1984] recorded planar human cup movements and determined the required joint movements for a planar, three degree of freedom (DoF) robot, so that it could follow the trajectories while visual feedback was used for error compensation. Both Sato et al. [1993] and Shone et al. [2000] used motion planning approaches which relied on very accurate models of the ball and the string while employing only one DoF in [Shone et al., 2000] or two DoF in [Sato et al., 1993] so that the complete state-space could be searched exhaustively. Interestingly, exploratory robot moves were used in [Sato et al., 1993] to estimate the parameters of the employed model. Probably the most advanced preceding work on learning Kendama was carried out by Miyamoto et al. [1996] who used a seven DoF anthropomorphic arm and recorded human motions to train a neural network to reconstruct via-points. Employing full kinematic knowledge, the authors optimize a desired trajectory.

The state of the system can be described by joint angles and joint velocities of the robot as well as the the Cartesian coordinates and velocities of the ball. The actions are the joint space accelerations where each of the seven joints is driven by a separate motor primitive, but with one common canonical system. The movement uses all seven degrees of freedom and is not in a plane. All motor primitives are perturbed separately but employ the same joint final reward given by

$$r(t) = \begin{cases} \exp(-\alpha(x_c - x_b)^2 - \alpha(y_c - y_b)^2) & \text{if } t = t_c, \\ 0 & \text{otherwise,} \end{cases}$$

where we denote the moment when the ball passes the rim of the cup with a downward direction by t_c , the cup position by $[x_c, y_c, z_c] \in \mathbb{R}^3$, the ball position by $[x_b, y_b, z_b] \in \mathbb{R}^3$, and a scaling parameter by $\alpha = 100$ (see also Figure 4.10). The algorithm is robust to changes of this parameter as long as the reward clearly discriminates good and suboptimal trials. The directional information is necessary as the algorithm

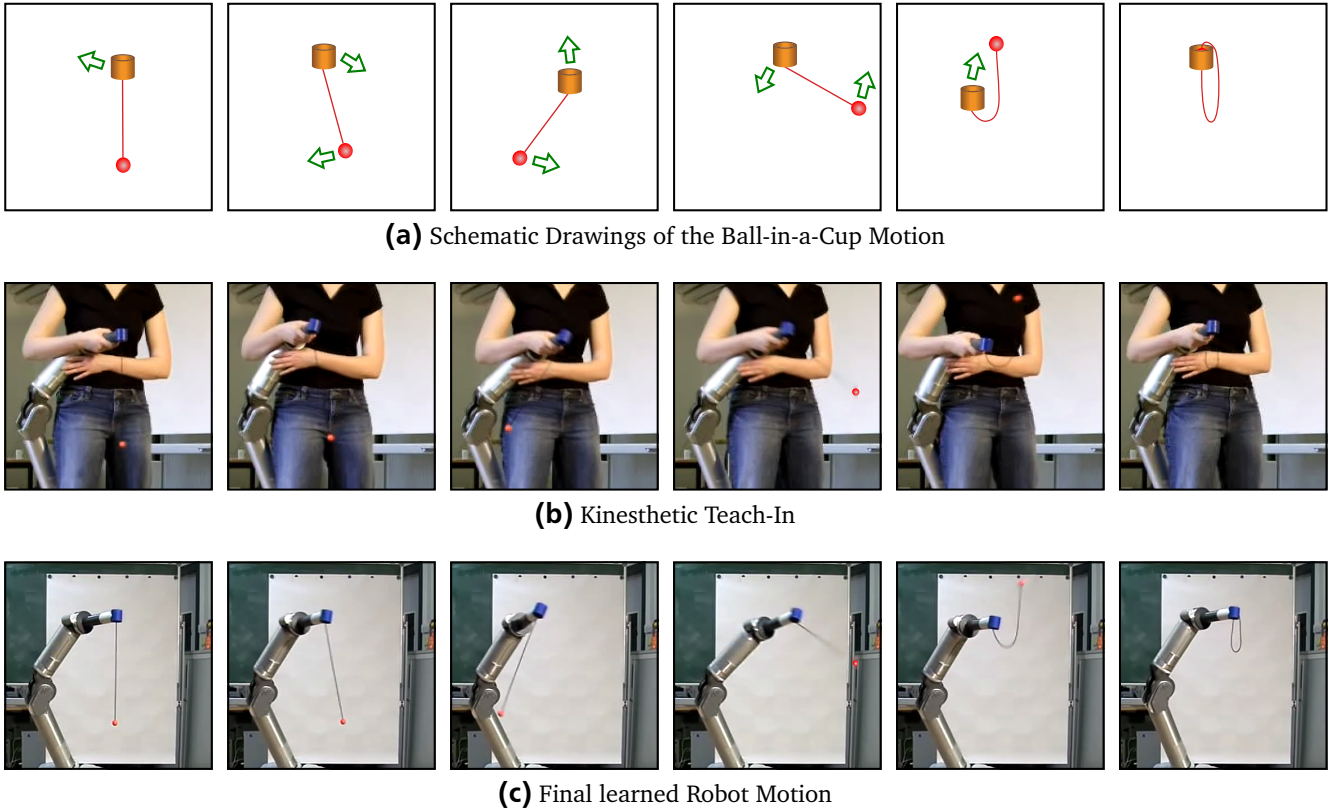


Figure 4.11: This figure shows schematic drawings of the Ball-in-a-Cup motion (a), the final learned robot motion (c), as well as a kinesthetic teach-in (b). The arrows show the directions of the current movements in that frame. The human cup motion was taught to the robot by imitation learning with 31 parameters per joint for an approximately 3 seconds long trajectory. The robot manages to reproduce the imitated motion quite accurately, but the ball misses the cup by several centimeters. After approximately 75 iterations of our Policy learning by Weighting Exploration with the Returns (PoWER) algorithm the robot has improved its motion so that the ball goes in the cup. Also see Figure 4.12.




Figure 4.12: This figure shows the expected return of the learned policy in the Ball-in-a-Cup evaluation averaged over 20 runs.

could otherwise learn to hit the bottom of the cup with the ball. This solution would correspond to a local maximum whose reward is very close to the optimal one, but the policy very far from the optimal one. The reward needs to include a term avoiding this local maximum. PoWER is based on the idea of considering the reward as an improper probability distribution. Transforming the reward using the exponential enforces this constraint. The reward is not only affected by the movements of the cup but foremost by the movements of the ball, which are sensitive to small changes in the cup’s movement. A small perturbation of the initial condition or during the trajectory can change the movement of the ball and hence the outcome of the complete movement. The position of the ball is estimated using a stereo vision system and is needed to determine the reward.

Due to the complexity of the task, Ball-in-a-Cup is a hard motor learning task for children, who usually only succeed at it by observing another person playing combined with a lot of improvement by trial-and-error. Mimicking how children learn to play Ball-in-a-Cup, we first initialize the motor primitives by imitation learning and, subsequently, improve them by reinforcement learning. We recorded the motions of a human player by kinesthetic teach-in to obtain an example for imitation as shown in Figure 4.11b. A single demonstration was used for imitation learning. Learning from multiple demonstrations did not improve the performance as the task is sensitive to small differences. As expected, the robot fails to reproduce the presented behavior even if we use all the recorded details for the imitation. Thus, reinforcement learning is needed for self-improvement. The more parameters used for the learning, the slower the convergence is. Due to the imitation, the ball must go above the rim of the cup such that the algorithm gets at least a small positive reward for all rollouts. This way exhaustive exploration is avoided as the algorithm can compare the performance of the different rollouts. We determined that 31 shape-parameters per motor primitive are needed. With less parameters the ball does not go above the rim of the cup in the initial trial and the algorithm does not receive any meaningful information about the trial using the aforementioned reward function. More shape-parameters will lead to a more accurate reproduction of the demonstrated movement and, thus, to a better initial policy. However, there is a trade-off between this better initial policy and a potentially lower learning speed. Using three times as many parameters the algorithm converged at roughly the same time. The hyper-parameters σ_{ij} are initially set in the same order of magnitude as the median of the parameters for each motor primitive and are then optimized alongside the shape-parameters by PoWER. The performance of the algorithm is fairly robust for values chosen in this range. Figure 4.12 shows the expected return over the number of rollouts where convergence to a maximum is clearly recognizable. The robot regularly succeeds at bringing the ball into the cup after approximately 75 iterations. Figure 4.13 shows the improvement of the policy over the rollouts. From our experience, nine year old children get the ball in the cup for the first time after about 35 trials while the robot gets the ball in for the first time after 42-45 rollouts. However, after




Figure 4.13: This figure shows the improvement of the policy over rollouts. The snapshots from the video show the position of the ball closest to the cup during a rollout. (1) Initial policy after imitation learning. (15) Policy after 15 rollouts, already closer. (25) Policy after 25 rollouts, going too far. (45) Policy after 45 rollouts, hitting the near rim. (60) Policy after 60 rollouts, hitting the far rim. (100) Final policy after 100 rollouts.

100 trials, the robot exhibits perfect runs in every single trial while children do not have a comparable success rate. Of course, such a comparison with children is contrived as a robot can precisely reproduce movements unlike any human being, and children can most likely adapt faster to changes in the setup.

4.4 Discussion & Conclusion

In Section 4.4.1, we will discuss robotics as a benchmark for reinforcement learning, in Section 4.4.2 we discuss different simulation to robot transfer scenarios, and we will draw our conclusions in Section 4.4.3.

4.4.1 Discussion: Robotics as Benchmark for Reinforcement Learning?

Most reinforcement learning algorithms are evaluated on synthetic benchmarks, often involving discrete states and actions. Agents in many simple grid worlds take millions of actions and episodes before convergence. As a result, many methods have focused too strongly on such problem domains. In contrast, many real world problems such as robotics are best represented with high-dimensional, continuous states and actions. Every single trial run is costly and as a result such applications force us to focus on problems that will often be overlooked accidentally in synthetic examples. Simulations are a helpful testbed for debugging algorithms. Continuous states and actions as well as noise can be simulated, however simulations pose the temptation of using unrealistically many iterations and also allow us to exploit the perfect models.

Our typical approach consists of testing the algorithm in a simulation of the robot and the environment. Once the performance is satisfactory we replace either the robot or the environment with its real counterpart depending on the potential hazards. Replacing the robot is always possible as we can still simulate the environment taking into account the state of the real robot. Learning with a simulated robot in the real environment is not always a possibility especially if the robot influences the observed environment, such as in the Ball-in-a-Cup task. The final evaluations are done on the real robot in the real environment.

Our experience in robotics show that the plant and the environment can often not be represented accurately enough by a physical model and that the learned policies are thus not entirely transferable. If sufficiently accurate models were available, we could resort to the large body of work on optimal control [Kirk, 1970], which offers alternatives to data driven reinforcement learning. However, when a model with large errors is used, the solution suffers severely from an optimization bias as already experienced by Atkeson [1994]. Here, the reinforcement learning algorithm exploits the imperfections of the simulator rather than yielding an optimal policy.

None of our learned policies could be transferred from a simulator to the real system without changes despite that the immediate errors of the simulator have been smaller than the measurement noise. Methods which jointly learn the models and the policy as well as perform some of the evaluations and

updates in simulation (such as Dyna-like architectures as in [Sutton, 1990]) may alleviate this problem. In theory, a simulation could also be used to eliminate obviously bad solutions quickly. However, the differences between the simulation and the real robot do accumulate over time and this approach is only feasible if the horizon is short or the simulation very accurate. Priming the learning process by imitation learning and optimizing in its vicinity achieves the desired effect better and has thus been employed in this chapter.

Parametrized policies greatly reduce the need of samples and evaluations. Choosing an appropriate representation like motor primitives renders the learning process even more efficient.


One major problem with robotics benchmarks is the repeatability of the experiments. The results are tied to the specific hardware and setup used. Even a comparison with simulated robots is often not possible as many groups rely on proprietary simulators that are not freely available. Most of the benchmarks presented in this chapter rely on initial demonstrations to speed up the learning process. These demonstrations play an important part in the learning process. However, these initial demonstrations are also tied to the hardware or the simulator and are, thus, not straightforward to share. Comparing new algorithms to results from different authors usually requires the reimplementation of their algorithms to have a fair comparison.

Reproducibility is a key requirement for benchmarking but also a major challenge for robot reinforcement learning. To overcome this problem there are two orthogonal approaches: (i) a central organizer provides an identical setup for all participants and (ii) all participants share the same setup in a benchmarking lab. The first approach has been majorly pushed by funding agencies in the USA and Europe. In the USA, there have been special programs on robot learning such as DARPA Learning Locomotion (L2), Learning Applied to Ground Robotics (LAGR) and the DARPA Autonomous Robot Manipulation (ARM) [DARPA, 2010c,b,a]. However, the hurdles involved in getting these robots to work have limited the participation to strong robotics groups instead of opening the robot reinforcement learning domain to more machine learning researchers. Alternative ways of providing identical setups are low cost standardized hardware or a system composed purely of commercially available components. The first suffers from reproducibility and reliability issues while the latter results in significant system integration problems. Hence, it may be more suitable for a robot reinforcement learning challenge to be hosted by a robot learning group with significant experience in both domains. The host lab specifies tasks that they have been able to accomplish on a real robot system. The hosts also need to devise a virtual robot laboratory for allowing the challenge participants to program, test and debug their algorithms. To limit the workload and the risk of the organizers, a first benchmarking round would be conducted using this simulated setup to determine the promising approaches. Successful participants will be invited by the host lab in order to test these algorithms in learning on the real system where the host lab needs to provide significant guidance. To our knowledge, no benchmark based on this approach has been proposed yet. We are currently evaluating possibilities to organize a challenge using such a shared setup in the context of the PASCAL2 Challenge Program [PASCAL2, 2010].


To successfully apply reinforcement learning to robotics, a fair level of knowledge on the limitations and maintenance of the robot hardware is necessary. These limitations include feasible actions, feasible run-time, as well as measurement delay and noise. Cooperation with a strong robotics group is still extremely important in order to apply novel reinforcement learning methods in robotics.

4.4.2 Discussion on Simulation to Robot Transfer Scenarios

In this chapter, we have discussed reinforcement learning for real robots with highly dynamic tasks. The opposite extreme in robotics would be, for example, a maze navigation problem where a mobile robot that has macro-actions such as “go left” and the lower level control moves the robot exactly by a well-defined distance to the left. In this scenario, it would probably be easier to transfer simulation results to real systems. For highly dynamic tasks or environments, accurate simulations are generically difficult. Besides fast moving objects and many interacting objects as well as deformable objects (often called soft bodies),



(a) Reinforcement learning required unrealistically many trials in simulation.



(b) Imitation learning only was sufficient on the real robot.

Figure 4.14: This figure illustrates the Ball-Paddling task in simulation and on the real robot. The difference between simulation and robotics can be particularly emphasized in this problem where unrealistically many trials were needed on the simulation for reinforcement learning while the real world behavior could be learned by imitation learning. It illustrates the energy-consuming scenario and the difficulties of realistic learning in the presence of contact forces.

like cloth, string, fluids, hydraulic tubes and other elastic materials are hard to simulate reliably and, thus, have an enormous impact on transferability. Additionally, the level and quality of measurement noise has a direct implication on the difficulty and the transferability of the learning task.

Better simulations often alleviate some of these problems. However, there is always a trade-off as more detailed simulations also require more precise model identification, higher temporal resolution, and, frequently even finite elements based simulations. Such detailed simulations may even be much slower than real-time, thus defeating one of the major advantages of simulations.

Aside from these clear difficulties in creating simulations that allow the transfer to real systems, we have observed three critically different scenarios for reinforcement learning in robotics. These scenarios are characterized by the energy flow between the policy and the system. In the energy-absorbing scenario, the task has passive dynamics and, hence, it is safer and easier to learn on a real robot. We are going to discuss the examples of Ball-Paddling, foothold selection in legged locomotion, and grasping (see Section 4.4.2). The second scenario has a border-line behavior: the system conserves most of the energy but the policy also only needs to inject energy into the system for a limited time. We will discuss Ball-in-a-Cup, Tetherball Target Hitting, and Mountain-Car as examples for this scenario (see Section 4.4.2). In the energy-emitting scenario energy is inserted due to the system dynamics even if the policy does not transfer energy into the system. The classical examples are Cart-Pole and inverted helicopters, and we also have the Underactuated Swing-Up which has to stabilize at the top (see Section 4.4.2). These different scenarios have implications on the relative utility of simulations and real robots.

As we are discussing our experience in performing such experiments, it may at times appear anecdotal. We hope the reader benefits from our insights nevertheless. However, the resulting classification bears similarities with insights on control law derivation [Fantoni and Lozano, 2001].

Energy-Absorbing Scenario

In this scenario, the system absorbs energy from the actions. As shown in Figure 4.14, we learned a Ball-Paddling task where a ping-pong ball is attached to a paddle by an elastic string and the ball has to be kept in the air by repeatedly hitting it from below. In this setup, the elastic string pulls the ball back towards the center of the paddle and the contact forces between the ball and the paddle are very

complex. We modeled the system in as much detail as possible, including friction models, restitution models, dampening models, models for the elastic string, and air drag. However, in simulation the paddling behavior was still highly unpredictable and we needed a few thousand iterations to learn an optimal frequency, amplitude, and movement shape. The number of simulated trials exceeded the feasible amount on a real system. In contrast, when learning on the real system, we obtained a stable paddling behavior by imitation learning using the initial demonstration only and no further reinforcement learning was needed.

In general, scenarios with complex contact forces often work better in a real-world experiment. This problem was particularly drastic in locomotion experiments on rough terrain where the real world was an easier learning environment due to favorable friction properties during foot contact [Peters, 2007]. In this experiment, learning was significantly harder in simulation and the learned policy could not be transferred. The same effect occurs in grasping when objects often cannot be grasped in simulation due to slip but the real world friction still allows them to be picked up. Hence, in this scenario, policies from simulations are frequently not helpful and learning in simulation is harder than on the real system. The results only transfer in a few cases. A simulation is therefore only recommended as a feasibility study and for software debugging. As most contacts differ significantly due to the current properties (which vary with time and temperature) of the two interacting objects, only a learned simulator is likely to grasp all relevant details.

Border-Line Scenario

In this scenario, adding too much energy to a system does not necessarily lead to a negative outcome. For example, in the Mountain-Car problem (see Section 4.3.3), inserting more energy and driving through the goal at a higher velocity does not affect task achievement. In contrast not inserting enough energy will result in a failure as the car cannot reach the top of the mountain. The Tetherball Target Hitting application presented in Section 4.3.4 exhibits a very similar behavior. The Ball-in-a-Cup experiment (see Section 4.3.7) highlights the resulting similarities between learning in good simulations and the real world for this scenario. Success is possible if more energy is inserted and the ball flies higher. However, when using too little energy the ball will stay below the opening of the cup. In this favorable scenario the “classical” strategy can be applied: learn how to learn in simulation. The policy learned in simulation does not necessarily transfer to the real world and the real-world scenario can be highly different but the learning speed and behavior are similar. Hence, hyper-parameters such as learning and exploration rates can be tuned in simulation. The learning algorithm may take longer due to increased errors, modeling problems and uncertainty. Still, good practice is to create a sufficiently accurate simulator and to adapt the learning strategy subsequently to the real system.

Energy-Emitting Scenario

Energy emission causes very different problems. Uncertainty in states will cause overreactions, hence, drastic failures are likely to occur when the system becomes unstable in a control theory sense. This system excitability often makes the task significantly harder to learn on a real robot in comparison to a simulated setup. Here, pre-studies in simulations are a necessary but not sufficient condition. Due to unmodeled nonlinearities, the exploration will affect various algorithms differently. Classical examples are helicopters in inverted flight [Ng et al., 2004b] and the pendulum in a Cart-Pole task in an upright position [Sutton and Barto, 1998] as these have to be constantly stabilized. Additionally the pendulum in the Swing-Up has to be stabilized in the final position or it will fall over and cause a failure. In this chapter, we take the example of the Swing-Up to illustrate how some methods unexpectedly do better in the real world as exhibited by Figures 4.6 and 4.8. The learning progress of all algorithms is noisier and the eRWR performs better on the real robot. The form of exploration employed by PoWER seems to give it an additional advantage in the first 20 rollouts as direct exploration on the actions is partially

obscured by measurement noise. In order to cope with differences to the real-world, simulations need to be more stochastic than the real system (as suggested by Ng et al. [2004b]) and should be learned to make transferring the results easier (as for example in [Schaal et al., 2002]).

4.4.3 Conclusion

In this chapter, we have presented a framework for deriving several policy learning methods that are applicable in robotics and an application to a highly complex motor learning task on a real Barrett WAM robot arm. We have shown that policy gradient methods are a special case of this framework. During initial experiments, we realized that the form of exploration highly influences the speed of the policy learning method. This empirical insight resulted in a novel policy learning algorithm, Policy learning by Weighting Exploration with the Returns (PoWER), an EM-inspired algorithm that outperforms several other policy search methods both on standard benchmarks as well as on a simulated Underactuated Swing-Up.

We have successfully applied this novel PoWER algorithm in the context of learning two tasks on a physical robot, namely the Underacted Swing-Up and Ball-in-a-Cup. Due to the curse of dimensionality, we cannot start with an arbitrary solution. Instead, we mimic the way children learn Ball-in-a-Cup and first present an example movement for imitation learning, which is recorded using kinesthetic teach-in. Subsequently, our reinforcement learning algorithm learns how to move the ball into the cup reliably. After only realistically few episodes, the task can be regularly fulfilled and the robot shows very good average performance. After 100 rollouts, the hyper-parameters, such as the exploration rate, have converged to negligible size and do not influence the outcome of the behavior any further. The experiments in this chapter use the original variant of the motor primitives which cannot deal with large perturbations. However, the extended variable-feedback variant presented in [Kober et al., 2008] can deal with a variety of changes directly (for example, in the length of the string or the size or weight of the ball) while the approach presented in this chapter will recover quickly by learning an adjusted policy in a few roll-outs. In [Kober et al., 2008], we have also shown that learning a strategy of pulling the ball up and moving the cup under the ball (as in Kendama) is possible in approximately the same number of trials. We have discovered a variety of different human strategies for Ball-in-a-Cup in movement recordings, see [Chiappa et al., 2009].

Our approach has already given rise to follow-up work in other contexts, for example, [Vlassis et al., 2009, Kormushev et al., 2010]. Theodorou et al. [2010] have shown that an algorithm very similar to PoWER can also be derived from a completely different perspective, that is, the path integral approach.

4.A Derivations

In this appendix, we provide the derivations of various algorithms in more details than in the main text. We first present, how the episodic REINFORCE [Williams, 1992] can be obtained (Section 4.A.1). Subsequently, we show how the episodic Reward Weighted Regression (eRWR) [Peters and Schaal, 2007] can be generalized for the episodic case (Section 4.A.2), and finally we derive EM Policy learning by Weighting Exploration with the Returns (PoWER) and show a number of simplifications (Section 4.A.3).

4.A.1 REINFORCE

If we choose a stochastic policy in the form $a = \theta^T \phi(s, t) + \varepsilon_t$ with $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$, we have

$$\pi(a_t | s_t, t) = \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2} (a - \theta^T \phi)^2\right),$$

and, thus,

$$\partial_{\theta} \log \pi = \sigma^{-2} (a - \theta^T \phi) \phi^T.$$

Therefore, the gradient, in Equation (4.2), becomes

$$\partial_{\theta'} L_{\theta}(\theta') = E \left\{ \sum_{t=1}^T \sigma^{-2} (a - \theta'^T \phi) \phi^T R \right\} = E \left\{ \sum_{t=1}^T \sigma^{-2} \varepsilon_t \phi^T R \right\}, \quad (4.5)$$

which corresponds to the episodic REINFORCE algorithm [Williams, 1992].

4.A.2 Episodic Reward Weighted Regression (eRWR)

Setting Equation (4.5) to zero

$$\partial_{\theta'} L_{\theta}(\theta') = E \left\{ \sum_{t=1}^T \sigma^{-2} (a - \theta'^T \phi) \phi^T R \right\} = 0,$$

we obtain

$$E \left\{ \sum_{t=1}^T \sigma^{-2} a R \phi^T \right\} = E \left\{ \sum_{t=1}^T \sigma^{-2} (\theta'^T \phi) R \phi^T \right\}.$$

Since σ is constant, we have $E \left\{ \sum_{t=1}^T a R \phi^T \right\} = \theta'^T E \left\{ \sum_{t=1}^T \phi R \phi^T \right\}$. The θ' minimizing the least squares error can be found by locally weighted linear regression (R as weights and ϕ as basis functions) considering each time-step and rollout separately

$$\theta' = (\Phi^T \mathbf{R} \Phi)^{-1} \Phi^T \mathbf{R} \mathbf{A},$$

with $\Phi = [\phi_1^1, \dots, \phi_T^1, \phi_1^2, \dots, \phi_T^2, \dots, \phi_1^H, \dots, \phi_T^H]^T$, $\mathbf{R} = \text{diag}(R^1, \dots, R^1, R^2, \dots, R^2, \dots, R^H, \dots, R^H)$, and $\mathbf{A} = [a_1^1, \dots, a_T^1, a_1^2, \dots, a_T^2, \dots, a_1^H, \dots, a_T^H]^T$ for H rollouts.

The same derivation holds if we use Equation (4.3) instead of Equation (4.2). Then \mathbf{R} in the regression is replaced by $\mathbf{Q}^{\pi} = \text{diag}(Q_1^{\pi,1}, \dots, Q_T^{\pi,1}, Q_1^{\pi,2}, \dots, Q_T^{\pi,2}, \dots, Q_1^{\pi,H}, \dots, Q_T^{\pi,H})$. Using the state-action value function Q^{π} yields slightly faster convergence than using the return R .

4.A.3 EM Policy learning by Weighting Exploration with the Returns (PoWER)

If we choose a stochastic policy in the form $a = (\theta + \varepsilon_t)^T \phi(s, t)$ with $\varepsilon_t \sim \mathcal{N}(\mathbf{0}, \hat{\Sigma})$, we have

$$\pi(a_t | s_t, t) = \mathcal{N}(a | \theta^T \phi(s, t), \phi(s, t)^T \hat{\Sigma} \phi(s, t)) = (2\pi \phi^T \hat{\Sigma} \phi)^{-1/2} \exp \left(-\frac{(a - \theta^T \phi)^2}{2\phi^T \hat{\Sigma} \phi} \right),$$

and, thus, $\partial_{\theta} \log \pi = (a - \theta^T \phi) \phi^T / (\phi^T \hat{\Sigma} \phi)^2$. Therefore Equation (4.3) becomes

$$\partial_{\theta'} L_{\theta}(\theta') = E \left\{ \sum_{t=1}^T \frac{(a - \theta'^T \phi) \phi^T}{\phi^T \hat{\Sigma} \phi} Q^{\pi} \right\}.$$

Setting this equation to zero is equivalent to

$$E \left\{ \sum_{t=1}^T \frac{a \phi^T}{\phi^T \hat{\Sigma} \phi} Q^{\pi} \right\} \equiv E \left\{ \sum_{t=1}^T \frac{((\theta + \varepsilon_t)^T \phi) \phi^T}{\phi^T \hat{\Sigma} \phi} Q^{\pi} \right\} = E \left\{ \sum_{t=1}^T \frac{(\theta^T \phi) \phi^T}{\phi^T \hat{\Sigma} \phi} Q^{\pi} \right\}.$$

This equation yields

$$\begin{aligned}\theta^T &= E \left\{ \sum_{t=1}^T \frac{((\theta + \varepsilon_t)^T \phi) \phi^T}{\phi^T \hat{\Sigma} \phi} Q^\pi \right\} E \left\{ \sum_{t=1}^T \frac{\phi \phi^T}{\phi^T \hat{\Sigma} \phi} Q^\pi \right\}^{-1} \\ &= \theta^T + E \left\{ \sum_{t=1}^T \frac{\varepsilon_t^T \phi \phi^T}{\phi^T \hat{\Sigma} \phi} Q^\pi \right\} E \left\{ \sum_{t=1}^T \frac{\phi \phi^T}{\phi^T \hat{\Sigma} \phi} Q^\pi \right\}^{-1}\end{aligned}$$

and finally with $\mathbf{W} = \phi \phi^T (\phi^T \hat{\Sigma} \phi)^{-1}$ we get

$$\theta' = \theta + E \left\{ \sum_{t=1}^T \mathbf{W} Q^\pi \right\}^{-1} E \left\{ \sum_{t=1}^T \mathbf{W} \varepsilon_t Q^\pi \right\}. \quad (4.6)$$

If $\hat{\Sigma}$ is diagonal, that is, the exploration of the parameters is pairwise independent, all parameters employ the same exploration, and the exploration is constant over rollouts, \mathbf{W} simplifies to $\mathbf{W} = \phi \phi^T (\phi^T \phi)^{-1}$. Normalized basis functions ϕ further simplify \mathbf{W} to $\mathbf{W}(\mathbf{s}, t) = \phi \phi^T$.

If only one parameter is active at each time step, $\mathbf{W}(\mathbf{s}, t)$ is diagonal and Equation (4.6) simplifies to

$$\theta'_i = \theta_i + \frac{E \left\{ \sum_{t=1}^T \phi_i^2 / (\phi^T \hat{\Sigma} \phi) \varepsilon_{i,t} Q^\pi \right\}}{E \left\{ \sum_{t=1}^T \phi_i^2 / (\phi^T \hat{\Sigma} \phi) Q^\pi \right\}} \quad (4.7)$$

$$= \theta_i + \frac{E \left\{ \sum_{t=1}^T \sigma_i^{-1} \varepsilon_{i,t} Q^\pi \right\}}{E \left\{ \sum_{t=1}^T \sigma_i^{-1} Q^\pi \right\}}, \quad (4.8)$$

where θ'_i is one individual parameter, ϕ_i and $\varepsilon_{i,t}$ are the corresponding elements of ϕ and ε_t , and σ_i is the respective entry of the diagonal of $\hat{\Sigma}$. If the σ_i are constant over the rollouts we get $\theta'_i = \theta_i + E \left\{ \sum_{t=1}^T \varepsilon_{i,t} Q^\pi \right\} / E \left\{ \sum_{t=1}^T Q^\pi \right\}$. The independence simplification in Equations (4.7, 4.8) works well in practice, even if there is some overlap between the activations, such as in the case of dynamical system motor primitives [Ijspeert et al., 2002a,b, Schaal et al., 2003, 2007]. Weighting the exploration with the basis functions, as in Equation (4.7), yields slightly better results than completely ignoring the interactions, as in Equation (4.8).

The policy can be equivalently expressed as

$$\pi(a_t | \mathbf{s}_t, t) = p(a_t | \mathbf{s}_t, t, \varepsilon_t) p(\varepsilon_t | \mathbf{s}_t, t) = p(a_t | \mathbf{s}_t, t, \varepsilon_t) \mathcal{N}(\varepsilon_t | \mathbf{0}, \hat{\Sigma}).$$

Applying Equation (4.3) to the variance $\hat{\Sigma}$ we get

$$\partial_{\hat{\Sigma}'} L_{\hat{\Sigma}}(\hat{\Sigma}') = E \left\{ \sum_{t=1}^T \partial_{\hat{\Sigma}'} \log \mathcal{N}(\varepsilon_t | \mathbf{0}, \hat{\Sigma}') Q^\pi \right\},$$

as $p(a_t | \mathbf{s}_t, t, \varepsilon_t)$ is independent from the variance. Setting this equation to zero and solving for $\hat{\Sigma}'$ yields

$$\hat{\Sigma}' = \frac{E \left\{ \sum_{t=1}^T \varepsilon_t \varepsilon_t^T Q^\pi \right\}}{E \left\{ \sum_{t=1}^T Q^\pi \right\}},$$

which is the same solution as we get in standard maximum likelihood problems.

The same derivation holds if we use Equation (4.2) instead of Equation (4.3). Then, the state-action value function Q^π is replaced everywhere by the return R .



5 Reinforcement Learning to Adjust Parametrized Motor Primitives to New Situations

Humans manage to adapt learned movements very quickly to new situations by generalizing learned behaviors from similar situations. In contrast, robots currently often need to re-learn the complete movement. In this chapter, we propose a method that learns to generalize parametrized motor plans by adapting a small set of global parameters, called meta-parameters. We employ reinforcement learning to learn the required meta-parameters to deal with the current situation, described by states. We introduce an appropriate reinforcement learning algorithm based on a kernelized version of the reward-weighted regression. To show its feasibility, we evaluate this algorithm on a toy example and compare it to several previous approaches. Subsequently, we apply the approach to three robot tasks, i.e., the generalization of throwing movements in darts, of hitting movements in table tennis, and of throwing balls where the tasks are learned on several different real physical robots, i.e., a Barrett WAM, a BioRob, the JST-ICORP/SARCOS CBI and a Kuka KR 6.

5.1 Introduction

Human movements appear to be represented using movement templates, also called motor primitives [Schmidt and Wrisberg, 2000]. Once learned, these templates allow humans to quickly adapt their movements to variations of the situation without the need of re-learning the complete movement. For example, the overall shape of table tennis forehands are very similar when the swing is adapted to varied trajectories of the incoming ball and a different targets on the opponent's court. To accomplish such behavior, the human player has learned by trial and error how the global parameters of a generic forehand need to be adapted due to changes in the situation [Mülling et al., 2010, 2011].

In robot learning, motor primitives based on dynamical systems [Ijspeert et al., 2002b, Schaal et al., 2007] can be considered a technical counterpart to these templates. They allow acquiring new behaviors quickly and reliably both by imitation and reinforcement learning. Resulting successes have shown that it is possible to rapidly learn motor primitives for complex behaviors such as tennis-like swings [Ijspeert et al., 2002b], T-ball batting [Peters and Schaal, 2008a], drumming [Pongas et al., 2005], biped locomotion [Nakanishi et al., 2004], ball-in-a-cup [Kober and Peters, 2011a], and even in tasks with potential industrial applications [Urbanek et al., 2004]. While the examples are impressive, they do not yet address how a motor primitive can be generalized to a different behavior by trial and error without re-learning the task. Such generalization of behaviors can be achieved by adapting the meta-parameters of the movement representation. Meta-parameters are defined as a small set of parameters that adapt the global movement behavior. The dynamical system motor primitives can be adapted both spatially and temporally without changing the overall shape of the motion [Ijspeert et al., 2002b]. In this chapter, we learn a mapping from a range of changed situations, described by states, to the meta-parameters to adapt the template's behavior. We consider movements where it is sufficient to reshape (e.g., rescale the motion spatially and/or temporally) the global movement by optimizing meta-parameters to adapt to a new situation instead of tuning the movement primitive's shape parameters that describe the fine details of the movement.

Dynamical systems motor primitives have the capability to adapt the movement to a changed end positions. Here, the end position is a meta-parameter. This was exploited in [Ijspeert et al., 2002b] for tennis-like swings with static ball targets and in [Pastor et al., 2009] for object manipulation. In these papers, the desired end position is given in Cartesian coordinates and the movement primitives operate in Cartesian coordinates as well. Thus, the meta-parameters of the motor primitives are straightforward to set. In this chapter, we are interested in non-intuitive connections, where the relation between the desired outcome and the meta-parameters is not straightforward. There is related prior work in the context of programming by demonstration by Ude et al. [2010] and Kronander et al. [2011] who employ

supervised learning to learn a mapping from desired outcomes to meta-parameters for tasks such as reaching, throwing, drumming, and mini-golf. They assume that a teacher has presented a number of demonstrations that cannot be contradictory and the task is to imitate and generalize these demonstrations. Lampariello et al. [2011] employ a global planner to provide demonstrations of optimal catching meta-parameters and use supervised learning approaches to generalize these in real-time. In contrast, in our setting the robot actively explores different movements and improves the behavior according to a cost function. It can deal with contradictory demonstrations and actively generate its own scenarios by exploration combined with self-improvement. As mentioned in [Ude et al., 2010], the two approaches may even be complimentary: reinforcement learning can provide demonstrations for supervised learning, and supervised learning can be used as a starting point for reinforcement learning.

Adapting movements to situations is also discussed in [Jetchev and Toussaint, 2009] in a supervised learning setting. Their approach is based on predicting a trajectory from a previously demonstrated set and refining it by motion planning. The authors note that kernel ridge regression performed poorly for the prediction if the new situation is far from previously seen ones as the algorithm yields the global mean. In our approach, we employ a cost weighted mean that overcomes this problem. If the situation is far from previously seen ones, large exploration will help to find a solution.

In machine learning, there have been many attempts to use meta-parameters in order to generalize between tasks [Caruana, 1997]. Particularly, in grid-world domains, significant speed-up could be achieved by adjusting policies by modifying their meta-parameters, e.g., re-using options with different subgoals [McGovern and Barto, 2001]. The learning of meta-parameters of the learning algorithm has been proposed as a model for neuromodulation in the brain [Doya, 2002]. In contrast, we learn the meta-parameters of a motor skill in this chapter. In robotics, such meta-parameter learning could be particularly helpful due to the complexity of reinforcement learning for complex motor skills with high dimensional states and actions. The cost of experience is high as sample generation is time consuming and often requires human interaction (e.g., in cart-pole, for placing the pole back on the robot's hand) or supervision (e.g., for safety during the execution of the trial). Generalizing a teacher's demonstration or a previously learned policy to new situations may reduce both the complexity of the task and the number of required samples. Hence, a reinforcement learning method for acquiring and refining meta-parameters of pre-structured primitive movements becomes an essential next step, which we will address in this chapter.

This chapter does not address the problem of deciding whether it is more advantageous to generalize existing generic movements or to learn a novel one. Similar to most reinforcement learning approaches, the states and meta-parameters (which correspond to actions in the standard reinforcement learning settings) as well as the cost or reward function need to be designed by the user prior to the learning process. Here, we can only provide a few general indications with regard to the choice of these setting. Cost functions need to capture the desired outcome of the reinforcement learning process. Often the global target can be described verbally - but it is not obvious how the cost needs to be scaled and how to take secondary optimization criteria into account. For example, when throwing at a target, the global goal is hitting it. However, it is not always obvious which distance metric should be used to score misses, which secondary criteria (e.g. required torques) should be included, and which weight each criterion should be assigned. These choices influence both the learning performance and the final policy. Even for human reaching movements, the underlying cost function is not completely understood [Bays and Wolpert, 2007]. In practice, informative cost functions (i.e., cost functions that contain a notion of closeness) often perform better than binary reward functions in robotic tasks. In this chapter, we used a number of cost functions both with and without secondary objectives. In the future, inverse reinforcement learning [Russell, 1998] may be a useful alternative to automatically recover underlying cost functions from data as done already in other settings.

The state of the environment needs to enable the robot to obtain sufficient information to react appropriately. The proposed algorithm can cope with superfluous states at a cost of slower learning. Similarly, the meta-parameters are defined by the underlying representation of the movement. For example, the dynamical systems motor primitives [Ijspeert et al., 2002b, Schaal et al., 2007] have meta-

parameters for scaling the duration and amplitude of the movement as well as the possibility to change the final position. Restricting the meta-parameters to task relevant ones, may often speed up the learning process.

We present current work on automatic meta-parameter acquisition for motor primitives by reinforcement learning. We focus on learning the mapping from situations to meta-parameters and how to employ these in dynamical systems motor primitives. We extend the motor primitives of Ijspeert et al. [2002b] with a learned meta-parameter function and re-frame the problem as an episodic reinforcement learning scenario. In order to obtain an algorithm for fast reinforcement learning of meta-parameters, we view reinforcement learning as a reward-weighted self-imitation [Peters and Schaal, 2008b, Kober and Peters, 2011a].

To have a general meta-parameter learning, we adopted a parametric method, the reward-weighted regression [Peters and Schaal, 2008b], and turned it into a non-parametric one. We call this method Cost-regularized Kernel Regression (CrKR), which is related to Gaussian process regression [Rasmussen and Williams, 2006] but differs in the key aspects of incorporating costs and exploration naturally. We compare the CrKR with a traditional policy gradient algorithm [Peters and Schaal, 2008a], the reward-weighted regression [Peters and Schaal, 2008b], and supervised learning [Ude et al., 2010, Kronander et al., 2011] on a toy problem in order to show that it outperforms available previously developed approaches. As complex motor control scenarios, we evaluate the algorithm in the acquisition of flexible motor primitives for dart games such as *Around the Clock* [Masters Games Ltd., 2010], for *table tennis*, and for *ball target throwing*.

5.2 Meta-Parameter Learning for Motor Primitives

The goal of this chapter is to show that elementary movements can be generalized by modifying only the meta-parameters of the primitives using learned mappings based on self-improvement. In Section 5.2.1, we first review how a single primitive movement can be represented and learned. We discuss how meta-parameters may be able to adapt the motor primitive spatially and temporally to the new situation. In order to develop algorithms that learn to automatically adjust such motor primitives, we model meta-parameter self-improvement as an episodic reinforcement learning problem in Section 5.2.2. While this problem could in theory be treated with arbitrary reinforcement learning methods, the availability of few samples suggests that more efficient, task appropriate reinforcement learning approaches are needed. To avoid the limitations of parametric function approximation, we aim for a kernel-based approach. When a movement is generalized, new parameter settings need to be explored. Hence, a predictive distribution over the meta-parameters is required to serve as an exploratory policy. These requirements lead to the method which we derive in Section 5.2.3 and employ for meta-parameter learning in Section 5.2.4.

5.2.1 Motor Primitives with Meta-Parameters

In this section, we review how the dynamical systems motor primitives [Ijspeert et al., 2002b, Schaal et al., 2007] discussed in Chapter 3 can be used for meta-parameter learning. The dynamical system motor primitives are a powerful movement representation that allows ensuring the stability of the movement¹, choosing between a rhythmic and a discrete movement and is invariant under rescaling of both duration and movement amplitude. These modification parameters can become part of the meta-parameters of the movement.

In this chapter, we focus on single stroke movements which appear frequently in human motor control [Wulf, 2007, Schaal et al., 2007]. Therefore, we will always focus on the discrete version of the dynamical system motor primitives in this chapter.

¹ Note that the dynamical systems motor primitives ensure the stability of the movement generation but cannot guarantee the stability of the movement execution [Ijspeert et al., 2002b, Schaal et al., 2007].




Figure 5.1: This figure illustrates a table tennis task. The situation, described by the state s , corresponds to the positions and velocities of the ball and the robot at the time the ball is above the net. The meta-parameters γ are the joint positions and velocity at which the ball is hit. The policy parameters represent the backward motion and the movement on the arc. The meta-parameter function $\tilde{\gamma}(s)$, which maps the state to the meta-parameters, is learned.

The motor primitive policy is invariant under transformations of the initial position \mathbf{x}_1^0 , the initial velocity \mathbf{x}_2^0 , the goal \mathbf{g} , the goal velocity $\dot{\mathbf{g}}$, the amplitude A , and the duration T as discussed in Chapter 3. These six modification parameters can be used as the meta-parameters γ of the movement. Obviously, we can make more use of the motor primitive framework by adjusting the meta-parameters γ depending on the current situation or state s according to a meta-parameter function $\tilde{\gamma}(s)$. The meta-parameter γ is treated as a random variable where the variance corresponds to the uncertainty. The state s can for example contain the current position, velocity and acceleration of the robot and external objects, as well as the target to be achieved. This paper focuses on learning the meta-parameter function $\tilde{\gamma}(s)$ by episodic reinforcement learning.

Illustrations of the Learning Problem

We discuss the resulting learning problem based on the two examples shown in Figures 5.1 and 5.2.

As a first illustration of the meta-parameter learning problem, we take a table tennis task which is illustrated in Figure 5.1 (in Section 5.3.3, we will expand this example to a robot application). Here, the desired skill is to return a table tennis ball. The motor primitive corresponds to the hitting movement. When modeling a single hitting movement with dynamical-systems motor primitives [Ijspeert et al., 2002b], the combination of retracting and hitting motions would be represented by one movement primitive and can be learned by determining the movement parameters θ . These parameters can either be estimated by imitation learning or acquired by reinforcement learning. The return can be adapted by changing the paddle position and velocity at the hitting point. These variables can be influenced by modifying the meta-parameters of the motor primitive such as the final joint positions and velocities. The state consists of the current positions and velocities of the ball and the robot at the time the ball is directly above the net. The meta-parameter function $\tilde{\gamma}(s)$ maps the state (the state of the ball and the robot before the return) to the meta-parameters γ (the final positions and velocities of the motor primitive). Its variance corresponds to the uncertainty of the mapping.

In a 2D dart throwing task with a dart on a launcher which is illustrated in Figure 5.2 (in Section 5.3.2, we will expand this example to a robot application) the desired skill is to hit a specified point on a wall with a dart. The dart is placed on the launcher and held there by friction. The motor primitive corresponds to the throwing of the dart. When modeling a single dart's movement with dynamical-systems motor primitives [Ijspeert et al., 2002b], the combination of retracting and throwing motions would be represented by the movement parameters θ of one movement primitive. The dart's impact position can be adapted to a desired target by changing the velocity and the angle at which the dart leaves the launcher.




Figure 5.2: This figure illustrates a 2D dart throwing task. The situation, described by the state s corresponds to the relative height. The meta-parameters γ are the velocity and the angle at which the dart leaves the launcher. The policy parameters represent the backward motion and the movement on the arc. The meta-parameter function $\bar{\gamma}(s)$, which maps the state to the meta-parameters, is learned.

These variables can be influenced by changing the meta-parameters of the motor primitive such as the final position of the launcher and the duration of the throw. The state consists of the current position of the hand and the desired position on the target. If the thrower is always at the same distance from the wall the two positions can be equivalently expressed as the vertical distance. The meta-parameter function $\bar{\gamma}(s)$ maps the state (the relative height) to the meta-parameters γ (the final position g and the duration of the motor primitive T).

The approach presented in this chapter is applicable to any movement representation that has meta-parameters, i.e., a small set of parameters that allows to modify the movement. In contrast to [Lampariello et al., 2011, Jetchev and Toussaint, 2009, Grimes and Rao, 2008, Bentivegna et al., 2004b] our approach does not require explicit (re-)planning of the motion.

In the next sections, we derive and apply an appropriate reinforcement learning algorithm.

5.2.2 Problem Statement: Meta-Parameter Self-Improvement

The problem of meta-parameter learning is to find a stochastic policy $\pi(\gamma|\mathbf{x}) = p(\gamma|\mathbf{s})$ that maximizes the expected return

$$J(\pi) = \int_{\mathbb{S}} p(\mathbf{s}) \int_{\mathbb{G}} \pi(\gamma|\mathbf{s}) R(\mathbf{s}, \gamma) d\gamma d\mathbf{s}, \quad (5.1)$$

where \mathbb{S} denotes the the space of states \mathbf{s} , \mathbb{G} denotes the the space of meta-parameters γ , and $R(\mathbf{s}, \gamma)$ denotes all the rewards following the selection of the meta-parameter γ according to a situation described by state \mathbf{s} . Such a policy $\pi(\gamma|\mathbf{x})$ is a probability distribution over meta-parameters given the current state. The stochastic formulation allows a natural incorporation of exploration, and the optimal time-invariant policy has been shown to be stochastic in the case of hidden state variables [Sutton et al., 1999, Jaakkola et al., 1993]. The return of an episode is $R(\mathbf{s}, \gamma) = T^{-1} \sum_{t=0}^T r^t$ with number of steps T and rewards r^t . For a parametrized policy π with parameters \mathbf{w} it is natural to first try a policy gradient approach such as finite-difference methods, vanilla policy gradient approaches and natural gradients. While we will denote the shape parameters by θ , we denote the parameters of the meta-parameter function by \mathbf{w} . Reinforcement learning of the meta-parameter function $\bar{\gamma}(s)$ is not straightforward as only few examples can be generated on the real system and trials are often quite expensive. The credit assignment problem is non-trivial as the whole movement is affected by every change in the meta-parameter function. Early attempts using policy gradient approaches resulted in tens of thousands of trials even for simple toy problems, which is not feasible on a real system.

Algorithm 5.1 Meta-Parameter Learning

Preparation steps:

Learn one or more motor primitives by imitation and/or reinforcement learning (yields shape parameters θ).

Determine initial state \mathbf{s}^0 , meta-parameters γ^0 , and cost C^0 corresponding to the initial motor primitive.

Initialize the corresponding matrices $\mathbf{S}, \mathbf{\Gamma}, \mathbf{C}$.

Choose a kernel \mathbf{k}, \mathbf{K} .

Set a scaling parameter λ .

for all iterations j **do**

Determine the state \mathbf{s}^j specifying the situation.

Calculate the meta-parameters γ^j by:

Determine the mean of each meta-parameter i $\bar{\gamma}_i(\mathbf{s}^j) = \mathbf{k}(\mathbf{s}^j)^T (\mathbf{K} + \lambda \mathbf{C})^{-1} \mathbf{\Gamma}_i$,

Determine the variance $\sigma^2(\mathbf{s}^j) = \mathbf{k}(\mathbf{s}^j, \mathbf{s}^j) - \mathbf{k}(\mathbf{s}^j)^T (\mathbf{K} + \lambda \mathbf{C})^{-1} \mathbf{k}(\mathbf{s}^j)$,

Draw the meta-parameters from a Gaussian distribution $\gamma^j \sim \mathcal{N}(\gamma | \bar{\gamma}(\mathbf{s}^j), \sigma^2(\mathbf{s}^j) \mathbf{I})$

Execute the motor primitive using the new meta-parameters.

Calculate the cost c^j at the end of the episode.

Update $\mathbf{S}, \mathbf{\Gamma}, \mathbf{C}$ according to the achieved result.

end for

Dayan and Hinton [1997] showed that an immediate reward can be maximized by instead minimizing the Kullback-Leibler divergence $D(\pi(\gamma|\mathbf{s})R(\mathbf{s}, \gamma) || \pi'(\gamma|\mathbf{s}))$ between the reward-weighted policy $\pi(\gamma|\mathbf{s})$ and the new policy $\pi'(\gamma|\mathbf{s})$. As we are in an episodic setting, this form of optimization solves the considered problem. Williams [1992] suggested to use Gaussian noise in this context; hence, we employ a policy of the form

$$\pi(\gamma|\mathbf{s}) = \mathcal{N}(\gamma | \bar{\gamma}(\mathbf{s}), \sigma^2(\mathbf{s}) \mathbf{I}),$$

where we have the deterministic mean policy $\bar{\gamma}(\mathbf{s}) = \phi(\mathbf{s})^T \mathbf{w}$ with basis functions $\phi(\mathbf{s})$ and parameters \mathbf{w} as well as the variance $\sigma^2(\mathbf{s})$ that determines the exploration $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2(\mathbf{s}) \mathbf{I})$ as e.g., in [Peters and Schaal, 2008a]. The parameters \mathbf{w} can then be adapted by reward-weighted regression in an immediate reward [Peters and Schaal, 2008b] or episodic reinforcement learning scenario (see Chapter 4). The reasoning behind this reward-weighted regression is that the reward can be treated as an improper probability distribution over indicator variables determining whether the action is optimal or not.

5.2.3 A Task-Appropriate Reinforcement Learning Algorithm

Designing good basis functions is challenging, a nonparametric representation is better suited in this context. There is an intuitive way of turning the reward-weighted regression into a Cost-regularized Kernel Regression. The kernelization of the reward-weighted regression can be done straightforwardly (similar to Section 6.1 of [Bishop, 2006] for regular supervised learning). Inserting the reward-weighted

regression solution $\mathbf{w} = (\Phi^T \mathbf{R} \Phi + \lambda \mathbf{I})^{-1} \Phi^T \mathbf{R} \Gamma_i$ and using the Woodbury formula² [Welling, 2010], we transform reward-weighted regression into a Cost-regularized Kernel Regression

$$\begin{aligned} \tilde{\gamma}_i &= \phi(\mathbf{s})^T \mathbf{w} = \phi(\mathbf{s})^T (\Phi^T \mathbf{R} \Phi + \lambda \mathbf{I})^{-1} \Phi^T \mathbf{R} \Gamma_i \\ &= \phi(\mathbf{s})^T \Phi^T (\Phi \Phi^T + \lambda \mathbf{R}^{-1})^{-1} \Gamma_i, \end{aligned} \quad (5.2)$$

where the rows of Φ correspond to the basis functions $\phi(\mathbf{s}_i) = \Phi_i$ of the training examples, Γ_i is a vector containing the training examples for meta-parameter component γ_i , and λ is a ridge factor. Next, we assume that the accumulated rewards R_k are strictly positive $R_k > 0$ and can be transformed into costs by $c_k = 1/R_k$. Hence, we have a cost matrix $\mathbf{C} = \mathbf{R}^{-1} = \text{diag}(R_1^{-1}, \dots, R_n^{-1})$ with the cost of all n data points. After replacing $\mathbf{k}(\mathbf{s}) = \phi(\mathbf{s})^T \Phi^T$ and $\mathbf{K} = \Phi \Phi^T$, we obtain the Cost-regularized Kernel Regression

$$\tilde{\gamma}_i = \tilde{\gamma}_i(\mathbf{s}) = \mathbf{k}(\mathbf{s})^T (\mathbf{K} + \lambda \mathbf{C})^{-1} \Gamma_i,$$

which gives us a deterministic policy. Here, costs correspond to the uncertainty about the training examples. Thus, a high cost is incurred for being further away from the desired optimal solution at a point. In our formulation, a high cost therefore corresponds to a high uncertainty of the prediction at this point.

In order to incorporate exploration, we need to have a stochastic policy and, hence, we need a predictive distribution. This distribution can be obtained by performing the policy update with a Gaussian process regression and we directly see from the kernel ridge regression that

$$\sigma^2(\mathbf{s}) = k(\mathbf{s}, \mathbf{s}) + \lambda - \mathbf{k}(\mathbf{s})^T (\mathbf{K} + \lambda \mathbf{C})^{-1} \mathbf{k}(\mathbf{s}),$$

where $k(\mathbf{s}, \mathbf{s}) = \phi(\mathbf{s})^T \phi(\mathbf{s})$ is the norm of the point in the kernel space. We call this algorithm Cost-regularized Kernel Regression. Algorithm 5.1 describes the complete learning procedure, where the rows of \mathbf{S} correspond to the states of the training examples $\mathbf{s}_i = \mathbf{S}_i$.

The algorithm corresponds to a Gaussian process regression where the costs on the diagonal are input-dependent noise priors. The parameter λ acts as a exploration-exploitation trade-off parameter as illustrated in Figure 5.5. Gaussian processes have been used previously for reinforcement learning [Engel et al., 2005] in value function based approaches while here we use them to learn the policy.

5.2.4 Meta-Parameter Learning by Reinforcement Learning

As a result of Section 5.2.3, we have a framework of motor primitives as introduced in Section 5.2.1 that we can use for reinforcement learning of meta-parameters as outlined in Section 5.2.2. We have generalized the reward-weighted regression policy update to instead become a Cost-regularized Kernel Regression (CrKR) update where the predictive variance is used for exploration. In Algorithm 1, we show the complete algorithm resulting from these steps.

The algorithm receives three inputs, i.e., (i) a motor primitive that has associated meta-parameters γ , (ii) an initial example containing state \mathbf{s}^0 , meta-parameter γ^0 and cost C^0 , as well as (iii) a scaling parameter λ . The initial motor primitive can be obtained by imitation learning [Ijspeert et al., 2002b] and, subsequently, improved by parametrized reinforcement learning algorithms such as policy gradients [Peters and Schaal, 2008a] or Policy learning by Weighting Exploration with the Returns (PoWER) [Kober and Peters, 2011a, Chapter 4]. The demonstration also yields the initial example needed for meta-parameter learning. While the scaling parameter is an open parameter, it is reasonable to choose it as a fraction of the average cost and the output noise parameter (note that output noise and other possible hyper-parameters of the kernel can also be obtained by approximating the unweighted meta-parameter function).

² The equality $(\Phi^T \mathbf{R} \Phi + \lambda \mathbf{I})^{-1} \Phi^T \mathbf{R} = \Phi^T (\Phi \Phi^T + \lambda \mathbf{R}^{-1})^{-1}$ is straightforward to verify by left and right multiplying the non-inverted terms: $\Phi^T \mathbf{R} (\Phi \Phi^T + \lambda \mathbf{R}^{-1}) = (\Phi^T \mathbf{R} \Phi + \lambda \mathbf{I}) \Phi^T$.




Figure 5.3: This figure illustrates the meaning of policy improvements with Cost-regularized Kernel Regression. Each sample consists of a state, a meta-parameter and a cost where the cost is indicated the blue error bars. The red line represents the improved mean policy, the dashed green lines indicate the exploration/variance of the new policy. For comparison, the gray lines show standard Gaussian process regression. As the cost of a data point is equivalent to having more noise, pairs of states and meta-parameter with low cost are more likely to be reproduced than others with high costs.

Illustration of the Algorithm

In order to illustrate this algorithm, we will use the example of the table tennis task introduced in Section 5.2.1. Here, the robot should hit the ball accurately while not destroying its mechanics. Hence, the cost could correspond to the distance between the ball and the paddle, as well as the squared torques. The initial policy is based on a prior, illustrated in Figure 5.3(a), that has a variance for initial exploration (it often makes sense to start with a uniform prior). This variance is used to enforce exploration. To return a ball, we sample the meta-parameters from the policy based on the current state. After the trial the cost is determined and, in conjunction with the employed meta-parameters, used to update the policy. If the cost is large (e.g., the ball was far from the racket), the variance of the policy is large as it may still be improved and therefore needs exploration. Furthermore, the mean of the policy is shifted only slightly towards the observed example as we are uncertain about the optimality of this action. If the cost is small, we know that we are close to an optimal policy (e.g., the racket hit the ball off-center) and only have to search in a small region around the observed trial. The effects of the cost on the mean and the variance are illustrated in Figure 5.3(b). Each additional sample refines the policy and the overall performance improves (see Figure 5.3(c)). If a state is visited several times and different meta-parameters are sampled, the policy update must favor the meta-parameters with lower costs. If several sets of meta-parameters have similarly low costs, where it converges depends on the order of samples. The cost function should be designed to avoid this behavior and to favor a single set. The exploration has to be restricted to safe meta-parameter ranges. Algorithm 1 exhibits this behavior as the exploration is only local and restricted by the prior (see Figure 5.3). If the initial policy is safe, exploring the neighboring regions is likely to be safe as well. Additionally, lower level controllers as well as the mechanics of the robot ensure that kinematic and dynamic constraints are satisfied and a term in the cost function can be used to discourage potentially harmful movements.

In the example of the 2D dart throwing task, the cost is similar. Here, the robot should throw darts accurately while not destroying its mechanics. Hence, the cost could correspond to the error between desired goal and the impact point, as well as the absolute velocity of the end-effector. Often the state is determined by the environment, e.g., the ball trajectory in table tennis depends on the opponent. However, for the dart setting, we could choose the next target and thus employ CrKR as an active learning approach by picking states with large variances. In the dart throwing example we have a correspondence between the state and the outcome similar to a regression problem. However, the mapping between the state and the meta-parameter is not unique. The same height can be achieved by different combinations




Figure 5.4: This figure shows the performance of the compared algorithms averaged over 10 complete learning runs. Cost-regularized Kernel Regression finds solutions with the same final performance two orders of magnitude faster than the finite difference gradient (FD) approach and twice as fast as the reward-weighted regression. At the beginning FD often is highly unstable due to our attempts of keeping the overall learning speed as high as possible to make it a stronger competitor. The lines show the median and error bars indicate standard deviation. The initialization and the initial costs are identical for all approaches. However, the omission of the first twenty rollouts was necessary to cope with the logarithmic rollout axis. The number of rollouts includes the rollouts not used to update the policy.

of velocities and angles. Averaging these combinations is likely to generate inconsistent solutions. The regression must hence favor the meta-parameters with the lower costs. CrKR can be employed as a regularized regression method in this setting.

5.3 Evaluations and Experiments

In Section 5.2, we have introduced both a framework for meta-parameter self-improvement as well as an appropriate reinforcement learning algorithm used in this framework. In this section, we will first show that the presented reinforcement learning algorithm yields higher performance than off-the shelf approaches. Hence, we compare it on a simple planar cannon shooting problem [Lawrence et al., 2003] with the preceding reward-weighted regression, an off-the-shelf finite difference policy gradient approach, and show the advantages over supervised learning approaches.

The resulting meta-parameter learning framework can be used in a variety of settings in robotics. We consider three scenarios here, i.e., (i) dart throwing with a simulated Barrett WAM, a real Kuka KR 6, and the JST-ICORP/SARCOS humanoid robot CBi [Cheng et al., 2007], (ii) table tennis with a simulated robot arm and a real Barrett WAM, and (iii) throwing a ball at targets with a MATLAB simulation and a real BioRob [Lens et al., 2010].

5.3.1 Benchmark Comparison: Toy Cannon Shots

In the first task, we only consider a simple simulated planar cannon shooting where we benchmark our Reinforcement Learning by Cost-regularized Kernel Regression approach against a finite difference gradient estimator and the reward-weighted regression. Additionally we contrast our reinforcement learning approach to a supervised one. Here, we want to learn an optimal policy for a 2D toy cannon environment similar to [Lawrence et al., 2003]. This benchmark example serves to illustrate our approach and to compare it to various previous approaches.




Figure 5.5: This figure illustrates the influence of the parameter λ for the Cost-regularized Kernel Regression. The red curve ($\lambda = 0.5$) corresponds to the red curve (Cost-regularized Kernel Regression) in Figure 5.4(c). The parameter λ trades off the exploration versus the exploitation. A higher λ leads to larger exploration and, thus, faster convergence to a suboptimal solution. The results are averaged over 10 complete learning runs. The lines show the median and error bars indicate standard deviation. The number of rollouts includes the rollouts not used to update the policy.




Figure 5.6: In this figure, we compare Gaussian process regression (GPR) in a supervised learning setting as proposed by [Ude et al., 2010, Kronander et al., 2011] to Cost-regularized Kernel Regression (CrKR) in a reinforcement learning setting. The red curve corresponds to the red curve (Cost-regularized Kernel Regression) in Figure 5.4. The GPR is trained with samples from the prior used for the CrKR (blue line) and with samples of the final CrKR policy (cyan line) respectively. The black line indicates the cost after CrKR has converged. GPR with samples drawn from the final policy performs best. Please note that this comparison is contrived as the role of CrKR is to discover the policy that is provided to “GPR structured policy”. GPR can only reproduce the demonstrated policy, which is achieved perfectly with 1000 samples. GPR can reproduce the demonstrated policy more accurately if more samples are available. However, it cannot improve the policy according to a cost function and it is impacted by contradictory demonstrations. The results are averaged over 10 complete learning runs. The lines show the median and error bars indicate standard deviation. The number of rollouts includes the rollouts not used to update the policy.

The setup is given as follows: A toy cannon is at a fixed location $[0.0, 0.1] m$. The trajectory of the cannon ball depends on the angle with respect to the ground and the speed at which it leaves the cannon. The flight of the canon ball is simulated as ballistic flight of a point mass with Stokes’s drag as wind model. The cannon ball is supposed to hit the ground at a desired distance. The desired distance $[1..3] m$ and the wind speed $[0..1] m/s$, which is always horizontal, are used as input states, the velocities in horizontal and vertical directions are the meta-parameters (which influences the angle and the speed of the ball leaving the cannon). In this benchmark we do not employ the motor primitives but set the meta-parameters directly. Lower speed can be compensated by a larger angle. Thus, there are different possible policies for hitting a target; we intend to learn the one which is optimal for a given cost function. This cost function is defined as

$$c = (b_x - s_x)^2 + 0.01 (\dot{b}_x^2 + \dot{b}_z^2),$$

where b_x is the impact position on the ground, s_x the desired impact position as indicated by the state, and $\dot{b}_{\{x,z\}}$ are the horizontal and vertical velocities of the cannon ball at the impact point respectively. It corresponds to maximizing the precision while minimizing the employed energy according to the chosen weighting. The input states (desired distance and wind speed) are drawn from a uniform distribution and directly passed to the algorithms. All approaches performed well in this setting, first driving the position error to zero and, subsequently, optimizing the impact velocity. The experiment was initialized with $[1, 10] m/s$ as initial ball velocities and $1 m/s$ as wind velocity. This setting corresponds to a very high parabola, which is far from optimal. For plots, we evaluate the policy on a test set of 25 uniformly randomly chosen points that remain the same throughout of the experiment and are never used in the learning process but only to generate Figure 5.4.

We compare our novel algorithm to a finite difference policy gradient (FD) method [Peters and Schaal, 2008a] and to the reward-weighted regression (RWR) [Peters and Schaal, 2008b]. The FD method uses a parametric policy that employs radial basis functions in order to represent the policy and perturbs the parameters. We used 25 Gaussian basis functions on a regular grid for each meta-parameter, thus a total of 50 basis functions. The number of basis functions, the learning rate, as well as the magnitude of the perturbations were tuned for best performance. We used 51 sets of uniformly perturbed parameters for each update step. The perturbed policies were evaluated on a batch of 25 input parameters to avoid over-fitting on specific input states. The FD algorithm converges after approximately 2000 batch gradient evaluations, which corresponds to 2,550,000 shots with the toy cannon.

The RWR method uses the same parametric policy as the finite difference gradient method. Exploration is achieved by adding Gaussian noise to the mean policy. All open parameters were tuned for best performance. The reward transformation introduced by Peters and Schaal [2008b] did not improve performance in this episodic setting. The RWR algorithm converges after approximately 40,000 shots with the toy cannon. For the Cost-regularized Kernel Regression (CrKR) the inputs are chosen randomly from a uniform distribution. We use Gaussian kernels and the open parameters were optimized by cross-validation on a small test set prior to the experiment. Each trial is added as a new training point if it landed in the desired distance range. The CrKR algorithm converges after approximately 20,000 shots with the toy cannon. The bandwidth of the kernels used for CrKR is in the same order of magnitude as the bandwidth of the basis functions. However, due to the non-parametric nature of CrKR, narrower kernels can be used to capture more details in order to improve performance. Figure 5.5 illustrates the influence of the parameter λ for the CrKR.

After convergence, the costs of CrKR are the same as for RWR and slightly lower than those of the FD method. The CrKR method needs two orders of magnitude fewer shots than the FD method. The RWR approach requires twice the shots of CrKR demonstrating that a non-parametric policy, as employed by CrKR, is better adapted to this class of problems than a parametric policy. The squared error between the actual and desired impact is approximately 5 times higher for the finite difference gradient method, see Figure 5.4.




Figure 5.7: This figure shows a dart throw in a physically realistic simulation.

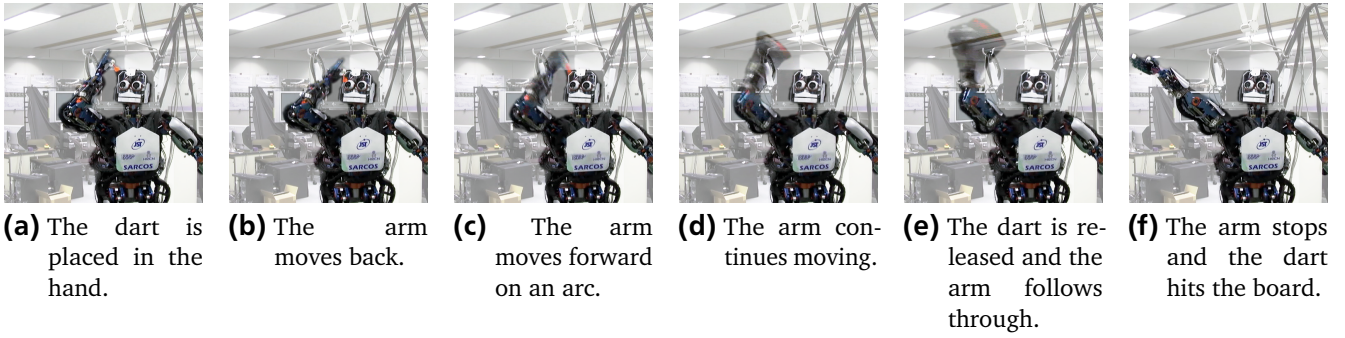


Figure 5.8: This figure shows a dart throw on the real JST-ICORP/SARCOS humanoid robot CBI.

Compared to standard Gaussian process regression (GPR) in a supervised setting, CrKR can improve the policy over time according to a cost function and outperforms GPR in settings where different combinations of meta-parameters yield the same result. For details, see Figure 5.6.

5.3.2 Robot Dart-Throwing Games

Now, we turn towards the complete framework, i.e., we intend to learn the meta-parameters for motor primitives in discrete movements. We compare the Cost-regularized Kernel Regression (CrKR) algorithm to the reward-weighted regression (RWR). As a sufficiently complex scenario, we chose a robot dart throwing task inspired by [Lawrence et al., 2003]. However, we take a more complicated scenario and choose dart games such as *Around the Clock* [Masters Games Ltd., 2010] instead of simple throwing at a fixed location. Hence, it will have an additional parameter in the state depending on the location on the dartboard that should come next in the sequence. The acquisition of a basic motor primitive is achieved using previous work on imitation learning [Ijspeert et al., 2002b]. Only the meta-parameter function is learned using CrKR or RWR. For the learning process, the targets (which are part of the state) are uniformly distributed on the dartboard. For the evaluation the targets are placed in the center of the fields. The reward is calculated based on the impact position observed by a vision system in the real robot experiments or the simulated impact position.

The dart is placed on a launcher attached to the end-effector and held there by stiction. We use the Barrett WAM robot arm in order to achieve the high accelerations needed to overcome the stiction. See Figure 5.7, for a complete throwing movement. The motor primitive is trained by imitation learning with kinesthetic teach-in. We use the Cartesian coordinates with respect to the center of the dart board as input states. In comparison with the benchmark example, we cannot directly influence the release velocity in this setup. Hence, we employ the parameter for the final position \mathbf{g} , the time scale of the motor primitive τ and the angle around the vertical axis (i.e., the orientation towards the dart board to which the robot moves before throwing) as meta-parameters instead. The popular dart game *Around the Clock* requires




Figure 5.9: This figure shows the cost function of the dart-throwing task for a whole game *Around the Clock* in each rollout. The costs are averaged over 10 runs with the error-bars indicating standard deviation. The number of rollouts includes the rollouts not used to update the policy.

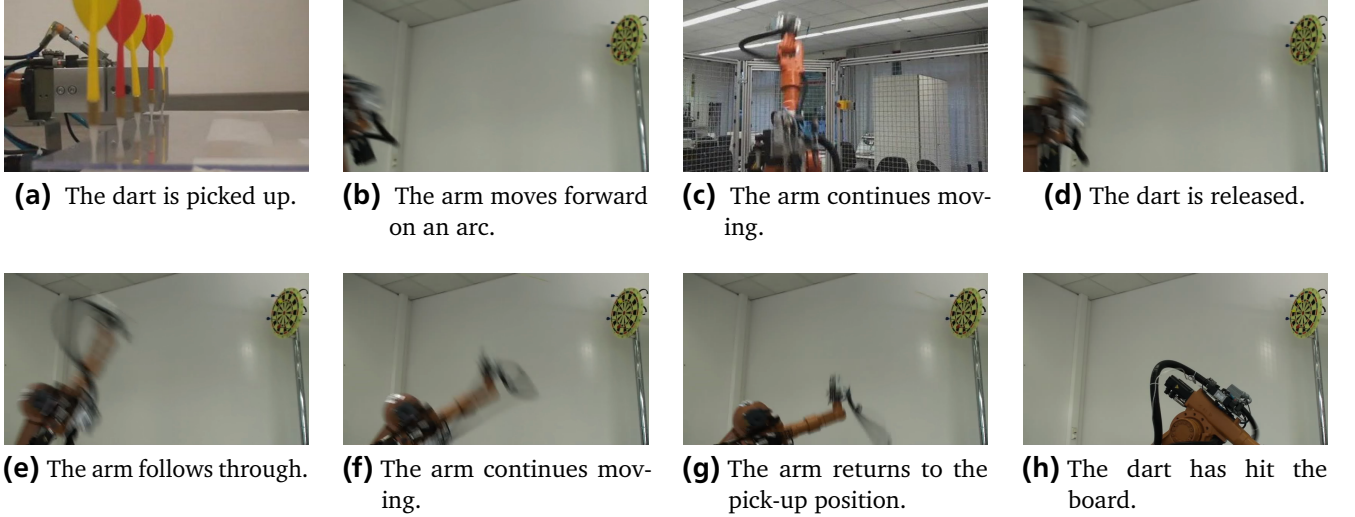


Figure 5.10: This figure shows a dart throw on the real Kuka KR 6 robot.

the player to hit the numbers in ascending order, then the bulls-eye. As energy is lost overcoming the stiction of the launching sled, the darts fly lower and we placed the dartboard lower than official rules require. The cost function is defined as

$$c = 10 \sqrt{\sum_{i \in \{x, z\}} (d_i - s_i)^2} + \tau,$$

where d_i are the horizontal and vertical positions of the dart on the dartboard after the throw, s_i are the horizontal and vertical positions of the target corresponding to the state, and τ corresponds to the velocity of the motion. After approximately 1000 throws the algorithms have converged but CrKR yields a high performance already much earlier (see Figure 5.9). We again used a parametric policy with radial basis functions for RWR. Here, we employed 225 Gaussian basis function on a regular grid per meta-parameter. Designing a good parametric policy proved very difficult in this setting as is reflected by the poor performance of RWR.

This experiment has also being carried out on three real, physical robots, i.e., a Barrett WAM, the humanoid robot CBi (JST-ICORP/SARCOS), and a Kuka KR 6. CBi was developed within the framework

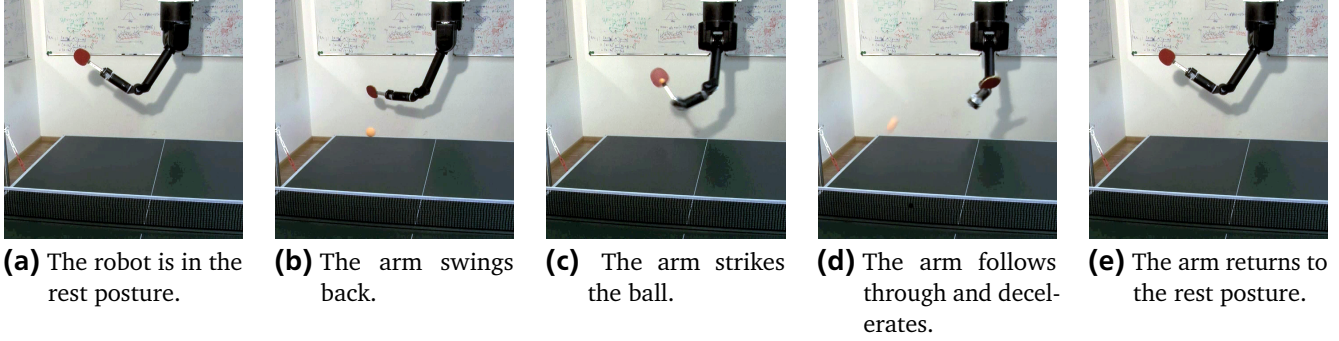


Figure 5.11: This figure shows the phases of a table tennis stroke on the real Barrett WAM.

of the JST-ICORP Computational Brain Project at ATR Computational Neuroscience Labs. The hardware of the robot was developed by the American robotic development company SARCOS. CBi can open and close the fingers which helps for more human-like throwing instead of the launcher employed by the Barrett WAM. See Figure 5.8 for a throwing movement.

We evaluated the approach on a setup using the Kuka KR 6 robot and a pneumatic gripper. The robot automatically picks up the darts from a stand. The position of the first degree of freedom (horizontal position) as well as the position of the fifth degree of freedom and the release timing (vertical position) were controlled by the algorithm. Due to inaccurate release timing the vertical position varied in a range of 10cm. Additionally the learning approach had to cope with non-stationary behavior as the outcome of the same set of parameters changed by one third of the dart board diameter upward. Despite these additional complications the robot learned to reliably (within the reproduction accuracy of 10cm as noted above) hit all positions on the dart board using only a total of 260 rollouts. See Figure 5.10 for a throwing movement.

5.3.3 Robot Table Tennis

In the second evaluation of the complete framework, we use the proposed method for hitting a table tennis ball in the air. The setup consists of a ball gun that serves to the forehand of the robot, a Barrett WAM and a standard sized table. The movement of the robot has three phases. The robot is in a rest posture and starts to swing back when the ball is launched. During this swing-back phase, the open parameters for the stroke are to be learned. The second phase is the hitting phase which ends with the contact of the ball and racket. In the final phase, the robot gradually ends the stroking motion and returns to the rest posture. See Figure 5.11 for an illustration of a complete episode and Chapter 3 for a more detailed description. The movements in the three phases are represented by three motor primitives obtained by imitation learning. We only learn the meta-parameters for the hitting phase.

The meta-parameters are the joint positions \mathbf{g} and velocities $\dot{\mathbf{g}}$ for all seven degrees of freedom at the end of the second phase (the instant of hitting the ball) and a timing parameter t_{hit} that controls when the swing back phase is transitioning to the hitting phase. For this task we employ a variant of the motor primitives that allows to set non-zero end velocities [Kober et al., 2010a]. We learn these 15 meta-parameters as a function of the state, which corresponds to the ball positions and velocities when it is directly over the net. We employed a Gaussian kernel and optimized the open kernel parameters according to typical values for the input and output beforehand. As cost function we employ

$$c = \sqrt{\sum_{i \in \{x, y, z\}} (b_i(t_{\text{hit}}) - p_i(t_{\text{hit}}))^2},$$

where $b_i(t_{\text{hit}})$ are the Cartesian positions of the ball and $p_i(t_{\text{hit}})$ are the Cartesian positions of the center of the paddle, both at the predicted hitting time t_{hit} . The policy is evaluated every 50 episodes with




Figure 5.12: This figure shows samples of the learned forehands. Note that this figure only illustrates the learned meta-parameter function in this context but cannot show timing (see Figure 5.13) and velocity and it requires a careful observer to note the important configuration differences resulting from the meta-parameters.




Figure 5.13: This figure illustrates the effect of the velocity of the ball towards the robot on the time it has until the ball needs to be hit. The plot was generated by sweeping through the velocity component towards the robot, keeping the other position and velocity values fixed. The line is the mean of 100 sweeps drawn from the same ball distribution as used during the learning.

25 ball launches picked randomly at the beginning of the learning. We initialize the behavior with five successful strokes observed from another player. After initializing the meta-parameter function with only these five initial examples, the robot misses approximately 95% of the balls as shown in Figure 5.15. Trials are only used to update the policy if the robot has successfully hit the ball as they did not significantly improve the learning performance and in order to keep the calculation sufficiently fast. Figures 5.12 and 5.14 illustrate different positions of the ball the policy is capable of dealing with after the learning. Figure 5.13 illustrates the dependence of the timing parameter on the ball velocity towards the robot and Figure 5.15 illustrates the costs over all episodes. For the results in Figure 5.15, we have simulated the flight of the ball as a simple ballistic point mass and the bouncing behavior using a restitution constant for the velocities. The state is directly taken from the simulated ball data with some added Gaussian noise. In the real robot experiment (Figure 5.16), the ball is shot with a ball cannon. The position of the ball is determined by two pairs of stereo cameras and the velocity is obtained by numerical differentiation. In this second setting, the state information is a lot less reliable due to noise in the vision system and even the same observed state can lead to different outcomes due to unobserved spin.




Figure 5.14: This figure shows samples of the learned forehands on the real robot.




Figure 5.15: This figure shows the cost function of the simulated table tennis task averaged over 10 runs with the error-bars indicating standard deviation. The red line represents the percentage of successful hits and the blue line the average cost. The number of rollouts includes the rollouts not used to update the policy. At the beginning the robot misses the ball 95% of the episodes and on average by 50 cm. At the end of the learning the robot hits almost all balls.

5.3.4 Active Learning of Ball Throwing

As an active learning setting, we chose a ball throwing task where the goal is to improve the throws while trying to perform well in a higher level game. For this scenario, it is important to balance learning of the individual actions by practicing them while at the same time, focusing on the overall performance in order to achieve the complete skill. Prominent examples are leisure time activities such as sports or motor skill games. For example, when playing darts with friends, you will neither always attempt the lowest risk action, nor always try to practice one particular throw, which will be valuable when mastered. Instead, you are likely to try plays with a reasonable level of risk and rely on safe throws in critical situations. This exploration is tightly woven into higher order dart games.

The higher level is modeled as a standard reinforcement learning problem with discrete states and actions. The lower level learning is done using CrKR. The higher level determines the target the robot is supposed to hit. The lower level has to learn how to hit this target. The transition probabilities of the higher level can be estimated from the learned meta-parameter function as explained in Section 5.3.4. We will discuss the rules of the game in Section 5.3.4, a simulated experiment in Section 5.3.4, and the results of an evaluation with a real BioRob in Section 5.3.4.




Figure 5.16: This figure shows the cost function of the table tennis task on the real robot. The policy was learned entirely on the real robot. The red line represents the percentage of successful hits and the blue line the average cost. The number of rollouts includes the rollouts not used to update the policy. At the beginning the robot misses the ball 70% of the episodes and on average by 15 cm. At the end of the learning the robot hits 80% of the balls.

Game used for the Evaluations

The game is reminiscent of blackjack as the goal is to collect as many points as possible without going over a threshold. The player throws a ball at three targets. The three rewards of one, two, and three are assigned to one target each. The setup of the game is illustrated in Figure 5.17. If the ball lands in the target, the player receives the corresponding number of points. The player starts with zero points if he gets more than 10 points he “busts” and incurs a loss of -10. The player has the option to “stand” (i.e., stop throwing and collect the accumulated number of points) at all times. Missing all targets does not entail a cost.

Two-Level Learning Approach

Our framework considers a hierarchy of two levels: a strategy level and a behavior level. The strategy level determines the strategy for the high-level moves, here termed “behaviors”, of the game. The behavior level deals with executing these behaviors in an optimal fashion. The strategy level chooses the next behavior, which is then executed by the behavior level. Upon completion of the behavior, the strategy level chooses the next behavior. The setup is illustrated in Figure 5.18.

We assume that the game has discrete states $s \in \mathbb{S}$ and discrete behaviors $b \in \mathbb{B}$. In the dart setting a behavior could be attempting to hit a specific field and the state could correspond to the current score. Given the current state, each behavior has an associated expected outcome $o \in \mathbb{O}$. For example, the behavior “throw at target X” has the outcome “change score by X” as a result of hitting target X. The transition probabilities \mathcal{P}_{so}^b of the strategy level would express how likely it is to hit a different field. The game can be modeled as an Markov decision process or MDP [Sutton and Barto, 1998], where the states consist of the number of accumulated points (zero to ten) and two additional game states (“bust” and “stand”). The behaviors correspond to attempting to throw at a specific target or to “stand” and are fixed beforehand. We assume to have an episodic game with a finite horizon, which can be expressed equivalently as an infinite horizon problem where we define an absorbing terminal state in which all actions receive an immediate reward of 0.

On the behavior level, we augment the state space with continuous states that describe the robot and the environment to form the combined state space \mathbf{s} . This state space could, for example, include the




Figure 5.17: This figure illustrates the side-stall game. The player throws the ball and if it lands in the target (illustrated by a wall with target holes) gets the number of points written next to it. Missing the targets is not punished, however, going over ten points leads to a loss of ten points.

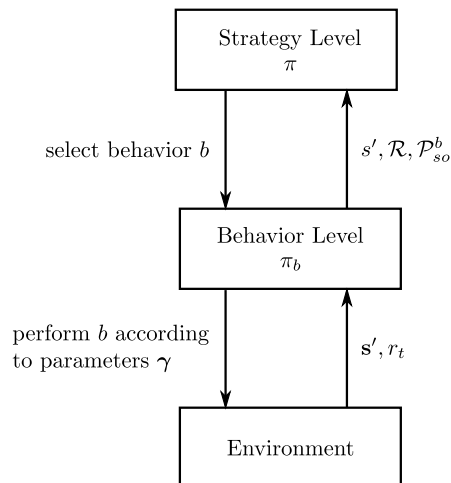


Figure 5.18: This figure illustrates the setup of the roles of the different levels.

position and velocity of the arm, the position of the targets as well as the current score. The actions are considered to be continuous and could, for example, be the accelerations of the arm. As the strategy level has to wait until the behavior is completed, the behaviors need to be of episodic nature as well. We have a single motor primitive representing the three behaviors of aiming at the three targets. Hitting the desired target is learned using CrKR. We employ Policy Iteration [Sutton and Barto, 1998] to learn on the strategy level.

The rewards for the strategy learning are fixed by the rules of the game. The possible states and behaviors also result from the way the game is played. The missing piece for the strategy learning is the transition probabilities \mathcal{P}_{so}^b . The behavior learning by CrKR associates each behavior with a variance. Each of these behaviors correspond to an expected change in state, the outcome o . For example “aim at 2” corresponds to “increase score by 2”. However, the meta-parameter function does not explicitly include information regarding what happens if the expected change in state is not achieved. We assume that there is a discrete set of outcomes $o \in \mathbb{O}$ (i.e., change in state) for all behaviors b for a certain state s . For example in this game hitting each target, and missing, is associated with either increasing the player’s score, winning or to bust (i.e., going over ten). With the meta-parameter function, we can calculate the overlaps of the ranges of possible meta-parameters for the different behaviors. These overlaps can then be used to determine how likely it is to end up with a change of state associated with a behavior different from the desired one. This approach relies on the assumption that we know for each behavior the associated range of meta-parameters and their likelihood.

The meta-parameters are drawn according to a normal distribution, thus the overlap has to be weighted accordingly. The probability of the outcome o when performing behavior b can be calculated as follows:

$$\mathcal{P}_{so}^b = \int p^b(\gamma) \frac{p^o(\gamma)}{\sum_{k \in \mathbb{O}} p^k(\gamma)} d\gamma,$$

where γ is the meta-parameters, $p^b(\gamma)$ is the probability of picking the meta-parameter γ when performing behavior b , $p^o(\gamma)$ is the probability of picking the meta-parameter γ when performing the action associated to the considered outcome o , and $\sum_{k \in \mathbb{O}} p^k(\gamma)$ is the normalizing factor. This scenario has first been treated in [Kober and Peters, 2011b].

Evaluation in Simulation

We first evaluated our approach using a MATLAB based simulation. The throw is modeled as a two dimensional ballistic flight of a point mass. The targets correspond to segments of the ground line. The meta-parameters are the initial horizontal and vertical velocities of the ball. The meta-parameters used to initialize the learning make the ball drop in front of the first target. The cost function for the behavior level is

$$c = \sum_{i \in \{x, z\}} \dot{b}_i^2 + (b_x - s_x)^2,$$

where \dot{b}_i are the initial velocities, b_x is the impact position and s_x the desired impact position. The state corresponds to the three targets and is determined by the higher level. Figure 5.19 illustrates how the player learns to throw more accurately while playing. Figure 5.20 illustrates how learning to perform the lower level actions more reliably enables the player to perform better in the game.

Evaluation on a real BioRob

We employ a BioRob to throw balls in a catapult like fashion. The arm is approximately 0.75m long, and it can reach 1.55m above the ground. The targets are located at a distance of 2.5m from the robot at a height

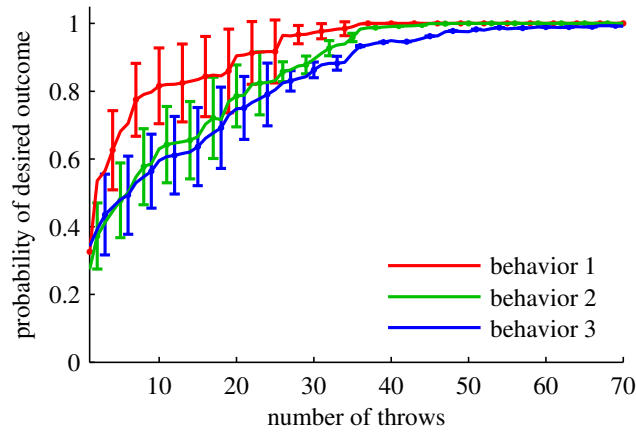


Figure 5.19: This figure illustrates the transition probabilities of the three behaviors to their associated outcome in simulation. For example, the red line indicates the probability of gaining one point when throwing at target 1. After approximately 50 throws the player has improved his accuracy level such that he always hits the desired target. The plots are averaged over 10 runs with the error-bars indicating standard deviations.




Figure 5.20: This figure illustrates the improvement of the player over the number of games in simulation. Due to the large penalty for busting the framework always uses a safe strategy. Already after five completed games the player reaches almost always the maximum possible score of 10. As the number of throws is not punished there are initially many throws that miss the target. After 7 games the number of throws has converged to 4, which is the minimum required number. The plots are averaged over 10 runs with the error-bars indicating standard deviations.




Figure 5.21: This figure illustrates the setup of the robot evaluation.

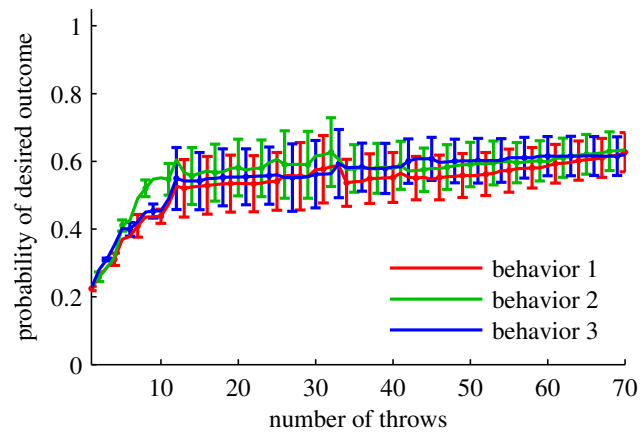


Figure 5.22: This figure illustrates the transition probabilities of the three behaviors to their associated outcome like in Figure 5.19. The skill improves a lot in the first 15 throws after that the improvement levels off. Initially behavior 2, associated with target 2 (which lies in the center) is most likely to succeed. The success rate of 60% corresponds to the level of reproducibility of our setup. The framework manages to handle this large uncertainty by choosing to “stand” early on. The plots are averaged over 4 runs with the error-bars indicating standard deviations.

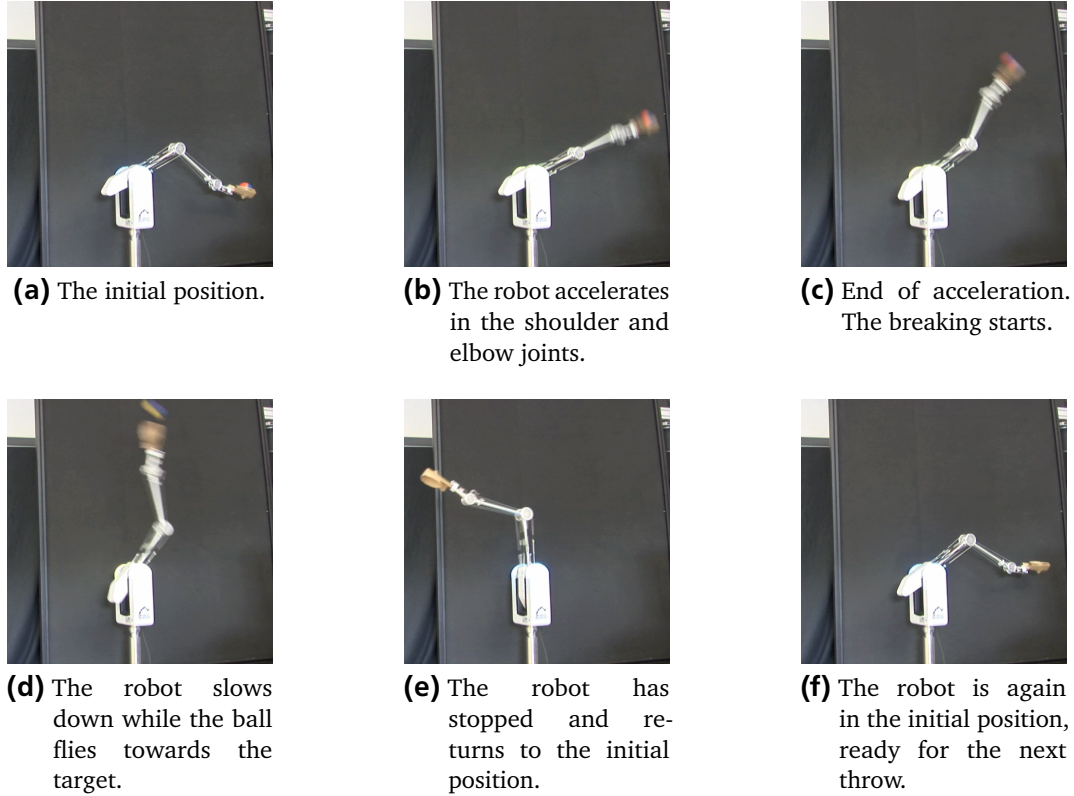


Figure 5.23: These frames illustrate one throwing motion with the BioRob.

of $0.9m$, $1.2m$, and $1.5m$ respectively. The ball is placed in a funnel-shaped receptacle. In this setup, the initial horizontal and vertical velocities of the ball cannot directly be set. Instead, the meta-parameters are defined as the duration and amount of acceleration for two joints that are in the throwing plane. The robot starts in a fixed initial position, accelerates the two joints according to the meta-parameter indicating the magnitude, and accelerates in the opposite direction after the time determined by the other meta-parameter in order to break. Finally the robot returns to the initial position. See Figure 5.23 for an illustration of one throwing motion. The state corresponds to the three targets and is determined by the higher level. The outcome of the throw is observed by a vision system.

Executing the throw with identical parameters will only land at the same target in approximately 60% of the throws, due to the high velocities involved and small differences in putting the ball in the holder. Thus, the algorithm has to deal with large uncertainties. The cost function for the behavior level is

$$c = \sum_{i \in \{1,2\}} \ddot{\theta}_i^2 + t_{\text{acc}}^2 + (b_x - s_x)^2,$$

where $\ddot{\theta}_i$ is the acceleration magnitude, t_{acc} the acceleration duration, b_x is the impact position and s_x the desired impact position. The setup makes it intentionally hard to hit target 3. The target can only be hit with a very restricted set of parameters. For targets 1 and 2 increasing the amount of acceleration or the duration will result in a higher hit. Target 3 is at the limit where higher accelerations or longer durations will lead to a throw in a downward direction with a high velocity.

The typical behavior of one complete experiment is as follows: At the beginning the robot explores in a very large area and stands as soon as it reaches a score of 8, 9, or 10. Due to the large punishment it is not willing to attempt to throw at 1 or 2 while having a large uncertainty, and, thus, a high chance of busting. Later on, it has learned that attempting to throw at 2 has a very low chance of ending up in 3 and hence will attempt to throw 2 points if the current score is 8. We setup the policy iteration to

favor behaviors with a higher number, if the values of the behaviors are identical. The first throws of a round will often be aimed at 3, even if the probability of hitting target 2 using this action is actually higher than hitting the associated target 3. Until 8 or more points have been accumulated, action 3 is safe (i.e., cannot lead to busting), does not entrain a punishment if missing or hitting a lower target, and has a large learning potential. Figure 5.22 illustrates how the robot learns to throw more accurately within the physical limits of the system.

5.4 Conclusion & Future Work

In this chapter, we have studied the problem of meta-parameter learning for motor primitives. It is an essential step towards applying motor primitives for learning complex motor skills in robotics more flexibly. We have discussed an appropriate reinforcement learning algorithm for mapping situations to meta-parameters.

We show that the necessary mapping from situation to meta-parameter can be learned using a Cost-regularized Kernel Regression (CrKR) while the parameters of the motor primitive can still be acquired through traditional approaches. The predictive variance of CrKR is used for exploration in on-policy meta-parameter reinforcement learning. We compare the resulting algorithm in a toy scenario to a policy gradient algorithm with a well-tuned policy representation and the reward-weighted regression. We show that our CrKR algorithm can significantly outperform these preceding methods. We also illustrate the advantages of our reinforcement learning approach over supervised learning approaches in this setting. To demonstrate the system in a complex scenario, we have chosen the *Around the Clock* dart throwing game, table tennis, and ball throwing implemented both on simulated and real robots. In these scenarios we show that our approach performs well in a wide variety of settings, i.e. on four different real robots (namely a Barrett WAM, a BioRob, the JST-ICORP/SARCOS CBi and a Kuka KR 6), with different cost functions (both with and without secondary objectives), and with different policies in conjunction with their associated meta-parameters.

In the ball throwing task, we have discussed first steps towards a supervisory layer that deals with sequencing different motor primitives. This supervisory layer is learned by an hierarchical reinforcement learning approach [Huber and Grunp, 1998, Barto and Mahadevan, 2003]. In this framework, the motor primitives with meta-parameter functions could also be seen as robotics counterpart of options [McGovern and Barto, 2001] or macro-actions [McGovern et al., 1997]. The presented approach needs to be extended to deal with different actions that do not share the same underlying parametrization. For example in a table tennis task the supervisory layer would decide between a forehand motor primitive and a backhand motor primitive, the spatial meta-parameter and the timing of the motor primitive would be adapted according to the incoming ball, and the motor primitive would generate the trajectory. Future work will require to automatically detect which parameters can serve as meta-parameters as well as to discovering new motor primitives.



6 Learning Prioritized Control of Motor Primitives

Many tasks in robotics can be decomposed into sub-tasks that are performed simultaneously. In many cases, these sub-tasks cannot all be achieved jointly and a prioritization of such sub-tasks is required to resolve this issue. In this chapter, we discuss a novel learning approach that allows to learn a prioritized control law built on a set of sub-tasks represented by motor primitives. The primitives are executed simultaneously but have different priorities. Primitives of higher priority can override the commands of the conflicting lower priority ones. The dominance structure of these primitives has a significant impact on the performance of the prioritized control law. We evaluate the proposed approach with a ball bouncing task on a Barrett WAM.

6.1 Introduction

When learning a new skill, it is often easier to practice the required sub-tasks separately and later on combine them to perform the task – instead of attempting to learn the complete skill as a whole. For example, in sports sub-tasks can often be trained separately. Individual skills required in the sport are trained in isolation to improve the overall performance, e.g., in volleyball a serve can be trained without playing the whole game.

Sub-tasks often have to be performed simultaneously and it is not always possible to completely fulfill all at once. Hence, the sub-tasks need to be prioritized. An intuitive example for this kind of prioritizing sub-tasks happens during a volleyball game: a player considers hitting the ball (and hence avoiding it touching the ground and his team loosing a point) more important than locating a team mate and playing the ball precisely to him. The player will attempt to fulfill both sub-tasks. If this is not possible it is often better to “safe” the ball with a high hit and hope that another player recovers it rather than immediately loosing a point.

In this chapter, we learn different sub-tasks that are represented by motor primitives that combined can perform a more complicated task. For doing so, we will stack controls corresponding to different primitives that represent movements in task space. These primitives are assigned different priorities and the motor commands corresponding to primitives with higher priorities can override the motor commands of lower priority ones. The proposed approach is outlined in Section 6.1.1 and further developed in Section 6.2. We evaluate our approach with a ball-bouncing task (see Figure 6.1 and Section 6.3).

As the sub-tasks describe the movements in task space, we have to learn a control that is mapping to the robot joint space. Unfortunately, this mapping is not a well-defined function for many robots. For example, if the considered task space has less degrees of freedom than the robot, multiple solutions




Figure 6.1: This figure illustrates the ball-bouncing task on a Barrett WAM. The goal is to keep the ball bouncing on the racket.

are possible. This redundancy can be resolved by introducing a null-space control, i.e., a behavior that operates on the redundant degrees of freedom. Such a null-space control can for example pull the robot towards a rest posture [Peters and Schaal, 2008a], prevent getting close to joint limits [Chaumette and Marchand, 2001], avoid obstacles [Khatib, 1986] or singularities [Yoshikawa, 1985]. Computing the task space control often corresponds to an optimization problem, that can for example be solved by a gradient based approach. A well known approach is the pseudo-inverse solution [Khatib, 1986, Peters and Schaal, 2008a]. An alternative is to learn an operational space control law that implicitly includes the null-space behavior [Peters and Schaal, 2008b]. Once learned, it corresponds to a unique mapping from desired actions in operational space to required actions in joint space.

The problem studied in this chapter is related to hierarchical control problems as discussed in [Findeisen et al., 1980]. Using prioritized primitives in classical control has been explored in [Sentis and Khatib, 2005] by using analytical projections into the null-space. In this chapter, we propose a learning approach that does not require complete knowledge of the system, the constraints, and the task. In the reinforcement learning community, the compositions of options (i.e., concurrent options), which is related to the concurrent execution of primitives, has been studied by Precup et al. [1998]. Learning null-space control has been explored in [Towell et al., 2010]. In contrast, we do not attempt to recover the implicit null-space policy but build a hierarchical operational space control law from user demonstrated primitives.

6.1.1 Proposed Approach

Based on the observation that many tasks can be described as a superposition of sub-tasks, we want to have a set of controls that can be executed simultaneously. As a representation for the sub-tasks, we chose the dynamical systems motor primitives, which are discussed in more detail in Chapter 3. Such primitives are well suited as representation for the sub-tasks as they ensure the stability of the movement generation. They are invariant under transformations of initial position and velocity, the final position and velocity, the duration as well as the movement amplitude.

In this chapter, these primitives are described in different task spaces, e.g., in the form

$$\ddot{\mathbf{x}}_i = \pi_i(\mathbf{x}_i, \dot{\mathbf{x}}_i, z)$$

where z denotes a share canonical system while \mathbf{x}_i are positions in task-space i . For example, if we have a primitive “move end-effector up and down” its task space would correspond to the Cartesian position indicating the height (as well as the corresponding velocities and accelerations) but not include the sideways movement or the orientation of the end-effector. The dynamical systems motor primitives are well suited to represent different kinds of vertical movements starting and ending at various states and of different duration.

These primitives are prioritized such that

$$i \succeq i - 1,$$

which reads a “task i dominates task $i - 1$ ”. If both sub-tasks can be fulfilled at the same time, our system will do so – but if this should not be possible, sub-task i will be fulfilled at the expense of sub-task $i - 1$. We attempt to reproduce a complex task that consists of several sub-tasks, represented by motor primitives,

$$\{\pi_1, \pi_2, \dots, \pi_N\}$$

that are concurrently executed at the same time following the prioritization scheme

$$N \succeq N - 1 \succeq \dots \succeq 2 \succeq 1.$$

This approach requires a prioritized control law that composes the motor command out of the primitives π_i , i.e.,

$$\mathbf{u} = \mathbf{f}(\pi_1, \pi_2, \dots, \pi_N, \mathbf{q}, \dot{\mathbf{q}})$$

where $\mathbf{q}, \dot{\mathbf{q}}$ are the joint position and joint velocity, \mathbf{u} are the generated motor commands (torques or accelerations).

We try to acquire the prioritized control law in three steps, which we will illustrate with the ball-bouncing task:

1. We observe $\ddot{\mathbf{x}}_i(t), \dot{\mathbf{x}}_i(t), \mathbf{x}_i(t)$ individually for each of the primitives that will be used for the task. For the ball-bouncing example, we may have the following sub-tasks: “move under the ball”, “hit the ball”, and “change racket orientation”. The training data is collected by executing only one primitive at a time without considering the global strategy, e.g., for the “change racket orientation” primitive by keeping the position of the racket fixed and only changing its orientation without a ball being present. This training data is used to acquire the task by imitation learning under the assumption that these tasks did not need to overrule each other in the demonstration (Sect. 6.2).
2. We enumerate all possible dominance structures and learn a prioritized control law for each dominance list that fusions the motor primitives. For the three ball-bouncing primitives there are six possible orders, as listed in Table 6.1.
3. We choose the most successful of these approaches. The activation and adaptation of the different primitives is handled by a strategy layer (Section 6.3.2). In the ball-bouncing task, we evaluate how long each of the prioritized control laws keeps the ball in the air and pick the best performing one (Section 6.3.3).

Clearly, enumerating all possible dominance structures only works for small systems (as the number of possibilities grows with $n!$, i.e., exponentially fast).

6.2 Learning the Prioritized Control Law

By learning the prioritized control, we want to obtain a control law

$$\mathbf{u} = \ddot{\mathbf{q}} = \mathbf{f}(\pi_1, \pi_2, \dots, \pi_N, \mathbf{q}, \dot{\mathbf{q}}),$$

i.e., we want to obtain the required control \mathbf{u} that executes the primitives $\pi_1, \pi_2, \dots, \pi_N$. Here, the controls correspond to the joint accelerations $\ddot{\mathbf{q}}$. The required joint accelerations not only depend on the primitives but also on the current state of the robot, i.e., the joint positions \mathbf{q} and joint velocities $\dot{\mathbf{q}}$. Any control law can be represented locally as a linear control law. In our setting, these linear control laws can be represented as

$$\mathbf{u} = \begin{bmatrix} \ddot{\mathbf{x}}_i \\ \dot{\mathbf{q}} \\ \mathbf{q} \end{bmatrix}^T \boldsymbol{\theta} = \boldsymbol{\phi}^T \boldsymbol{\theta},$$

where $\boldsymbol{\theta}$ are the parameters we want to learn and $\boldsymbol{\phi} = \begin{bmatrix} \ddot{\mathbf{x}}_i & \dot{\mathbf{q}} & \mathbf{q} \end{bmatrix}$ acts as features. Often the actions of the primitive $\ddot{\mathbf{x}}_i$ can be achieved in multiple different ways due to the redundancies in the robot degrees of freedom. To ensure consistency, a null-space control is introduced. The null-space control can, for example, be defined to pull the robot towards a rest posture \mathbf{q}_0 , resulting in the null-space control

$$\mathbf{u}_0 = -\mathbf{K}_D \dot{\mathbf{q}} - \mathbf{K}_P (\mathbf{q} - \mathbf{q}_0),$$

where \mathbf{K}_D and \mathbf{K}_P are gains for the velocities and positions respectively.

To learn the prioritized control law, we try to generalize the learning operational space control approach from [Peters and Schaal, 2008b] to a hierarchical control approach [Sentis and Khatib, 2005, Peters and Schaal, 2008a].

6.2.1 Single Primitive Control Law

A straightforward approach to learn the motor commands \mathbf{u} , represented by the linear model $\mathbf{u} = \boldsymbol{\phi}^T \boldsymbol{\theta}$, is using linear regression. This approach minimizes the squared error

$$E^2 = \sum_{t=1}^T (\mathbf{u}_t^{\text{ref}} - \boldsymbol{\phi}_t^T \boldsymbol{\theta})^2$$

between the demonstrated control of the primitive u_t^{ref} and the recovered linear policy $\mathbf{u}_t = \boldsymbol{\phi}_t^T \boldsymbol{\theta}$, where T is the number of samples. The parameters minimizing this error are

$$\boldsymbol{\theta} = (\boldsymbol{\Phi}^T \boldsymbol{\Phi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Phi}^T \mathbf{U}, \quad (6.1)$$

with $\boldsymbol{\Phi}$ and \mathbf{U} containing the values of the demonstrated $\boldsymbol{\phi}$ and \mathbf{u} for all time-steps t respectively, and a ridge factor λ . If the task space and the joint-space coincide, the controls $\mathbf{u} = \ddot{\mathbf{q}}$ are identical to the action of the primitive $\ddot{\mathbf{x}}_i$. We also know that *locally* any control law that can be learned from data is a viable control law [Peters and Schaal, 2008b]. The error with respect to the training data is minimized, *however*, if the training data is not consistent, the plain linear regression will average the motor commands, which is unlikely to fulfill the actions of the primitive.

In order to enforce consistency, the learning approach has to resolve the redundancy and incorporate the null-space control. We can achieve this by using the program

$$\begin{aligned} \min_{\mathbf{u}} J &= (\mathbf{u} - \mathbf{u}_0)^T \mathbf{N} (\mathbf{u} - \mathbf{u}_0) \\ \text{s.t. } \ddot{\mathbf{x}} &= \pi(\mathbf{x}, \dot{\mathbf{x}}, z) \end{aligned} \quad (6.2)$$

as discussed in [Peters and Schaal, 2008a]. Here the cost J is defined as the weighted squared difference of the control \mathbf{u} and the null-space control \mathbf{u}_0 , where the metric \mathbf{N} is a positive semi-definite matrix. The idea is to find controls \mathbf{u} that are as close as possible to the null-space control \mathbf{u}_0 while still fulfilling the constraints of the primitive π . This program can also be solved as discussed in [Peters and Schaal, 2008b]. Briefly speaking, the regression in Equation (6.1) can be made consistent by weighting down the error by weights w_t and hence obtaining

$$\boldsymbol{\theta} = (\boldsymbol{\Phi}^T \mathbf{W} \boldsymbol{\Phi} + \lambda \mathbf{I})^{-1} \boldsymbol{\Phi}^T \mathbf{W} \mathbf{U} \quad (6.3)$$

with $\mathbf{W} = \text{diag}(w_1, \dots, w_{Tn})$ for T samples. This approach works well for linear models and can be gotten to work with multiple locally linear control laws. Nevertheless, it maximizes a reward instead of minimizing a cost. The cost J can be transformed into weights w_t by passing it through an exponential function

$$w_t = \exp(-\alpha \tilde{\mathbf{u}}_t^T \mathbf{N} \tilde{\mathbf{u}}_t),$$

where $\tilde{\mathbf{u}}_t = (\mathbf{u}_t - \mathbf{u}_0)$. The scaling factor α acts as a monotonic transformation that does not affect the optimal solution but can increase the efficiency of the learning algorithm.

Using the Woodbury formula [Welling, 2010] Equation (6.3) can be transformed into

$$\mathbf{u} = \boldsymbol{\phi}(x)^T \boldsymbol{\Phi}^T (\boldsymbol{\Phi} \boldsymbol{\Phi}^T + \mathbf{W}_U)^{-1} \mathbf{U} \quad (6.4)$$

with $\mathbf{W}_U = \text{diag}(\tilde{\mathbf{u}}_1^T \mathbf{N} \tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_n^T \mathbf{N} \tilde{\mathbf{u}}_n)$. By introducing the kernels $\mathbf{k}(\mathbf{s}) = \boldsymbol{\phi}(\mathbf{s})^T \boldsymbol{\Phi}^T$ and $\mathbf{K} = \boldsymbol{\Phi} \boldsymbol{\Phi}^T$ we obtain

$$\mathbf{u} = \mathbf{k}(\mathbf{s})^T (\mathbf{K} + \mathbf{W}_U)^{-1} \mathbf{U},$$

which is related to the kernel regression [Bishop, 2006] and corresponds to the Cost-regularized Kernel Regression introduced in Chapter 5. This kernelized form of Equation (6.4) overcomes the limitations of the linear model at a cost of higher computational complexity.

Algorithm 6.1 Learning the Prioritized Control Law

```
define null-space control  $\mathbf{u}_0$ , metric  $\mathbf{N}$ , scaling factor  $\alpha$ 

collect controls  $\mathbf{u}_{i,t}$  and features  $\phi_{i,t}$  for all primitives  $i \in \{1, \dots, N\}$  and all time-steps  $t \in \{1, \dots, T\}$ 
separately

for primitives  $i = 1 \dots N$  ( $N$ : highest priority) do
  for time-steps  $t = 1 \dots T$  do
    calculate offset controls
     $\hat{\mathbf{u}}_{i,t} = \mathbf{u}_{i,t} - \sum_{j=1}^{i-1} \phi_{i,t}^T \boldsymbol{\theta}_j - \mathbf{u}_{0,t}$ 

    calculate weights  $\hat{w}_{i,t} = \exp(-\alpha \hat{\mathbf{u}}_{i,t}^T \mathbf{N} \hat{\mathbf{u}}_{i,t})$ 
  end for

  build control matrix  $\hat{\mathbf{U}}_i$  containing  $\hat{\mathbf{u}}_{i,1} \dots \hat{\mathbf{u}}_{i,T}$ 
  build feature matrix  $\hat{\Phi}_i$  containing  $\phi_{i,1} \dots \phi_{i,T}$ 
  build weight matrix  $\hat{\mathbf{W}}_i = \text{diag}(\hat{w}_{i,1}, \dots, \hat{w}_{i,T})$ 
  calculate parameters
   $\boldsymbol{\theta}_i = (\hat{\Phi}_i^T \hat{\mathbf{W}}_i \hat{\Phi}_i + \lambda \mathbf{I})^{-1} \hat{\Phi}_i^T \hat{\mathbf{W}}_i \hat{\mathbf{U}}_i$ 
end for
end for
```

6.2.2 Prioritized Primitives Control Law

In the previous section, we have described how the control law for a single primitive can be learned. To generalize this approach to multiple primitives with different priorities, we want a control law that always fulfills the primitive with the highest priority and follows the remaining primitives as much as possible according to their place in the hierarchy. Our idea is to represent the higher priority control laws as correction term with respect to the lower priority primitives. The control of the primitive with the lowest priority is learned first. This control is subsequently considered to be a baseline and the primitives of higher priority only learn the difference to this baseline control. The change between the motor commands resulting from primitives of lower priority is minimized. The approach is reminiscent of online passive-aggressive algorithms [Crammer et al., 2006]. Hence, control laws of higher priority primitives only learn the offset between their desired behavior and the behavior of the lower priority primitives. This structure allows them to override the actions of the primitives of lesser priority and, therefore, add more detailed control in the regions of the state space they are concerned with. The combined control of all primitives is

$$\mathbf{u} = \mathbf{u}_0 + \sum_{n=1}^N \Delta \mathbf{u}_n,$$

where \mathbf{u}_0 is the null-space control and $\Delta \mathbf{u}_n$ are the offset controls of the N primitives.

Such control laws can be expressed by changing the program in Equation (6.2) to

$$\begin{aligned} \min_{\mathbf{u}_i} J &= \left(\mathbf{u}_i - \sum_{j=1}^{i-1} \Delta \mathbf{u}_j - \mathbf{u}_0 \right)^T \mathbf{N} \left(\mathbf{u}_i - \sum_{j=1}^{i-1} \Delta \mathbf{u}_j - \mathbf{u}_0 \right) \\ \text{s.t. } \ddot{\mathbf{x}}_i &= \pi_i(\mathbf{x}_i, \dot{\mathbf{x}}_i, \mathbf{z}), \end{aligned}$$

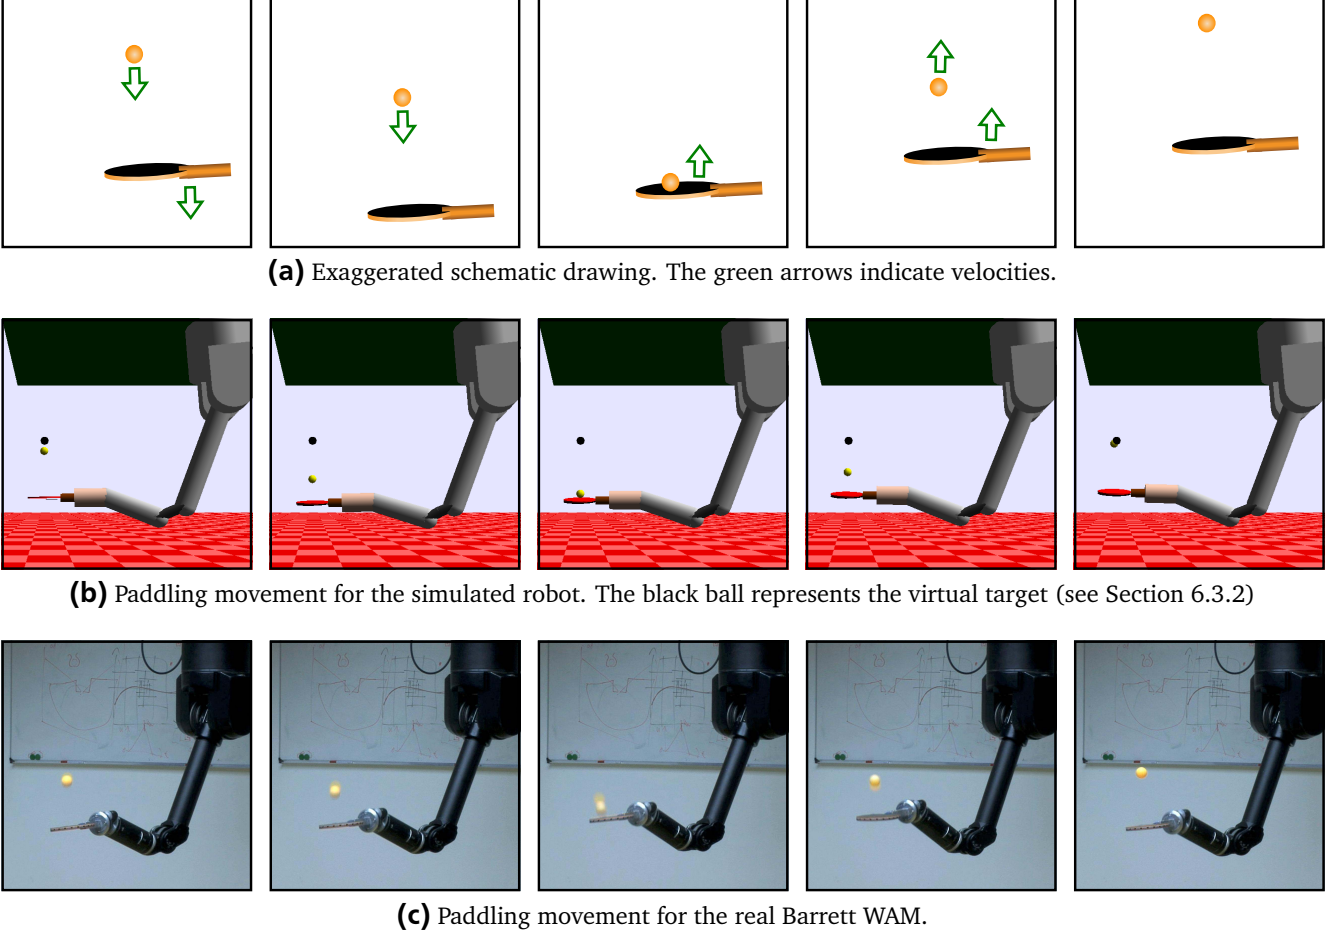


Figure 6.2: This figure illustrates a possible sequence of bouncing the ball on the racket in a schematic drawing, in simulation, and on the real robot.

where the primitives need to be learned in the increasing order of their priority, the primitive with the lowest priority is learned first, the primitive with the highest priority is learned last. The regression in Equation (6.3) changes to

$$\theta_i = (\Phi_i^T \hat{\mathbf{W}}_i \Phi_i + \lambda \mathbf{I})^{-1} \Phi_i^T \hat{\mathbf{W}}_i \hat{\mathbf{U}}_i,$$

where $\hat{\mathbf{U}}_i$ contains the offset controls $\hat{\mathbf{u}}_{i,t} = \mathbf{u}_{i,t} - \sum_{j=1}^{i-1} \Delta \mathbf{u}_{j,t} - \mathbf{u}_{0,t}$ for all time-steps t , where $\Delta \mathbf{u}_{j,t} = \phi_{i,t}^T \theta_j$. The weighting matrix $\hat{\mathbf{W}}_i$ now has the weights $\hat{w}_t = \exp(-\alpha \hat{\mathbf{u}}_{i,t}^T \mathbf{N} \hat{\mathbf{u}}_{i,t})$ on its diagonal and matrix $\hat{\mathbf{U}}_i$ contains offset controls $\hat{\mathbf{u}}_{i,t}$. The kernelized form of the prioritized control law can be obtained analogously. The complete approach is summarized in Algorithm 6.1.

6.3 Evaluation: Ball-Bouncing

In order to evaluate the proposed prioritized control approach, we chose a ball bouncing task. We describe the task in Section 6.3.1, explain a possible higher level strategy in Section 6.3.2, and discuss how the proposed framework can be applied in Section 6.3.3.




Figure 6.3: This figure illustrates the employed strategy for bouncing the ball on the racket. The highest point of the ball trajectory is supposed to coincide with the red target. The racket is always hitting the ball in a fixed height, i.e., the hitting plane. The strategy is to play the ball in a way that the next intersection with the hitting plane is directly below the target and the maximum height of the ball trajectory corresponds to the height of the target. If the bounce works exactly as planned, the ball needs to be hit only once to return to a strictly vertical bouncing behavior.

6.3.1 Task Description

The goal of the task is to bounce a table tennis ball above a racket. The racket is held in the player's hand, or in our case attached to the end-effector of the robot. The ball is supposed to be kept bouncing on the racket. A possible movement is illustrated in Figure 6.2.

It is desirable to stabilize the bouncing movement to a strictly vertical bounce, hence, avoiding the need of the player to move a lot in space and, thus, leaving the work space of the robot. The hitting height is a trade-off between having more time until the next hit at the expense of the next hitting position possibly being further away. The task can be sub-divided into three intuitive primitives: hitting the ball upward, moving the racket under the ball before hitting, and changing the orientation of the racket to move the ball to a desired location. A possible strategy is outlined in the next section.

The ball is tracked using a stereo vision setup and its positions and velocities are estimated by a Kalman filter. To initialize the ball-bouncing task, the ball is thrown towards the racket.

6.3.2 Bouncing Strategy

The strategy employed to achieve the desired bouncing behavior is based on an imagined target that indicates the desired bouncing height. This target is above the default posture of the racket. The top point of the ball trajectory is supposed to hit this target, and the stable behavior should be a strictly vertical bounce. This behavior can be achieved by defining a hitting plane, i.e., a height at which the ball is always hit (which corresponds to the default posture of the racket). On this hitting plane, the ball is always hit in a manner that the top point of its trajectory corresponds to the height of the target and the next intersection of the ball trajectory with the hitting plane is directly under the target. See Figure 6.3 for an illustration.

To achieve this desired ball behavior, the racket is always moved to the intersection point of the ball trajectory and the hitting plane. By choosing the hitting velocity and the orientation of the racket, the velocity and direction of the ball after being hit can be changed. The required hitting velocity and orientation are calculated using a model of the ball and the racket. The ball is modeled as a point mass that moves according to the ballistic flight equations. For the relatively low speeds and small distances air resistance is negligible. The contact with the racket is modeled as a reflection with a restitution factor.

Dominance Structure	Number of Hits	
	in Simulation	on Real Robot
single model	5.70 ± 0.73	1.10 ± 0.99
hit \succeq move \succeq orient	11.35 ± 2.16	2.30 ± 0.67
hit \succeq orient \succeq move	10.85 ± 1.46	1.70 ± 0.95
move \succeq hit \succeq orient	9.05 ± 0.76	1.40 ± 0.70
move \succeq orient \succeq hit	7.75 ± 1.48	1.40 ± 0.84
orient \succeq hit \succeq move	5.90 ± 0.85	1.30 ± 0.67
orient \succeq move \succeq hit	5.35 ± 0.49	1.30 ± 0.48

Table 6.1: This table shows the suitability of the possible dominance structures (mean \pm std). The “hit the ball” primitive clearly is the dominant one, followed by the “move under the ball” primitive. The prioritized control laws work significantly better than a single model learned using the combined training data of the three primitives. Preliminary results on the real robot confirm this ordering.

Using this strategy the ball can be brought back to a strictly vertical bouncing behavior with a single hit. However, this method requires the knowledge of the ball position and velocity, as well as a model of the ball behavior. An alternative strategy that stabilizes the behavior in a completely open loop behavior by employing a slightly concave paddle shape has been suggested in [Reist and D’Andrea, 2009]. A method similar to the proposed strategy has been employed by Kulchenko [2011] and Müller et al. [2011]. Buehler et al. [1994] proposed the mirror law for this task. The ball bouncing task has also been employed to study how humans stabilize a rhythmic task [Schaal et al., 1996].

6.3.3 Learning Results

As discussed in Section 6.3.1, the task can be described by three primitives: “move under the ball”, “hit the ball”, and “change racket orientation”. Training data is collected in the relevant state space independently for each primitive. For doing so, the parameters corresponding to the other primitives are kept fixed and variants of the primitive are hence executed from various different starting positions. The primitive “move under the ball” corresponds to movements in the horizontal plane, the primitive “hit the ball” to up and down movements, and the primitive “change racket orientation” only changes the orientation of the end-effector. We collected 30 seconds of training data for each primitive, corresponding to approximately 60 bounces.

Having only three primitives allows it to enumerate all six possible dominance structures, to learn the corresponding prioritized control law, and to evaluate the controller. As intuitive quality measure we counted the number of bounces until the robot missed, either due to imprecise control or due to the ball being outside of the safely reachable work-space.

Table 6.1 illustrates the resulting dominance structures. The most relevant primitive is the “hit the ball” primitive, followed by the “move under the ball” primitive. In the table it is clearly visible that inverting the order of two neighboring primitives that are in the preferred dominance order always results in a lower number of hits. Compared to a single model, that was trained using the combined training data of the three primitives, all but two prioritized control laws work significantly better. The ordering may appear slightly counter-intuitive as moving under the ball seems to be the most important primitive in order to keep the ball in the air, allowing for later corrections. However, the robot has a fixed base position and the ball moves quickly out of the safely reachable work-space, resulting in a low number of hits. Additionally, the default position of the racket is almost vertical, hence covering a fairly large area of the horizontal plane resulting in robustness with respect to errors in this primitive.

6.4 Conclusion

In this chapter, we have presented a prioritized control learning approach that is based on the superposition of movement primitives. We have introduced a novel framework for learning prioritized control. The controls of the lower priority primitives are fulfilled as long as they lay in the null-space of the higher priority ones and get overridden otherwise. As representation for the primitives, we employ the dynamical systems motor primitives [Ijspeert et al., 2002a, Schaal et al., 2007], which yield controls in the form of desired accelerations. These primitives are executed separately to collect training data. Local linear models are trained using a weighted regression technique incorporating the various possible dominance structures. In the presented ball bouncing task, the movement is restricted to a space where the controls are approximately linear. Hence, a single linear model per primitive was sufficient. This limitation can be overcome by either considering multiple local linear models [Peters and Schaal, 2008b] or by kernelizing the weighted regression, as described in Section 6.2.1 and 6.2.2.

The dominance structure of the task was determined by testing all possible structures exhaustively. Intuitively, the lower priority primitives represent a global behavior and the high priority primitives represent specialized corrections, hence overriding the lower priority controls. In most cases, the resulting prioritized control works significantly better than a single layer one that was trained with the combined training data of all primitives. As illustrated by the evaluations, the dominance structure can have a significant influence on the global success of the prioritized control. Enumerating all possible dominance structures is factorial in the number of primitives and hence unfeasible in practice for more than four primitives. In this case, smarter search strategies are needed.

The success of the different dominance structures not only depends on the task but also on the employed strategy of activating and adapting the different primitives. An interesting area for future research could be to jointly learn the prioritized control and the strategy.

The presented approach has been evaluated both in simulation and on a real Barrett WAM and we have demonstrated that our novel approach can successfully learn a ball-bouncing task.



7 Conclusion

In this thesis, we have discussed reinforcement learning approaches for motor skills represented by motor primitives. In the next section, we provide an overview of the key contributions in this thesis and then we discuss possible directions for extending the presented research.

7.1 Summary

The contributions of this paper are for the state-of-the-art in both machine learning and in robotics.

7.1.1 Algorithms

In this thesis, we have proposed a framework of policy search based on reward-weighted imitation. The resulting EM-inspired algorithms are applicable both to parametrized and non-parametric policies. The policy search algorithms presented in this thesis perform local optimization which results in fast convergence but also poses the risk of converging to bad local optima. For all the presented application scenarios, good initial policies are available, which mitigates the problem of bad local optima and renders the approaches applicable in practice.

In Section 1.2.1, we have discussed requirements for robotic reinforcement learning approaches, i.e., they should avoid damage to the robot, and should be fast, both in terms of convergence and computation time. Having a sample-efficient algorithm, only very few open parameters, and the ability to incorporate prior knowledge all are essential for fast convergence. In the following we will discuss how the proposed algorithms meet these requirements.

Policy learning by Weighting Exploration with the Returns (PoWER)

PoWER (Chapter 4) is an EM-inspired policy search algorithm relying on a parametrized policy and structured exploration in the parameter space. The algorithm is particularly suitable for motor primitives (Chapter 3). We introduce a framework of reward weighted imitation that yields several well known policy search algorithms: episodic REINFORCE [Williams, 1992], the policy gradient theorem [Sutton et al., 1999], episodic natural actor critic [Peters and Schaal, 2008c], a generalization of the reward-weighted regression [Peters and Schaal, 2008b].

PoWER is unlikely to damage the robot as it only explores locally and usually is initialized by an initial policy. Using importance sampling, the policy can be updated after each rollout, hence it is sample efficient. Prior knowledge can be incorporated both in the form of the initial policy and by defining the policy structure via the parametrization. The main open parameter is the exploration magnitude. This parameter only needs to be set once initially and can then be automatically adjusted during the learning process as discussed in Section 4.A.3. The policy structure is also an open parameter. The dynamical systems motor primitives essentially only have a single open parameter, i.e., the number of basis functions as an open parameter which corresponds to the amount of detail that the policy can capture. The bottleneck of the update calculation is a matrix inversion that can be avoided by making some additional independence assumptions that work well in conjunction with the dynamical systems motor primitives.

PoWER outperformed various other policy search approaches (i.e., finite difference gradients, episodic REINFORCE, ‘vanilla’ policy gradients with optimal baselines, episodic natural actor critic, and episodic reward-weighted regression) on benchmarks and robotic tasks. Our approach has inspired follow-up work in other contexts, for example [Vlassis et al., 2009, Kormushev et al., 2010]. Theodorou et al. [2010]

have derived algorithms based on the path integral approach that are very similar to PoWER and have also been successfully employed for robotic tasks [Buchli et al., 2011, Kalakrishnan et al., 2011, Pastor et al., 2011, Stulp et al., 2011, Tamosiunaite et al., 2011].

Cost-regularized Kernel Regression (CrKR)

CrKR (Chapter 5) is an EM-inspired policy search algorithm employing a non-parametric policy representation. CrKR is based on a kernelized version of the episodic reward-weighted regression [Peters and Schaal, 2008b, Chapter 4]. Similar to Gaussian process regression, CrKR yields a predictive variance that can be employed to guide the exploration. In this thesis it is employed to generalize movements to new situations. This type of learning relies on a policy representation that has a number of meta-parameters that allow to generalize the movement globally while retaining the details of the motion. The motor primitives (Chapter 3) inherently have six modification parameters (the initial position \mathbf{x}_1^0 , the initial velocity \mathbf{x}_2^0 , the goal \mathbf{g} , the goal velocity $\dot{\mathbf{g}}$, the amplitude \mathbf{A} , and the duration T) which serve as meta-parameters.

CrKR explores locally similar to PoWER. It can include an initial policy in the form of a prior or initial samples. The policy is updated after each rollout. Instead of fixing a parametrized policy, the kernel hyper-parameters need to be determined, which results in a more flexible policy representation as demonstrated in comparisons with finite difference gradients and episodic reward-weighted regression. The hyper-parameters can be estimated from initial demonstrations. The update equation also includes a matrix inversion, which theoretically could be replaced by a sparse approximation. However, in the proposed setting the sample size was sufficiently small to allow full matrix inversion.

7.1.2 Applications

The applications focus mainly on dynamical systems motor primitives as policy representations but we also demonstrated the approaches using various different parametrization. All approaches have been extensively evaluated on benchmark tasks as well as with simulated and real robots, namely a Barrett WAM, a BioRob, the JST-ICORP/SARCOS CBi and a Kuka KR 6.

We successfully learned the single movements of Tetherball Target Hitting, Casting, Underactuated Swing-Up, and Ball-in-a-Cup using imitation learning and subsequent reinforcement learning. Compared to alternative approaches the proposed algorithm consistently converged to a better solution in less rollouts. In the Ball-in-a-Cup task the number of optimized parameters would result in an unrealistic amount of rollouts for gradient based policy search approaches.

We learned to generalize the motor skills of dart throwing and table tennis strikes. In contrast to previous work on generalizing dynamical systems motor primitives, we employ reinforcement learning to discover a non-intuitive mapping between states and actions. We compared the proposed non-parametric approach to parametric ones and showed that it is hard to find a good parametrization in this setting.

We have demonstrated initial steps towards hierarchical learning with a ball target throwing task. In this task we learn how to hit targets while keeping in mind a higher level strategy. Finally, we employed a ball-bouncing task to explore first ideas towards learning to perform a task based on the concurrent execution of several motor primitives.

7.2 Open Problems

In this thesis, we have contributed to the state of art of autonomous acquisition of motor skill by robots by developing approaches for learning motor primitives via reinforcement learning in robotics. In the following, we briefly discuss possible extensions of the presented work and future lines of research.

The presented learning approaches have been applied to episodic tasks that correspond to sports and games. We present several ideas to render them applicable to a wider variety of tasks and how to overcome some of their current limitations.

Rhythmic Motor Tasks

This thesis focused on discrete motor tasks in an episodic setting. However, there are also many interesting rhythmic motor task that could be learned, e.g., the ball paddling movement, briefly discussed in Section 4.4.2, could be optimized to consume less energy. Several periods of such rhythmic tasks could be grouped into rollouts and the presented algorithms would be directly applicable. However, it is not obvious how such tasks could be learned online without frequent restarts. Deciding when complete restarts are required is an interesting challenge. Rhythmic tasks often require a start-up phase until transitioning into a stable rhythmic behavior. Similarly, even if a stable behavior with a high reward is achieved after online learning and exploration, it is not obvious to ensure that this state is directly reachable.

Compliance

The presented policy parametrizations are all based on positions, velocities and accelerations in joint or task space. Force controlled robots additionally allow to adapt their compliance and stiffness during the movement. When humans learn a new motor skill, they are often fairly tense initially but soon start to figure out during which parts of the movement they can relax their muscles. Decreasing the stiffness of the robot renders direct interactions with humans and its environment safer and has the additional benefit of being more energy efficient. Hence, the robot needs to remain compliant unless the specific part of the movement requires higher forces or precise positioning. Extending the dynamical systems motor primitives to include global or directional stiffness would render the presented algorithms applicable in this setting. A related approach has already been discussed by Kalakrishnan et al. [2011] for this problem.

Sparse Approximations

We employed CrKR (Chapter 5) to generalize motor skills to new situations. In the current implementation we rely on the fact that the sample size is typically small enough to compute the updates in real-time. In Chapter 6, we employed a related approach to learn prioritized control, which is straightforward to kernelize. However, only a few seconds of training data could be used for real-time computations due to the high sampling rate. The matrix inversion is the main bottleneck, but also calculating the kernel between the current state and all stored samples can pose problems. As an alternative a local approximation or a sparse approximation could be considered [Nguyen-Tuong and Peters, 2011]. For the local models, the state-space is divided into smaller local ones, hence reducing the number of training points. In a sparse approximation, only the most informative points are considered. The most challenging question for these approaches remains how to construct the local models or to decide which points to discard, include, or merge.

Continuous Adaptation

In Chapter 5, we evaluated tasks that only required a one-time adaptation to the situation. Especially for the table tennis task, continuously adapting the hitting movement to the ball's movement would render the task more reliable. Ball spin is hard to determine from the vision data before the ball bounces. However, due to acceleration limits of the robot, the hitting movement needs to be initiated before the bounce. Conceptually CrKR can handle this setting as well, however considering the whole state-space will result in significantly more samples, rendering real-time computations more challenging (see above), and also make generalizations more difficult due to the curse of dimensionality. Experiments need to

be done in order to evaluate the feasibility and to determine whether additional components such as dimensionality reduction or models would be beneficial.

Manipulation

Compared to the presented applications (related to games and sports), grasping and manipulation often require less dynamic movements but in contrast have a higher dimensional state space due to the more complicated interactions between the robot and the objects. Grasping and manipulation are essential building blocks for many household tasks and have to be able to generalize to novel objects and tools. For example, opening a door needs to take into account a wide variety of handles and would benefit from adaptive compliance levels (as discussed above) for approaching the handle, pushing the handle, and opening the door.

Using Models

In Chapter 2, we identified the use of models as one of the promising approaches to make reinforcement learning for robotics tractable. We mainly employed a model of the robot and environment for debugging purposes, before testing the algorithms on the real robot. The approaches indirectly employ a model of the robot by using, e.g., the inverse dynamics controller of the robot. Pre-training in simulation can be helpful if the models are sufficiently accurate. An interesting extension of the proposed approaches would be to employ mental rehearsal with learned models in between real robot rollouts.

Hierarchical Reinforcement Learning

The hierarchical reinforcement learning approach presented in Chapter 5 is very specific to the problem setting. A more general hierarchical reinforcement learning approach would need the ability to choose between different motor primitives, the ability to discover new options, i.e., decide when it is better to generate a new one than to generalize an old one, and the ability to learn the shape parameters jointly with the meta-parameters. In a table tennis scenario we could have a library of already learned motor primitives. The hierarchical reinforcement learning approach would need to pick the most appropriate primitive (a forehand, a backhand, a smash, etc.) from this library according to the incoming ball, the current state of the robot, as well as strategic considerations. The template primitives will have to be generalized according to the current situation, either based on a single primitive or by combining several primitives. An approach based on the mixture of motor primitives has been recently proposed by Mülling et al. [2010]. Especially at the beginning of the learning process it will often be advantageous to learn and add a new primitive to the library instead of generalizing an old one.

Sequencing Motor Primitives

Another interesting direction for research is learning to optimally transition between motor primitives, e.g., to change from a forehand to a backhand. The transition needs to take into account external constraints such as self-collision and decide whether there is sufficient time for a successful transition. This problem is highly related to motion blending [Kovar, 2004]. It would be very interesting to see how well the approaches developed by the computer graphics community transfer to robotics.

Superposition of Motor Primitives

In our daily lives we often perform several motor skills simultaneously like balancing a tray while walking or sidestepping and keeping balance while hitting a table tennis ball. In order to perform multiple motor skills simultaneously a system could for example overly, combine and prioritize motor primitives. First ideas have been explored in Chapter 6. Similar to inverse reinforcement learning, it would be interesting to try to recover a dominance structure of primitives from human demonstrations. Alternatively, reinforcement learning could be employed to determine the best ordering for a given reward.

Learning Layers Jointly

For the ball-bouncing task (Chapter 6), we assumed that both the motor primitives and the strategy layer are fixed. In this task, better performance could be achieved if not only the prioritized control would be adapted but the primitives and the strategy as well. For example, the primitives could be adapted to compensate for shortcomings of the learned control or the strategy could be adapted. A straightforward idea would be to run a reinforcement learning algorithm after the prioritized control has been learned. PoWER could be employed to learn the parameters of the underlying primitives. The approach presented in Chapter 5 is a first attempt to tackle this kind of hierarchical problems.

Multiple Strategies

Learning single motor skills with PoWER (Chapter 4) relied on a single initial demonstration. Combining several demonstrations can potentially provide a better understanding of the important features of the movement. However, using a weighted average of multiple demonstrations of different strategies is unlikely to succeed. Keeping multiple strategies in mind and figuring out which parts of the movement can be learned jointly and which parts have to be learned separately might lead to faster and more robust learning. Daniel et al. [2012] employ a hierarchical reinforcement learning framework to learn different strategies for a tetherball task jointly.

7.3 Publications

Excerpts of the research presented in this thesis have led to the following publications:

Journal Papers

- J. Kober and J. Peters. Reinforcement learning für motor-primitive. *Künstliche Intelligenz*, 09(3):38–40, 2009.
- J. Kober and J. Peters. Imitation and reinforcement learning - practical algorithms for motor primitive learning in robotics. *IEEE Robotics and Automation Magazine*, 17(2):55–62, 2010.
- J. Kober and J. Peters. Policy search for motor primitives in robotics. *Machine Learning*, 84(1-2):171–203, 2011.
- J. Kober, J. A. Bagnell, and J. Peters. Reinforcement learning in robotics: a survey. *International Journal of Robotics Research*, In Preparation, 2012a.
- J. Kober, A. Wilhelm, E. Oztop, and J. Peters. Reinforcement learning to adjust parametrized motor primitives to new situations. *Autonomous Robots*, Under Review, 2012b.
- K. Mülling, J. Kober, and J. Peters. A biomimetic approach to robot table tennis. *Adaptive Behavior*, 9(5):359–376, 2011.
- J. Peters, J. Kober, and S. Schaal. Algorithmen zum automatischen erlernen von motor-fähigkeiten. *at - Automatisierungstechnik*, 58(12):688–694, 2010.

Conference and Workshop Papers

- A. Boularias, J. Kober, and J. Peters. Relative entropy inverse reinforcement learning. *Journal of Machine Learning Research - Proceedings Track*, 15:182–189, 2011.

-
- S. Chiappa, J. Kober, and J. Peters. Using bayesian dynamical systems for motion template libraries. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 297–304, 2009.
- J. Kober and J. Peters. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 849–856, 2008.
- J. Kober and J. Peters. Learning motor primitives for robotics. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2112–2118, 2009a.
- J. Kober and J. Peters. Learning new basic movements for robotics. In *Proceedings of Autonome Mobile Systeme (AMS)*, 2009b.
- J. Kober and J. Peters. Learning elementary movements jointly with a higher level task. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 338–343, 2011.
- J. Kober and J. Peters. Learning prioritized control of motor primitives. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Under Review, 2012.
- J. Kober, B. Mohler, and J. Peters. Learning perceptual coupling for motor primitives. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 834–839, 2008.
- J. Kober, K. Mülling, O. Krömer, C. H. Lampert, B. Schölkopf, and J. Peters. Movement templates for learning of hitting and batting. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 853–858, 2010a.
- J. Kober, E. Oztop, and J. Peters. Reinforcement learning to adjust robot movements to new situations. In *Proceedings of Robotics: Science and Systems Conference (R:SS)*, 2010b.
- J. Kober, E. Oztop, and J. Peters. Reinforcement learning to adjust robot movements to new situations. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), Best Paper Track*, pages 2650–2655, 2011.
- K. Mülling, J. Kober, and J. Peters. Learning table tennis with a mixture of motor primitives. In *Proceedings of the 10th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 411–416, 2010a.
- K. Mülling, J. Kober, and J. Peters. A biomimetic approach to robot table tennis. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1921–1926, 2010b.
- K. Mülling, J. Kober, and J. Peters. Simulating human table tennis with a biomimetic robot setup. In *Proceedings of From Animals to Animats 11 - International Conference on the Simulation of Adaptive Behavior (SAB)*, volume 6226 of *Lecture Notes in Computer Science*, pages 273–282, 2010c.
- J. Peters and J. Kober. Using reward-weighted imitation for robot reinforcement learning. In *Proceedings of the IEEE International Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 226–232, 2009.
- J. Peters, J. Kober, and D. Nguyen-Tuong. Policy learning - a unified perspective with applications in robotics. In *Proceedings of the European Workshop on Reinforcement Learning (EWRL)*, 2008.
- J. Peters, K. Mülling, J. Kober, D. Nguyen-Tuong, and O. Kroemer. Towards motor skill learning for robotics. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2010.

Book Chapters

- J. Kober and J. Peters. *Reinforcement Learning: State-of-the-Art (Adaptation, Learning, and Optimization)*, chapter Reinforcement Learning in Robotics: a Survey, pages 579–610. Springer Berlin / Heidelberg, 2012.
- J. Kober, B. Mohler, and J. Peters. *From Motor Learning to Interaction Learning in Robots*, volume 264 of *Studies in Computational Intelligence*, chapter Imitation and Reinforcement Learning for Motor Primitives with Perceptual Coupling, pages 209–225. Springer Berlin / Heidelberg, 2010.



Bibliography

- P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning (ICML)*, page 1, 2004.
- P. Abbeel, M. Quigley, and A. Y. Ng. Using inaccurate models in reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2006.
- P. Abbeel, A. Coates, M. Quigley, and A. Y. Ng. An application of reinforcement learning to aerobatic helicopter flight. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- C. H. An, C. G. Atkeson, and J. M. Hollerbach. *Model-based control of a robot manipulator*. MIT Press, Cambridge, MA, USA, 1988.
- R. Andersson. *A robot ping-pong player: experiment in real-time intelligent control*. MIT Press, 1988.
- C. Andrieu, N. de Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50(1):5–43, 2003.
- B. D. Argall, B. Browning, and M. Veloso. Learning robot motion control with demonstration and advice-operators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008.
- B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57:469–483, 2009.
- S. Arimoto, S. Kawamura, and F. Miyazaki. Bettering operation of robots by learning. *Journal of Robotic Systems*, 1(2):123–140, 1984.
- M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine Learning*, 23(2-3):279–303, 1996.
- C. G. Atkeson. Using local trajectory optimizers to speed up global optimization in dynamic programming. In J. D. Cowan, G. Tesauro, and J. Alspector, editors, *Advances in Neural Information Processing Systems 6 (NIPS)*, pages 503–521, Denver, CO, USA, 1994.
- C. G. Atkeson. Nonparametric model-based reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, 1998.
- C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1997.
- C. G. Atkeson, A. Moore, and S. Stefan. Locally weighted learning for control. *AI Review*, 11:75–113, April 1997.
- H. Attias. Planning by probabilistic inference. In *Proceedings of the 9th International Workshop on Artificial Intelligence and Statistics (AISTATS)*, Key West, FL, USA, 2003.
- J. A. Bagnell. *Learning Decisions: Robustness, Uncertainty, and Approximation*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 2004.
- J. A. Bagnell and J. C. Schneider. Autonomous helicopter control using reinforcement learning policy search methods. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2001.

-
- J. A. Bagnell and J. C. Schneider. Covariant policy search. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1019–1024, 2003.
- J. A. Bagnell, S. Kadade, A. Ng, and J. C. Schneider. Policy search by dynamic programming. In *Advances in Neural Information Processing Systems 16 (NIPS)*, 2004.
- B. Bakker, V. Zhumatiy, G. Gruener, and J. Schmidhuber. A robot that reinforcement-learns to identify and memorize important previous observations. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- B. Bakker, V. Zhumatiy, G. Gruener, and J. Schmidhuber. Quasi-online reinforcement learning for robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2006.
- A. G. Barto and S. Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):341–379, 2003.
- P. M. Bays and D. M. Wolpert. Computational principles of sensorimotor control that minimise uncertainty and variability. *Journal of Physiology*, 578:387–396, 2007.
- R. E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- R. E. Bellman. *Introduction to the Mathematical Theory of Control Processes: Linear Equations and Quadratic Criteria*, volume I. Academic Press, New York, NY, 1967.
- R. E. Bellman. *Introduction to the Mathematical Theory of Control Processes: Nonlinear Processes*, volume II. Academic Press, New York, NY, 1971.
- H. Benbrahim and J. A. Franklin. Biped dynamic walking using reinforcement learning. *Robotics and Autonomous Systems*, 22(3-4):283–302, 1997.
- H. Benbrahim, J. S. Doleac, J. A. Franklin, and O. Selfridge. Real-time learning: a ball on a beam. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 1992.
- D. C. Bentivegna, C. G. Atkeson, and G. Cheng. Learning from observation and practice using behavioral primitives: Marble maze, 2004a.
- D. C. Bentivegna, A. Ude, C. G. Atkeson, and G. Cheng. Learning to act from observation and practice. *International Journal of Humanoid Robotics*, 1(4):585–611, 2004b.
- J. v. d. Berg, S. Miller, D. Duckworth, H. Hu, A. Wan, X.-Y. Fu, K. Goldberg, and P. Abbeel. Superhuman performance of surgical tasks by robots using iterative learning from human-guided demonstrations. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.
- J. T. Betts. *Practical methods for optimal control using nonlinear programming*, volume 3 of *Advances in Design and Control*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.
- J. Binder, D. Koller, S. J. Russell, and K. Kanazawa. Adaptive probabilistic networks with hidden variables. *Machine Learning*, 29(2-3):213–244, 1997.
- N. Birdwell and S. Livingston. Reinforcement learning in sensor-guided aibo robots. Technical report, University of Tennessee, Knoxville, June-August 2007. advised by Dr. Itamar Elhanany.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer Verlag, 2006.
- S. Bitzer, M. Howard, and S. Vijayakumar. Using dimensionality reduction to exploit constraints in reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.

-
- R. Bootsma and P. van Wieringen. Timing an attacking forehand drive in table tennis. *Journal of Experimental Psychology: Human Perception and Performance*, 16:21–29, 1990.
- BotJunkie. Botjunkie interview: Nancy dussault smith on irobot’s roomba, Feb. 2012. URL <http://www.botjunkie.com/2010/05/17/botjunkie-interview-nancy-dussault-smith-on-irobots-roomba/>.
- A. Boularias, J. Kober, and J. Peters. Relative entropy inverse reinforcement learning. *Journal of Machine Learning Research - Proceedings Track*, 15:182–189, 2011.
- J. Buchli, F. Stulp, E. Theodorou, and S. Schaal. Learning variable impedance control. *International Journal of Robotic Research*, 30(7):820–833, 2011.
- M. Buehler, D. E. Koditschek, and P. J. Kindlmann. Planning and control of robotic juggling and catching tasks. *International Journal of Robotic Research*, 13(2):101–118, 1994.
- B. H. M. Bukkems, D. Kostic, A. G. de Jager, and M. Steinbuch. Learning-based identification and iterative learning control of direct-drive robots. *IEEE Transactions on Control Systems Technology*, 13(4):537–549, 2005.
- L. Buşoniu, R. Babuška, B. De Schutter, and D. Ernst. *Reinforcement Learning and Dynamic Programming Using Function Approximators*. CRC Press, Boca Raton, Florida, 2010.
- R. Caruana. Multitask learning. *Machine Learning*, 28:41–75, 1997.
- H. W.-k. Chan, I. King, and J. Lui. Performance analysis of a new updating rule for TD(λ) learning in feedforward networks for position evaluation in go game. In *Proceedings of the IEEE International Conference on Neural Networks*, volume 3, pages 1716–1720, 1996.
- F. Chaumette and E. Marchand. A redundancy-based iterative approach for avoiding joint limits: application to visual servoing. *IEEE Transactions on Robotics and Automation*, 17(5):719–730, 2001.
- G. Cheng, S. Hyon, J. Morimoto, A. Ude, J. G. Hale, G. Colvin, W. Scroggin, and S. C. Jacobsen. CB: A humanoid research platform for exploring neuroscience. *Journal of Advance Robotics*, 21(10):1097–1114, 2007.
- S. Chiappa, J. Kober, and J. Peters. Using bayesian dynamical systems for motion template libraries. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 297–304, 2009.
- A. Coates, P. Abbeel, and A. Y. Ng. Apprenticeship learning for helicopter control. *Communications of the ACM*, 52(7):97–105, 2009.
- A. Cocora, K. Kersting, C. Plagemann, W. Burgard, and L. D. Raedt. Learning relational navigation policies. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006.
- K. Conn and R. A. Peters II. Reinforcement learning with a supervisor for a mobile robot in a real-world environment. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2007.
- K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer. Online passive-aggressive algorithms. *Journal of Machine Learning Research*, 7:551–585, 2006.
- C. Daniel, G. Neumann, and J. Peters. Hierarchical relative entropy policy search. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- DARPA. Autonomous robot manipulation (ARM), Oct. 2010a. URL <http://www.darpa.mil/ipto/programs/arm/arm.asp>.

-
- DARPA. Learning applied to ground robotics (LAGR), Oct. 2010b. URL <http://www.darpa.mil/ipto/programs/lagr/lagr.asp>.
- DARPA. Learning locomotion (L2), Oct. 2010c. URL <http://www.darpa.mil/ipto/programs/l1/l1.asp>.
- P. Dayan and G. E. Hinton. Using expectation-maximization for reinforcement learning. *Neural Computation*, 9(2):271–278, 1997.
- M. P. Deisenroth and C. E. Rasmussen. A practical and conceptual framework for learning in control. Technical Report UW-CSE-10-06-01, Department of Computer Science & Engineering, University of Washington, 185 Stevens Way, Seattle, WA, 98195, USA, June 2010.
- M. P. Deisenroth, C. E. Rasmussen, and D. Fox. Learning to control a low-cost manipulator using data-efficient reinforcement learning. In *Proceedings of the 9th International Conference on Robotics: Science & Systems (R:SS)*, 2011.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39:1–38, 1977.
- J.-Y. Donnat and J.-A. Meyer. Learning reactive and planning rules in a motivationally autonomous animat. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 26(3):381–395, 1996.
- M. Dorigo and M. Colombetti. Robot shaping: Developing situated agents through learning. Technical report, International Computer Science Institute, Berkeley, CA, 1993.
- K. Doya. Metalearning and neuromodulation. *Neural Networks*, 15(4-6):495 – 506, 2002.
- Y. Duan, Q. Liu, and X. Xu. Application of reinforcement learning in robot soccer. *Engineering Applications of Artificial Intelligence*, 20(7):936–950, 2007.
- Y. Duan, B. Cui, and H. Yang. Robot navigation based on fuzzy RL algorithm. In *Proceedings of the International Symposium on Neural Networks (ISNN)*, 2008.
- A. El-Fakdi, M. Carreras, and P. Ridao. Towards direct policy search reinforcement learning for robot control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006.
- G. Endo, J. Morimoto, T. Matsubara, J. Nakanishi, and G. Cheng. Learning CPG-based biped locomotion with a policy gradient method: Application to a humanoid robot. *International Journal of Robotic Research*, 27(2):213–228, 2008.
- Y. Engel, S. Mannor, and R. Meir. Reinforcement learning with gaussian processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 201–208, 2005.
- M. S. Erden and K. Leblebicioğlu. Free gait generation with reinforcement learning for a six-legged robot. *Robotics and Autonomous Systems (RAS)*, 56(3):199–212, 2008.
- A. H. Fagg, D. L. Lotspeich, Hoff, and G. A. J. Bekey. Rapid reinforcement learning for reactive control policy design for autonomous robots. In *Proceedings of Artificial Life in Robotics*, 1998.
- I. Fantoni and R. Lozano. *Non-Linear Control for Underactuated Mechanical Systems*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2001.
- B. G. Farley and W. A. Clark. Simulation of self-organizing systems by digital computer. *IRE Transactions on Information Theory*, 4:76–84, 1954.

- H. Fässler, H. A. Beyer, and J. T. Wen. A robot ping pong player: optimized mechanics, high performance 3d vision, and intelligent sensor control. *Robotersysteme*, 6(3):161–170, 1990.
- W. Findeisen, F. N. Bailey, M. Brdeys, K. Malinowski, P. Tatjewski, and J. Wozniak. *Control and coordination in hierarchical systems*. International series on applied systems analysis. J. Wiley, Chichester, 1980.
- C. Freeman, P. Lewin, E. Rogers, and J. Ratcliffe. Iterative learning control applied to a gantry robot and conveyor system. *Transactions of the Institute of Measurement and Control*, 32(3):251–264, June 2010.
- A. Gams and A. Ude. Generalization of example movements with dynamic systems. In *Proceedings of the IEEE/RSJ International Conference on Humanoid Robots (HUMANOIDS)*, 2009.
- C. Gaskett, L. Fletcher, and A. Zelinsky. Reinforcement learning for a vision based mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2000.
- T. Geng, B. Porr, and F. Wörgötter. Fast biped walking with a reflexive controller and real-time policy searching. In *Advances in Neural Information Processing Systems (NIPS)*, 2006.
- P. Glynn. Likelihood ratio gradient estimation: an overview. In *Proceedings of the Winter Simulation Conference (WSC)*, 1987.
- D. E. Goldberg. *Genetic algorithms*. Addison Wesley, 1989.
- K. Gräve, J. Stücker, and S. Behnke. Learning motion skills from expert demonstrations and own experience using gaussian process regression. In *Proceedings of the Joint International Symposium on Robotics (ISR) and German Conference on Robotics (ROBOTIK)*, 2010.
- D. B. Grimes and R. P. N. Rao. Learning nonparametric policies by imitation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and System (IROS)*, pages 2022–2028, 2008.
- F. Guenter, M. Hersch, S. Calinon, and A. Billard. Reinforcement learning for imitating constrained reaching movements. *Advanced Robotics, Special Issue on Imitative Robots*, 21(13):1521–1544, 2007.
- V. Gullapalli, J. Franklin, and H. Benbrahim. Acquiring robot skills via reinforcement learning. *IEEE Control Systems Journal, Special Issue on Robotics: Capturing Natural Motion*, 4(1):13–24, February 1994.
- R. Hafner and M. Riedmiller. Reinforcement learning on a omnidirectional mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- R. Hafner and M. Riedmiller. Neural reinforcement learning controllers for a real robot application. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2007.
- G. Hailu and G. Sommer. Integrating symbolic knowledge in reinforcement learning. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*, 1998.
- T. Hester, M. Quinlan, and P. Stone. Generalized model learning for reinforcement learning on a humanoid robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- M. Hoffman, A. Doucet, N. de Freitas, and A. Jasra. Bayesian policy learning with trans-dimensional MCMC. In *Advances in Neural Information Processing Systems 20 (NIPS)*, 2007.
- X. Huang and J. Weng. Novelty and reinforcement learning in the value system of developmental robots. In *Proceedings of the 2nd International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 2002.
- A. Hubbard and C. Seng. Visual movements of batters. *Research Quarterly*, 25:42–57, 1954.

- M. Huber and R. Grupen. Learning robot control – using control policies as abstract actions. In *Proceedings of the NIPS'98 Workshop: Abstraction and Hierarchy in Reinforcement Learning*, 1998.
- A. J. Ijspeert, J. Nakanishi, and S. Schaal. Movement imitation with nonlinear dynamical systems in humanoid robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1398–1403, Washington, DC, May 11-15 2002a.
- A. J. Ijspeert, J. Nakanishi, and S. Schaal. Learning attractor landscapes for learning motor primitives. In *Advances in Neural Information Processing Systems 15 (NIPS)*, pages 1523–1530, 2002b.
- W. Ilg, J. Albiez, H. Jeede, K. Berns, and R. Dillmann. Adaptive periodic movement control for the four legged walking machine BISAM. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- T. Jaakkola, M. I. Jordan, and S. P. Singh. Convergence of stochastic iterative dynamic programming algorithms. In *Advances in Neural Information Processing Systems (NIPS)*, volume 6, pages 703–710, 1993.
- N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In *Proceedings of the 3rd European Conference on Artificial Life*, pages 704–720, 1995.
- N. Jetchev and M. Toussaint. Trajectory prediction: learning to map situations to robot trajectories. In *Proceedings of the International Conference on Machine Learning (ICML)*, page 57, 2009.
- L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- M. Kalakrishnan, L. Righetti, P. Pastor, and S. Schaal. Learning force control policies for compliant manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4639–4644, 2011.
- R. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.
- R. E. Kalman. When is a linear control system optimal? *Journal of Basic Engineering*, 86(1):51–60, 1964.
- Z. Kalmár, C. Szepesvári, and A. Lőrincz. Modular reinforcement learning: An application to a real robot task. *Lecture Notes in Artificial Intelligence*, 1545:29–45, 1998.
- H. J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 11:P11011, 2005.
- 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 Conference (R:SS)*, 2008.
- M. Kawato. Feedback-error-learning neural network for supervised motor learning. *Advanced neural computers*, 6(3):365–372, 1990.
- R. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. J. Wiley, New York, 1976.
- O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5(1):90–98, 1986.
- H. Kimura, T. Yamashita, and S. Kobayashi. Reinforcement learning of walking behavior for a four-legged robot. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2001.

-
- F. Kirchner. Q-learning of complex behaviours on a six-legged walking machine. In *Proceedings of the EUROMICRO Workshop on Advanced Mobile Robots*, 1997.
- D. E. Kirk. *Optimal control theory*. Prentice-Hall, Englewood Cliffs, New Jersey, 1970.
- J. Ko, D. J. Klein, D. Fox, and D. Hähnel. Gaussian processes and reinforcement learning for identification and control of an autonomous blimp. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2007.
- J. Kober and J. Peters. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems 21 (NIPS)*, pages 849–856, 2008.
- J. Kober and J. Peters. Learning motor primitives for robotics. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2112–2118, 2009.
- J. Kober and J. Peters. Imitation and reinforcement learning - practical algorithms for motor primitive learning in robotics. *IEEE Robotics and Automation Magazine*, 17(2):55–62, 2010.
- J. Kober and J. Peters. Policy search for motor primitives in robotics. *Machine Learning*, 84(1-2):171–203, 2011a.
- J. Kober and J. Peters. Learning elementary movements jointly with a higher level task. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 338–343, 2011b.
- J. Kober and J. Peters. Learning prioritized control of motor primitives. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Under Review, 2012.
- J. Kober, B. Mohler, and J. Peters. Learning perceptual coupling for motor primitives. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 834–839, 2008.
- J. Kober, K. Mülling, O. Krömer, C. H. Lampert, B. Schölkopf, and J. Peters. Movement templates for learning of hitting and batting. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 853–858, 2010a.
- J. Kober, E. Oztop, and J. Peters. Reinforcement learning to adjust robot movements to new situations. In *Proceedings of Robotics: Science and Systems Conference (R:SS)*, 2010b.
- J. Kober, E. Oztop, and J. Peters. Reinforcement learning to adjust robot movements to new situations. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), Best Paper Track*, pages 2650–2655, 2011.
- J. Kober, J. A. Bagnell, and J. Peters. Reinforcement learning in robotics: a survey. *International Journal of Robotics Research*, In Preparation, 2012a.
- J. Kober, A. Wilhelm, E. Oztop, and J. Peters. Reinforcement learning to adjust parametrized motor primitives to new situations. *Autonomous Robots*, Under Review, 2012b.
- N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.
- J. Z. Kolter and A. Y. Ng. Policy search via the signed derivative. In *Proceedings of Robotics: Science and Systems Conference (R:SS)*, 2009.
- J. Z. Kolter, P. Abbeel, and A. Y. Ng. Hierarchical apprenticeship learning with application to quadruped locomotion. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.

- J. Z. Kolter, A. Coates, A. Y. Ng, Y. Gu, and C. DuHadway. Space-indexed dynamic programming: learning to follow trajectories. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2008.
- J. Z. Kolter, C. Plagemann, D. T. Jackson, A. Y. Ng, and S. Thrun. A probabilistic approach to mixed open-loop and closed-loop control, with application to extreme autonomous driving. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- P. Kormushev, S. Calinon, and D. G. Caldwell. Robot motor skill coordination with em-based reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- L. Kovar. *Automated methods for data-driven synthesis of realistic and controllable human motion*. PhD thesis, University of Wisconsin at Madison, 2004.
- O. Kroemer, R. Detry, J. Piater, and J. Peters. Active learning using mean shift optimization for robot grasping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009.
- O. Kroemer, R. Detry, J. Piater, and J. Peters. Combining active learning and reactive control for robot grasping. *Robotics and Autonomous Systems*, 58(9):1105–1116, 09 2010.
- K. Kronander, M. S. M. Khansari-Zadeh, and A. Billard. Learning to control planar hitting motions in a minigolf-like task. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 710–717, 2011.
- H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability*, 1950.
- P. Kulchenko. Robot juggles two ping-pong balls, Jan. 2011. URL <http://notebook.kulchenko.com/juggling/robot-juggles-two-ping-pong-balls>.
- I. Kwee, M. Hutter, and J. Schmidhuber. Gradient-based reinforcement planning in policy-search methods. In *Proceedings of the 5th European Workshop on Reinforcement Learning (EWRL)*, number 27, pages 27–29, 2001.
- R. Lampariello, D. Nguyen-Tuong, C. Castellini, G. Hirzinger, and J. Peters. Trajectory planning for optimal robot catching in real-time. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3719–3726, 2011.
- T. Latzke, S. Behnke, and M. Bennewitz. Imitative reinforcement learning for soccer playing robots. In *RoboCup 2006: Robot Soccer World Cup X*, pages 47–58. Springer-Verlag, Berlin, Heidelberg, 2007.
- G. Lawrence, N. Cowan, and S. Russell. Efficient gradient estimation for motor control learning. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 354–361, 2003.
- T. Lens, J. Kunz, C. Trommer, A. Karguth, and O. von Stryk. Biorob-arm: A quickly deployable and intrinsically safe, light-weight robot arm for service robotics applications. In *Proceedings of the 41st International Symposium on Robotics / 6th German Conference on Robotics*, pages 905–910, 2010.
- D. Lizotte, T. Wang, M. Bowling, and D. Schuurmans. Automatic gait optimization with gaussian process regression. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- S. Mahadevan and J. Connell. Automatic programming of behavior-based robots using reinforcement learning. *Artificial Intelligence*, 55(2-3):311 – 365, 1992.

- J. A. Martín H., J. de Lope, and D. Maravall. The kNN-TD reinforcement learning algorithm. In *Proceedings of the 3rd International Work-Conference on The Interplay Between Natural and Artificial Computation (IWINAC)*, pages 305–314, 2009.
- T. Martínez-Marín and T. Duckett. Fast reinforcement learning for vision-guided mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- Masters Games Ltd. The rules of darts, July 2010. URL <http://www.mastersgames.com/rules/darts-rules.htm>.
- M. J. Mataric. Reward functions for accelerated learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1994.
- M. J. Mataric. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4:73–83, 1997.
- M. Matsushima, T. Hashimoto, M. Takeuchi, and F. Miyazaki. A learning approach to robotic table tennis. *IEEE Transactions on Robotics*, 21(4):767–771, 2005.
- A. McGovern and A. G. Barto. Automatic discovery of subgoals in reinforcement learning using diverse density. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 361–368, 2001.
- A. McGovern, R. S. Sutton, and A. H. Fagg. Roles of macro-actions in accelerating reinforcement learning. In *Proceedings of Grace Hopper Celebration of Women in Computing*, 1997.
- G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics. John Wiley & Sons, 1997.
- P. McLeod, N. Reed, and Z. Dienes. The generalized optic acceleration cancellation theory of catching. *Journal of Experimental Psychology. Human Perception and Performance*, 32(1):139–148, 2006.
- J. Michels, A. Saxena, and A. Y. Ng. High speed obstacle avoidance using monocular vision and reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2005.
- M. L. Minsky. *Theory of Neural-Analog Reinforcement Systems and its Application to the Brain-Model Problem*. PhD thesis, Princeton University, 1954.
- N. Mitsunaga, C. Smith, T. Kanda, H. Ishiguro, and N. Hagita. Robot behavior adaptation for human-robot interaction based on policy gradient reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2005.
- H. Miyamoto, S. Schaal, F. Gandolfo, H. Gomi, Y. Koike, R. Osu, E. Nakano, Y. Wada, and M. Kawato. A kendama learning robot based on bi-directional theory. *Neural Networks*, 9(8):1281–1302, 1996. ISSN 0893-6080.
- J. Morimoto and K. Doya. Acquisition of stand-up behavior by a real robot using hierarchical reinforcement learning. *Robotics and Autonomous Systems*, 36(1):37–51, 2001.
- M. Müller, S. Lupashin, and R. D’Andrea. Quadrocopter ball juggling. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 5113–5120, 2011.
- K. Mülling and J. Peters. A computational model of human table tennis for robot application. In *Proceedings of Autonome Mobile Systeme (AMS)*, 2009.
- K. Mülling, J. Kober, and J. Peters. Learning table tennis with a mixture of motor primitives. In *Proceedings of the 10th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 411–416, 2010.

- K. Mülling, J. Kober, and J. Peters. A biomimetic approach to robot table tennis. *Adaptive Behavior*, 9(5): 359–376, 2011.
- J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, and M. Kawato. Learning from demonstration and adaptation of biped locomotion. *Robotics and Autonomous Systems (RAS)*, 47(2-3):79–91, 2004.
- J. Nakanishi, R. Cory, M. Mistry, J. Peters, and S. Schaal. Operational space control: a theoretical and empirical comparison. *International Journal of Robotics Research*, 27:737–757, 2008.
- B. Nemec, M. Tamošiūnaitė, F. Wörgötter, and A. Ude. Task adaptation through exploration and action sequencing. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2009.
- B. Nemec, M. Zorko, and L. Zlajpah. Learning of a ball-in-a-cup playing robot. In *Proceedings of the International Workshop on Robotics in Alpe-Adria-Danube Region (RAAD)*, 2010.
- A. Y. Ng and M. Jordan. Pegasus: A policy search method for large mdps and pomdps. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 406–415, 2000.
- A. Y. Ng, A. Coates, M. Diel, V. Ganapathi, J. Schulte, B. Tse, E. Berger, and E. Liang. Autonomous inverted helicopter flight via reinforcement learning. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2004a.
- A. Y. Ng, H. J. Kim, M. I. Jordan, and S. Sastry. Inverted autonomous helicopter flight via reinforcement learning. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*. MIT Press, 2004b.
- D. Nguyen-Tuong and J. Peters. Incremental online sparsification for model learning in real-time robot control. *Neurocomputing*, 74(11):1859–1867, 2011.
- M. Norrlöf. An adaptive iterative learning control algorithm with experiments on an industrial robot. *IEEE Transactions on Robotics*, 18(2):245–251, 2002.
- S. Oßwald, A. Hornung, and M. Bennewitz. Learning reliable and efficient navigation with a humanoid. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- L. Paletta, G. Fritz, F. Kintzler, J. Irran, and G. Dorffner. Perception and developmental learning of affordances in autonomous robots. In J. Hertzberg, M. Beetz, and R. Englert, editors, *KI 2007: Advances in Artificial Intelligence*, volume 4667 of *Lecture Notes in Computer Science*, pages 235–250. Springer Berlin / Heidelberg, 2007.
- D.-H. Park, H. Hoffmann, P. Pastor, and S. Schaal. Movement reproduction and obstacle avoidance with dynamic movement primitives and potential fields. In *Proceedings of the IEEE International Conference on Humanoid Robots (HUMANOIDS)*, pages 91–98, 2008.
- PASCAL2. Challenges, Oct. 2010. URL <http://pascallin2.ecs.soton.ac.uk/Challenges/>.
- P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal. Learning and generalization of motor skills by learning from demonstration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1293–1298, 2009.
- P. Pastor, M. Kalakrishnan, S. Chitta, E. Theodorou, and S. Schaal. Skill learning and task outcome prediction for manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011.
- M. Pendrith. Reinforcement learning in situated agents: Some theoretical problems and practical solutions. In *Proceedings of the European Workshop on Learning Robots (EWRL)*, 1999.

-
- L. Peshkin. *Reinforcement Learning by Policy Search*. PhD thesis, Brown University, Providence, RI, 2001.
- J. Peters. *Machine Learning of Motor Skills for Robotics*. PhD thesis, University of Southern California, Los Angeles, CA, 90089, USA, April 2007.
- J. Peters and S. Schaal. Policy gradient methods for robotics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2219 – 2225, 2006.
- J. Peters and S. Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2007.
- J. Peters and S. Schaal. Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21(4):682–697, 05 2008a.
- J. Peters and S. Schaal. Learning to control in operational space. *International Journal of Robotics Research*, 27(2):197–212, 02 2008b.
- J. Peters and S. Schaal. Natural actor-critic. *Neurocomputing*, 71(7-9):1180–1190, 03 2008c.
- J. Peters, S. Vijayakumar, and S. Schaal. Reinforcement learning for humanoid robotics. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS)*, pages 103–123, 2003.
- J. Peters, S. Vijayakumar, and S. Schaal. Linear quadratic regulation as benchmark for policy gradient methods. Technical report, University of Southern California, 2004.
- J. Peters, S. Vijayakumar, and S. Schaal. Natural actor-critic. In *Proceedings of the European Conference on Machine Learning (ECML)*, pages 280–291, 2005.
- J. Peters, K. Mülling, and Y. Altun. Relative entropy policy search. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2010a.
- J. Peters, K. Mülling, J. Kober, D. Nguyen-Tuong, and O. Kroemer. Towards motor skill learning for robotics. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2010b.
- J. H. Piater, S. Jodogne, R. Detry, D. Kraft, N. Krüger, O. Kroemer, and J. Peters. Learning visual representations for perception-action systems. *International Journal of Robotic Research*, 30(3):294–307, 2011.
- R. Platt, R. A. Grupen, and A. H. Fagg. Improving grasp skills using schema structured learning. In *Proceedings of the International Conference on Development and Learning*, 2006.
- D. Pongas, A. Billard, and S. Schaal. Rapid synchronization and accurate phase-locking of rhythmic motor primitives. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2911–2916, 2005.
- D. Precup, R. S. Sutton, and S. Singh. Theoretical results on reinforcement learning with temporally abstract options. In *Proceedings of the European Conference on Machine Learning (ECML)*, 1998.
- M. Ramanantsoa and A. Durey. Towards a stroke construction model. *International Journal of Table Tennis Science*, 2:97–114, 1994.
- C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- K. J. Åström and B. Wittenmark. *Adaptive Control*. Addison-Wesley, Reading, MA, 1989.
- N. Ratliff, D. Bradley, J. A. Bagnell, and J. Chestnutt. Boosting structured prediction for imitation learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2006a.

- N. Ratliff, J. A. Bagnell, and S. Srinivasa. Imitation learning for locomotion and manipulation. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS)*, 2007.
- N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, 2006b.
- P. Reist and R. D’Andrea. Bouncing an unconstrained ball in three dimensions with a blind juggling robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2753–2760, 2009.
- M. Riedmiller, T. Gabel, R. Hafner, and S. Lange. Reinforcement learning for robot soccer. *Autonomous Robots*, 27(1):55–73, 2009.
- A. Rottmann, C. Plagemann, P. Hilgers, and W. Burgard. Autonomous blimp control using model-free reinforcement learning in a continuous state and action space. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007.
- R. Y. Rubinstein and D. P. Kroese. *The Cross Entropy Method: A Unified Approach To Combinatorial Optimization, Monte-carlo Simulation (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2004.
- T. Rückstieß, M. Felder, and J. Schmidhuber. State-dependent exploration for policy gradient methods. In *Proceedings of the European Conference on Machine Learning (ECML)*, pages 234–249, 2008.
- S. Russell. Learning agents for uncertain environments (extended abstract). In *Proceedings of the 11th Annual Conference on Computational Learning Theory (COLT)*, pages 101–103, 1998.
- M.-a. Sato, Y. Nakamura, and S. Ishii. Reinforcement learning for biped locomotion. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN)*, 2002.
- S. Sato, T. Sakaguchi, Y. Masutani, and F. Miyazaki. Mastering of a task with interaction between a robot and its environment : “kendama” task. *Transactions of the Japan Society of Mechanical Engineers. C*, 59 (558):487–493, 1993.
- S. Schaal. Learning from demonstration. In *Advances in Neural Information Processing Systems (NIPS)*, 1996.
- S. Schaal and C. G. Atkeson. Robot juggling: An implementation of memory-based learning. *Control Systems Magazine*, 14(1):57–71, 1994.
- S. Schaal, D. Sternad, and C. G. Atkeson. One-handed juggling: a dynamical approach to a rhythmic movement task. *Journal of Motor Behavior*, 28(2):165–183, 1996.
- S. Schaal, C. G. Atkeson, and S. Vijayakumar. Scalable techniques from nonparametric statistics for real-time robot learning. *Applied Intelligence*, 17(1):49–60, 2002.
- S. Schaal, J. Peters, J. Nakanishi, and A. J. Ijspeert. Control, planning, learning, and imitation with dynamic movement primitives. In *Proceedings of the Workshop on Bilateral Paradigms on Humans and Humanoids, IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- S. Schaal, P. Mohajerian, and A. J. Ijspeert. Dynamics systems vs. optimal control - a unifying view. *Progress in Brain Research*, 165(1):425–445, 2007.
- R. Schmidt and C. Wrisberg. *Motor Learning and Performance*. Human Kinetics, second edition, 2000.
- J. G. Schneider. Exploiting model uncertainty estimates for safe dynamic control learning. In *Neural Information Processing Systems 9 (NIPS)*, pages 1047–1053, 1996.

- F. Sehnke, C. Osendorfer, T. Rückstieß, A. Graves, J. Peters, and J. Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 21(4):551–559, 05 2010.
- T. Senoo, A. Namiki, and M. Ishikawa. Ball control in high-speed batting motion using hybrid trajectory generator. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1762–1767, 2006.
- L. Sentis and O. Khatib. Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robotics*, 2(4):505–518, 2005.
- T. Shone, G. Krudysz, and K. Brown. Dynamic manipulation of kendama. Technical report, Rensselaer Polytechnic Institute, 2000.
- D. Silver, J. A. Bagnell, and A. Stentz. High performance outdoor navigation from overhead data using imitation learning. In *Proceedings of Robotics: Science and Systems Conference (R:SS)*, 2008.
- H. A. Simon. *The Shape of Automation for Men and Management*. New York: Harper & Row, 1965.
- W. D. Smart and L. P. Kaelbling. A framework for reinforcement learning on real robots. In *Proceedings of the National Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence (AAAI/IAAI)*, 1998.
- W. D. Smart and L. P. Kaelbling. Effective reinforcement learning for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2002.
- V. Soni and S. Singh. Reinforcement learning of hierarchical skills on the sony aibo robot. In *Proceedings of the International Conference on Development and Learning (ICDL)*, 2006.
- M. Strens and A. Moore. Direct policy search using paired statistical tests. In *Proceedings of the 18th International Conference on Machine Learning (ICML)*, 2001.
- F. Stulp, E. Theodorou, M. Kalakrishnan, P. Pastor, L. Righetti, and S. Schaal. Learning motion primitive goals for robust manipulation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 325 –331, 2011.
- C. Sumners. *Toys in Space: Exploring Science with the Astronauts*. McGraw-Hill, 1997.
- R. S. Sutton. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the International Machine Learning Conference (ICML)*, pages 9–44, 1990.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, Boston, MA, 1998.
- R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12 (NIPS)*, pages 1057–1063, 1999.
- R. S. Sutton, A. Koop, and D. Silver. On the role of tracking in stationary environments. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2007.
- M. M. Svinin, K. Yamada, and K. Ueda. Emergent synthesis of motion patterns for locomotion robots. *Artificial Intelligence in Engineering*, 15(4):353 – 363, 2001.
- K. Takenaka. Dynamical control of manipulator with vision : “cup and ball” game demonstrated by robot. *Transactions of the Japan Society of Mechanical Engineers. C*, 50(458):2046–2053, 1984.

- T. Tamei and T. Shibata. Policy gradient learning of cooperative interaction with a robot using user's biological signals. In *Proceedings of the International Conference on Neural Information Processing (ICONIP)*, 2009.
- M. Tamosiunaite, B. Nemec, A. Ude, and F. Wörgötter. Learning to pour with a robot arm combining goal and shape learning for dynamic movement primitives. *Robotics and Autonomous Systems*, 59:910 – 922, 2011.
- M. E. Taylor, S. Whiteson, and P. Stone. Transfer via inter-task mappings in policy search reinforcement learning. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2007.
- R. Tedrake. Stochastic policy gradient reinforcement learning on a simple 3d biped. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2004.
- R. Tedrake, T. W. Zhang, and H. S. Seung. Stochastic policy gradient reinforcement learning on a simple 3d biped. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 2849–2854, 2004.
- R. Tedrake, T. W. Zhang, and H. S. Seung. Learning to walk in 20 minutes. In *Proceedings of the Yale Workshop on Adaptive and Learning Systems*, 2005.
- R. Tedrake, I. R. Manchester, M. M. Tobenkin, and J. W. Roberts. LQR-trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research*, 29:1038–1052, 2010.
- G. Tesauro. TD-gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994.
- E. A. Theodorou, J. Buchli, and S. Schaal. Reinforcement learning of motor skills in high dimensions: A path integral approach. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2397–2403, 2010.
- S. Thrun. An approach to learning mobile robot navigation. *Robotics and Autonomous Systems*, 15: 301–319, 1995.
- M. Tokic, W. Ertel, and J. Fessler. The crawler, a class room demonstrator for reinforcement learning. In *Proceedings of the International Florida Artificial Intelligence Research Society Conference (FLAIRS)*, 2009.
- M. Toussaint and C. Goerick. Probabilistic inference for structured planning in robotics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007.
- M. Toussaint, A. Storkey, and S. Harmeling. *Inference and Learning in Dynamic Models*, chapter Expectation-Maximization methods for solving (PO)MDPs and optimal control problems. Cambridge University Press, 2010.
- C. Touzet. Neural reinforcement learning for behaviour synthesis. *Robotics and Autonomous Systems, Special issue on Learning Robot: the New Wave*, 22(3-4):251–281, December 1997.
- C. Towell, M. Howard, and S. Vijayakumar. Learning nullspace policies. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- D. Tyldesley and H. Whiting. Operational timing. *Journal of Human Movement Studies*, 1:172–177, 1975.
- E. Uchibe, M. Asada, and K. Hosoda. Cooperative behavior acquisition in multi mobile robots environment by reinforcement learning based on state vector estimation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1998.

-
- A. Ude, A. Gams, T. Asfour, and J. Morimoto. Task-specific generalization of discrete and periodic dynamic movement primitives. *IEEE Transactions on Robotics*, 26(5):800–815, 2010.
- H. Urbanek, A. Albu-Schäffer, and P. van der Smagt. Learning from demonstration repetitive movements for autonomous service robotics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 4, pages 3495–3500, 2004.
- L. J. P. van der Maaten, E. O. Postma, and H. J. van den Herik. Dimensionality reduction: A comparative review. Technical Report TiCC-TR 2009-005, Tilburg University, 2009.
- N. Vlassis, M. Toussaint, G. Kontes, and S. Piperidis. Learning model-free robot control by a Monte Carlo EM algorithm. *Autonomous Robots*, 27(2):123–130, August 2009.
- B. Wang, J. W. Li, and H. Liu. A heuristic reinforcement learning for robot approaching objects. In *Proceedings of the IEEE Conference on Robotics, Automation and Mechatronics*, 2006.
- M. Welling. The Kalman filter. Lecture Notes, 2010.
- R. A. Willgoss and J. Iqbal. Reinforcement learning of behaviors in mobile robots using noisy infrared sensing. In *Proceedings of the Australian Conference on Robotics and Automation*, 1999.
- R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.
- G. Wulf. *Attention and motor skill learning*. Human Kinetics, Champaign, IL, 2007.
- J. Yamaguchi and A. Takanishi. Development of a biped walking robot having antagonistic driven joints using nonlinear spring mechanism. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 185–192, 1997.
- T. Yasuda and K. Ohkura. A reinforcement learning technique with an adaptive action generator for a multi-robot system. In *Proceedings of the International Conference on Simulation of Adaptive Behavior (SAB)*, 2008.
- T. Yoshikawa. Manipulability of robotic mechanisms. *The International Journal of Robotics Research*, 4(2): 3–9, 1985.
- S. M. Youssef. Neuro-based learning of mobile robots with evolutionary path planning. In *Proceedings of the ICGST International Conference on Automation, Robotics and Autonomous Systems (ARAS)*, 2005.
- B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1433–1438, 2008.



List of Figures

1.1	This figure illustrates the structure of this thesis. Chapter 2 reviews the field of reinforcement learning applied to robotics, Chapter 3 describes the employed motor skill representation. Chapter 4 provides a framework for policy search approaches. Chapter 5 describes how movements can be generalized to new situations. Chapter 6 outlines how several motor skills can be performed simultaneously to achieve a more complex task. Chapter 1 gives an introduction and Chapter 7 provides a summary and an outlook on future work.	4
2.1	This figure illustrates robots to which reinforcement learning has been applied. The robots cover the whole range of wheeled mobile robots, robotic arms, autonomous vehicles, to humanoid robots. (a) The OBELIX robot is a wheeled mobile robot that learned to push boxes [Mahadevan and Connell, 1992] with a value function-based approach (Picture reprint with permission of Sridhar Mahadevan). (b) The Zebra Zero robot is a robot arm that learned a peg-in-hole insertion task [Gullapalli et al., 1994] with a model-free policy gradient approach (Picture reprint with permission of Rod Grupen). (c) The control of such autonomous blimps (Picture reprint with permission of Axel Rottmann) was learned with both a value function based approach [Rottmann et al., 2007] and model-based policy search [Ko et al., 2007]. (d) The Sarcos humanoid DB learned a pole-balancing task [Schaal, 1996] using forward models (Picture reprint with permission of Stefan Schaal). .	9
2.2	This Figure illustrates the state space of a robot reinforcement learning task.	10
2.3	Boston Dynamics LittleDog jumping [Kolter and Ng, 2009] (Picture reprint with permission of Zico Kolter).	19
2.4	Autonomous inverted helicopter flight [Ng et al., 2004b] (Picture reprint with permission of Andrew Ng).	22
2.5	This figure shows schematic drawings of the ball-in-a-cup motion (a), the final learned robot motion (c), as well as a kinesthetic teach-in (b). The green arrows show the directions of the current movements in that frame. The human cup motion was taught to the robot by imitation learning with 31 parameters per joint for an approximately 3 seconds long movement. The robot manages to reproduce the imitated motion quite accurately, but the ball misses the cup by several centimeters. After approximately 75 iterations of the Policy learning by Weighting Exploration with the Returns (PoWER) algorithm the robot has improved its motion so that the ball regularly goes into the cup.	25
3.1	This figure illustrates the different phases of a table tennis stroke. The blue box on the left represents a ping-pong ball launcher, the table is shown in green and the different states of the robot are superposed. A typical racket trajectory is indicated by the dark gray solid line while the orange dashed line represents the ball's trajectory. The robot goes through four stages: it swings back (①→②), it strikes the ball at a virtual hitting point with a goal-oriented velocity at posture ③, it follows the strike through (③→④) and finally returns to the rest posture ①. While the complete arm movement can be modeled with the Ijspeert approach, the reformulation in this paper is required to be able to properly strike the ball.	32

3.2	In this figure, we convey the intuition of the presented reactive templates for learning striking movements. The Ijspeert formulation can be seen as a nonlinear spring damper system that pulls a degree of freedom to a stationary goal while exhibiting a learned movement shape. The presented approach allows hitting a target with a specified velocity without replanning if the target is adapted and, in contrast to the Ijspeert formulation, can be seen as a degree of freedom pulled towards a moving goal.	33
3.3	Target velocity adaptation is essential for striking movements. This figure illustrates how different versions of the dynamical system based motor primitives are affected by a change of the target velocity. Here, an artificial training example (i.e., $q = 2t^2 + \cos(4t\pi) - 1$) is generated. After learning, all motor primitive formulations manage to reproduce the movements accurately from the training example for the same target velocity and cannot be distinguished. When the target velocity is tripled, this picture changes drastically. For Ijspeert's original model the amplitude modifier \mathbf{a} had to be increased to yield the desired velocity. The increased amplitude of the trajectory is clearly visible for the positions and even more drastic for the velocities and accelerations. The reformulations presented in this paper, stay closer to the movement shape and amplitude. Particularly the velocities and accelerations exhibit that the new approach allows much better generalizing of the learned behavior. This figure furthermore demonstrates how a large initial step in acceleration appears for Ijspeert's original model (and the reformulation for hitting) even if a transformation function is used to partially suppress it for the training example.	35
3.4	An important aspect of the Ijspeert framework is that such primitives are guaranteed to be stable and, hence, safe for learning. A problem of the regular formulation highly unevenly distributed acceleration with a jump at the beginning of the movement of its unaltered dynamics. These high accelerations affect the movement when the behavior is either generalized to new goals or when during trial-and-error learning where the initial parameters are small. Some of these problem have previously been noted by Park et al. [2008], and are particularly bad in the context of fast striking movements. Here, we compare the different formulations with respect to their acceleration in the unaltered dynamics case (i.e., $\mathbf{w} = \mathbf{0}$). For a better comparison, we set the goal velocity to zero ($\dot{\mathbf{g}} = 0$). The Ijspeert formulation clearly shows the problem with the large acceleration, as does the reformulation for hitting (with a hitting speed of $\dot{\mathbf{g}} = 0$ both are identical). While the Park modification starts without the jump in acceleration, it requires almost as high accelerations shortly afterwards. The acceleration-safe reformulation for hitting also starts out without a step in acceleration and does not require huge accelerations.	36
3.5	This figure demonstrates the generalization of an imitated behavior to a different target that is 15cm away from the original target. Note that this trajectory is for a static target, hence the slow motion. The depicted degree of freedom (DoF) is shoulder adduction-abduction (i.e., the second DoF). The solid gray bars indicate the time before and after the main movement, the gray dashed lines indicate the phase borders also depicted in Figure 3.1 and the target is hit at the second border.	38
3.6	This figure presents a hitting sequence from the demonstration, a generalization on the robot with a ball attached by a string as well as a generalization hitting a ball shot by a ping-pong ball launcher. The demonstration and the flying ball generalization are captured by a 25Hz video camera, the generalization with the attached ball is captured with 200Hz through our vision system. From left to right the stills represent: rest posture, swing-back posture, hitting point, swing-through and rest posture. The postures (①-④) are the same as in Figure 3.2.	39
3.7	Generalization to various targets (five different forehands at posture ③) are shown approximately when hitting the ball.	40

4.1	This figure shows the initial and the final trajectories for the two motor control tasks. Both start at 0 and have to go to 1 with minimal accelerations. From $T/2 = 0.75$ on the trajectory has to be as close to 1 as possible. For the passing through task the trajectory additionally has to pass through $p_M = 0.5$ at time $M = 7/40T$ indicated by the circle. . . .	51
4.2	This figure shows the mean performance of all compared methods in two benchmark tasks averaged over twenty learning runs with the error bars indicating the standard deviation. Policy learning by Weighting Exploration with the Returns (PoWER) clearly outperforms Finite Difference Gradients (FDG), ‘Vanilla’ Policy Gradients (VPG), the episodic Natural Actor Critic (eNAC), and the adapted Reward-Weighted Regression (eRWR) for both tasks. Note that this plot has logarithmic scales on both axes, thus a unit difference corresponds to an order of magnitude. The omission of the first twenty rollouts was necessary to cope with the log-log presentation.	52
4.3	This figure shows an illustration of the mountain-car task and the mean accumulated returns of the compared methods.	53
4.4	This figure shows an illustration of the Tetherball Target Hitting task and the mean returns of the compared methods.	54
4.5	This figure shows the time series of the Underactuated Swing-Up where only a single joint of the robot is moved with a torque limit ensured by limiting the maximal motor current of that joint. The resulting motion requires the robot to (ii) first move away from the target to limit the maximal required torque during the swing-up in (iii-v) and subsequent stabilization (vi).	55
4.6	This figure shows the performance of all compared methods for the swing-up in simulation and the mean performance averaged over 20 learning runs with the error bars indicating the standard deviation. PoWER outperforms the other algorithms from 50 rollouts on and finds a significantly better policy.	55
4.7	This figure shows the improvement of the policy over rollouts. The snapshots from the video show the final positions. (0) Initial policy after imitation learning (without torque limit). (1) Initial policy after imitation learning (with active torque limit). (20) Policy after 20 rollouts, going further up. (30) Policy after 30 rollouts, going too far. (40) Policy after 40 rollouts, going only a bit too far. (65) Final policy after 65 rollouts.	56
4.8	This figure shows the performance of all compared methods for the swing-up on the real robot and the mean performance averaged over 3 learning runs with the error bars indicating the standard deviation. PoWER outperforms the other algorithms and finds a significantly better policy.	56
4.9	This figure illustrates the Casting task and shows the mean returns of the compared methods.	57
4.10	This figure illustrates how the reward is calculated. The plane represents the level of the upper rim of the cup. For a successful rollout the ball has to be moved above the cup first and is then flying in a downward direction into the opening of the cup. The reward is calculated as the distance d of the center of the cup and the center of the ball on the plane at the moment the ball is passing the plane in a downward direction. If the ball is flying directly into the center of the cup, the distance is 0 and through the transformation $\exp(-d^2)$ yields the highest possible reward of 1. The further the ball passes the plane from the cup, the larger the distance and thus the smaller the resulting reward.	58

4.11	This figure shows schematic drawings of the Ball-in-a-Cup motion (a), the final learned robot motion (c), as well as a kinesthetic teach-in (b). The arrows show the directions of the current movements in that frame. The human cup motion was taught to the robot by imitation learning with 31 parameters per joint for an approximately 3 seconds long trajectory. The robot manages to reproduce the imitated motion quite accurately, but the ball misses the cup by several centimeters. After approximately 75 iterations of our Policy learning by Weighting Exploration with the Returns (PoWER) algorithm the robot has improved its motion so that the ball goes in the cup. Also see Figure 4.12.	59
4.12	This figure shows the expected return of the learned policy in the Ball-in-a-Cup evaluation averaged over 20 runs.	60
4.13	This figure shows the improvement of the policy over rollouts. The snapshots from the video show the position of the ball closest to the cup during a rollout. (1) Initial policy after imitation learning. (15) Policy after 15 rollouts, already closer. (25) Policy after 25 rollouts, going too far. (45) Policy after 45 rollouts, hitting the near rim. (60) Policy after 60 rollouts, hitting the far rim. (100) Final policy after 100 rollouts.	61
4.14	This figure illustrates the Ball-Paddling task in simulation and on the real robot. The difference between simulation and robotics can be particularly emphasized in this problem where unrealistically many trials were needed on the simulation for reinforcement learning while the real world behavior could be learned by imitation learning. It illustrates the energy-consuming scenario and the difficulties of realistic learning in the presence of contact forces.	63
5.1	This figure illustrates a table tennis task. The situation, described by the state \mathbf{s} , corresponds to the positions and velocities of the ball and the robot at the time the ball is above the net. The meta-parameters $\boldsymbol{\gamma}$ are the joint positions and velocity at which the ball is hit. The policy parameters represent the backward motion and the movement on the arc. The meta-parameter function $\tilde{\boldsymbol{\gamma}}(\mathbf{s})$, which maps the state to the meta-parameters, is learned. .	72
5.2	This figure illustrates a 2D dart throwing task. The situation, described by the state \mathbf{s} corresponds to the relative height. The meta-parameters $\boldsymbol{\gamma}$ are the velocity and the angle at which the dart leaves the launcher. The policy parameters represent the backward motion and the movement on the arc. The meta-parameter function $\tilde{\boldsymbol{\gamma}}(\mathbf{s})$, which maps the state to the meta-parameters, is learned.	73
5.3	This figure illustrates the meaning of policy improvements with Cost-regularized Kernel Regression. Each sample consists of a state, a meta-parameter and a cost where the cost is indicated the blue error bars. The red line represents the improved mean policy, the dashed green lines indicate the exploration/variance of the new policy. For comparison, the gray lines show standard Gaussian process regression. As the cost of a data point is equivalent to having more noise, pairs of states and meta-parameter with low cost are more likely to be reproduced than others with high costs.	76
5.4	This figure shows the performance of the compared algorithms averaged over 10 complete learning runs. Cost-regularized Kernel Regression finds solutions with the same final performance two orders of magnitude faster than the finite difference gradient (FD) approach and twice as fast as the reward-weighted regression. At the beginning FD often is highly unstable due to our attempts of keeping the overall learning speed as high as possible to make it a stronger competitor. The lines show the median and error bars indicate standard deviation. The initialization and the initial costs are identical for all approaches. However, the omission of the first twenty rollouts was necessary to cope with the logarithmic rollout axis. The number of rollouts includes the rollouts not used to update the policy.	77

5.5	This figure illustrates the influence of the parameter λ for the Cost-regularized Kernel Regression. The red curve ($\lambda = 0.5$) corresponds to the red curve (Cost-regularized Kernel Regression) in Figure 5.4(c). The parameter λ trades off the exploration versus the exploitation. A higher λ leads to larger exploration and, thus, faster convergence to a suboptimal solution. The results are averaged over 10 complete learning runs. The lines show the median and error bars indicate standard deviation. The number of rollouts includes the rollouts not used to update the policy.	78
5.6	In this figure, we compare Gaussian process regression (GPR) in a supervised learning setting as proposed by [Ude et al., 2010, Kronander et al., 2011] to Cost-regularized Kernel Regression (CrKR) in a reinforcement learning setting. The red curve corresponds to the red curve (Cost-regularized Kernel Regression) in Figure 5.4. The GPR is trained with samples from the prior used for the CrKR (blue line) and with samples of the final CrKR policy (cyan line) respectively. The black line indicates the cost after CrKR has converged. GPR with samples drawn from the final policy performs best. Please note that this comparison is contrived as the role of CrKR is to discover the policy that is provided to “GPR structured policy”. GPR can only reproduce the demonstrated policy, which is achieved perfectly with 1000 samples. GPR can reproduce the demonstrated policy more accurately if more samples are available. However, it cannot improve the policy according to a cost function and it is impacted by contradictory demonstrations. The results are averaged over 10 complete learning runs. The lines show the median and error bars indicate standard deviation. The number of rollouts includes the rollouts not used to update the policy. . . .	78
5.7	This figure shows a dart throw in a physically realistic simulation.	80
5.8	This figure shows a dart throw on the real JST-ICORP/SARCOS humanoid robot CBi. . . .	80
5.9	This figure shows the cost function of the dart-throwing task for a whole game <i>Around the Clock</i> in each rollout. The costs are averaged over 10 runs with the error-bars indicating standard deviation. The number of rollouts includes the rollouts not used to update the policy.	81
5.10	This figure shows a dart throw on the real Kuka KR 6 robot.	81
5.11	This figure shows the phases of a table tennis stroke on the real Barrett WAM.	82
5.12	This figure shows samples of the learned forehands. Note that this figure only illustrates the learned meta-parameter function in this context but cannot show timing (see Figure 5.13) and velocity and it requires a careful observer to note the important configuration differences resulting from the meta-parameters.	83
5.13	This figure illustrates the effect of the velocity of the ball towards the robot on the time it has until the ball needs to be hit. The plot was generated by sweeping through the velocity component towards the robot, keeping the other position and velocity values fixed. The line is the mean of 100 sweeps drawn from the same ball distribution as used during the learning.	83
5.14	This figure shows samples of the learned forehands on the real robot.	84
5.15	This figure shows the cost function of the simulated table tennis task averaged over 10 runs with the error-bars indicating standard deviation. The red line represents the percentage of successful hits and the blue line the average cost. The number of rollouts includes the rollouts not used to update the policy. At the beginning the robot misses the ball 95% of the episodes and on average by 50 cm. At the end of the learning the robot hits almost all balls.	84
5.16	This figure shows the cost function of the table tennis task on the real robot. The policy was learned entirely on the real robot. The red line represents the percentage of successful hits and the blue line the average cost. The number of rollouts includes the rollouts not used to update the policy. At the beginning the robot misses the ball 70% of the episodes and on average by 15 cm. At the end of the learning the robot hits 80% of the balls. . . .	85

5.17	This figure illustrates the side-stall game. The player throws the ball and if it lands in the target (illustrated by a wall with target holes) gets the number of points written next to it. Missing the targets is not punished, however, going over ten points leads to a loss of ten points.	86
5.18	This figure illustrates the setup of the roles of the different levels.	86
5.19	This figure illustrates the transition probabilities of the three behaviors to their associated outcome in simulation. For example, the red line indicates the probability of gaining one point when throwing at target 1. After approximately 50 throws the player has improved his accuracy level such that he always hits the desired target. The plots are averaged over 10 runs with the error-bars indicating standard deviations.	88
5.20	This figure illustrates the improvement of the player over the number of games in simulation. Due to the large penalty for busting the framework always uses a safe strategy. Already after five completed games the player reaches almost always the maximum possible score of 10. As the number of throws is not punished there are initially many throws that miss the target. After 7 games the number of throws has converged to 4, which is the minimum required number. The plots are averaged over 10 runs with the error-bars indicating standard deviations.	88
5.21	This figure illustrates the setup of the robot evaluation.	89
5.22	This figure illustrates the transition probabilities of the three behaviors to their associated outcome like in Figure 5.19. The skill improves a lot in the first 15 throws after that the improvement levels off. Initially behavior 2, associated with target 2 (which lies in the center) is most likely to succeed. The success rate of 60% corresponds to the level of reproducibility of our setup. The framework manages to handle this large uncertainty by choosing to “stand” early on. The plots are averaged over 4 runs with the error-bars indicating standard deviations.	89
5.23	These frames illustrate one throwing motion with the BioRob.	90
6.1	This figure illustrates the ball-bouncing task on a Barrett WAM. The goal is to keep the ball bouncing on the racket.	93
6.2	This figure illustrates a possible sequence of bouncing the ball on the racket in a schematic drawing, in simulation, and on the real robot.	98
6.3	This figure illustrates the employed strategy for bouncing the ball on the racket. The highest point of the ball trajectory is supposed to coincide with the red target. The racket is always hitting the ball in a fixed height, i.e., the hitting plane. The strategy is to play the ball in a way that the next intersection with the hitting plane is directly below the target and the maximum height of the ball trajectory corresponds to the height of the target. If the bounce works exactly as planned, the ball needs to be hit only once to return to a strictly vertical bouncing behavior.	99

List of Algorithms

4.1	'Vanilla' Policy Gradients (VPG)	45
4.2	episodic Natural Actor Critic (eNAC)	46
4.3	episodic Reward Weighted Regression (eRWR)	48
4.4	EM Policy learning by Weighting Exploration with the Returns (PoWER)	49
4.5	Finite Difference Gradients (FDG)	52
5.1	Meta-Parameter Learning	74
6.1	Learning the Prioritized Control Law	97



List of Tables

2.1	This table illustrates different value function based reinforcement learning methods employed for robotic tasks and associated publications.	14
2.2	This table illustrates different policy search reinforcement learning methods employed for robotic tasks and associated publications.	15
2.3	This table illustrates different methods of making robot reinforcement learning tractable by employing a suitable representation.	17
2.4	This table illustrates different methods of making robot reinforcement learning tractable by incorporating prior knowledge.	21
2.5	This table illustrates different methods of making robot reinforcement learning tractable using models.	24
4.1	Overview of the Experiments: 4.3.2 Basic Motor Learning, 4.3.3 Mountain-Car, 4.3.4 Tetherball Target Hitting, 4.3.5 Underactuated Swing-Up, 4.3.6 Casting, and 4.3.7 Ball-in-a-Cup	50
6.1	This table shows the suitability of the possible dominance structures (mean \pm std). The “hit the ball” primitive clearly is the dominant one, followed by the “move under the ball” primitive. The prioritized control laws work significantly better than a single model learned using the combined training data of the three primitives. Preliminary results on the real robot confirm this ordering.	100



Abbreviations

In this PhD thesis we use the following mathematical notation throughout this thesis:

Notation	Description
$\{x_1, x_2, \dots, x_n\}$	set with elements x_1, x_2, \dots, x_n
\mathbb{R}	real numbers
$\mathbf{x} = [x_1, x_2, \dots, x_n]$	a vector
x_i	the i^{th} component of the vector \mathbf{x}
\mathbf{x}^T	transpose of vector
$\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n]$	a matrix
\mathbf{a}_i	the i^{th} vector of the matrix \mathbf{A}
a_{ij}	the i, j^{th} component of the matrix \mathbf{A}
\mathbf{A}^T	transpose of matrix
\mathbf{A}^{-1}	matrix inverse
\mathbf{A}^+	matrix pseudo-inverse
$\mathbf{A}^{\frac{1}{2}}$	matrix root
$\nabla_{\theta_i} f$	derivative with respect to parameter θ_i
$\nabla_{\boldsymbol{\theta}} f$	derivative with respect to parameters θ_i
$\frac{\partial f}{\partial q}$	partial derivative
$p(\mathbf{x})$	probability density of \mathbf{x}
$E\{\mathbf{x}\}$	expectation of \mathbf{x}
$\bar{\mathbf{x}} = \langle \mathbf{x} \rangle$	sample average of \mathbf{x}

As symbols in this PhD thesis, the following symbols are used in several sections:

Symbol	Description
t	time
$\mathbf{x}, \dot{\mathbf{x}}, \ddot{\mathbf{x}}$	task space position, velocity, acceleration
$\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}$	joint space position, velocity, acceleration
\mathbf{s}_t	state (at time t)
\mathbf{a}_t	action (at time t)
$\mathbf{s}_{1:T+1}$	series of states \mathbf{s}_t with $t \in \{1, 2, \dots, T+1\}$
$\mathbf{a}_{1:T}$	series of actions \mathbf{a}_t with $t \in \{1, 2, \dots, T\}$
$\tau = [\mathbf{s}_{1:T+1}, \mathbf{a}_{1:T}]$	rollout, episode, trial
T	rollout length
$\pi(\mathbf{a}_t \mathbf{s}_t, t)$	policy
\mathbb{T}	set of all possible paths
r_t	reward (at time t)
$\mathbf{r}_{1:T}$	series of rewards r_t with $t \in \{1, 2, \dots, T\}$
$R(\tau)$	return
$J(\boldsymbol{\theta})$	expected return
$D(p q)$	Kullback-Leibler divergence
$Q^\pi(\mathbf{s}, \mathbf{a}, t)$	value function of policy π
$\mathbf{F}(\boldsymbol{\theta})$	Fisher information matrix
ε	exploration
H	number of rollouts
n	index of parameter
k	index of iteration
$\mathbf{g} = \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta})$	gradient
α	update rate
$\boldsymbol{\psi}(\cdot)$	weights
\mathbf{c}	centers
\mathbf{h}	widths
Δt	time step
λ	ridge factor
$k(\mathbf{s}, \mathbf{s}), \mathbf{K}$	kernel, Gram matrix
\mathbf{u}	motor command

Curriculum Vitæ

Educational Background

- 2008-2012 **PhD Student**
Empirical Inference, MPI for Intelligent Systems, Tübingen, Germany
- 2007-2008 **Diplomarbeit**
Empirical Inference, MPI for Biological Cybernetics, Tübingen, Germany
- 2006-2008 **Diplom-Ingenieur in Engineering Cybernetics**
University of Stuttgart, Germany
- 2004-2006 **Ingénieur Centralien**
Ecole Centrale Paris, France
- 2002-2004 **Vordiplom in Engineering Cybernetics**
University of Stuttgart, Germany
- 2001 **Abitur**
Schlossgymnasium Künzelsau, Germany
- 1999 **High School Diploma**
Tulelake, CA, USA

Research Positions

- 2011 **Summer Intern**
Disney Research Pittsburgh, PA, USA
- 2009 **Visiting Researcher**
Advanced Telecommunication Research Center, Kyoto, Japan
- 2005-2006 **Student Researcher**
Project Quadricoptère, Ecole Centrale Paris & Ecole des Mines, France
- 2003-2004 **Student Researcher**
Lehrstuhl für Bauphysik, University of Stuttgart, Germany

Teaching

- 2008 **Lab Assistant, Graduate course: Praktikum Regelungstechnik I**
Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany
- 2011-2012 **Teaching Assistant, Projektpraktikum: Lernende Roboter**
Intelligent Autonomous Systems, Technische Universität Darmstadt, Germany

Work Experience

- 2006 **Internship at Onboard-Controls**
Aérodrome de Blois le Breuil, France
- 2005 **Internship at Ingénieurs Sans Frontières / SOS Sahel**
Bamako & Sanankoroba, Mali
- 2005 **Part-time Work at Aviation Sans Frontières**
Paris, France
- 2001-2002 **Civil Service at Kivuli Centre**
Nairobi, Kenya