

Feature Discovery in Reinforcement Learning using Genetic Programming

Sertan Girgin, Philippe Preux

► **To cite this version:**

Sertan Girgin, Philippe Preux. Feature Discovery in Reinforcement Learning using Genetic Programming. [Research Report] INRIA. 2007. inria-00187997v2

HAL Id: inria-00187997

<https://hal.inria.fr/inria-00187997v2>

Submitted on 19 Nov 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Feature Discovery in Reinforcement Learning using Genetic Programming

Sertan Girgin — Philippe Preux

N° 6358

November 2007

Thème COG



*Rapport
de recherche*

Feature Discovery in Reinforcement Learning using Genetic Programming

Sertan Girgin* , Philippe Preux†

Thème COG — Systèmes cognitifs
Projets SequeL

Rapport de recherche n° 6358 — November 2007 —16 pages

Abstract: The goal of reinforcement learning is to find a policy that maximizes the expected reward accumulated by an agent over time based on its interactions with the environment; to this end, a function of the state of the agent has to be learned. It is often the case that states are better characterized by a set of features. However, finding a “good” set of features is generally a tedious task which requires a good domain knowledge. In this paper, we propose a genetic programming based approach for feature discovery in reinforcement learning. A population of individuals, each representing a set of features is evolved, and individuals are evaluated by their average performance on short reinforcement learning trials. The results of experiments conducted on several benchmark problems demonstrate that the resulting features allow the agent to learn better policies in a reduced amount of episodes.

Key-words: feature discovery, reinforcement learning, genetic programming

* serta.girgin@inria.fr

† philippe.preux@inria.fr

Découverte de caractéristiques dans le problème d'apprentissage par renforcement par programmation génétique

Résumé : L'objectif de l'apprentissage par renforcement est de trouver une politique qui maximise l'espérance des retours obtenus au cours du temps par un agent, en se basant sur ses interactions avec son environnement. Pour cela, une fonction de son état doit être apprise par l'agent. Il est courant qu'un état soit mieux spécifié par un ensemble de caractéristiques (features en anglais). Cependant, trouver un "bon" ensemble de caractéristiques est généralement une tâche très difficile et pénible qui requiert une bonne connaissance du domaine d'application. Dans cet article, nous proposons d'utiliser la programmation génétique pour cela. Une population d'individus, chacun représentant un ensemble de caractéristiques, évolue et les individus sont évalués par la performance d'un agent apprenant par renforcement en utilisant ces caractéristiques, sur de courts épisodes d'apprentissage. Les résultats expérimentaux obtenus sur plusieurs problèmes-tests montrent la faisabilité et l'intérêt pratique de l'approche, une meilleure politique étant apprise plus vite.

Mots-clés : découverte de caractéristiques, apprentissage par renforcement, programmation génétique

1 Introduction

Reinforcement learning (RL) is the problem faced by an agent that is situated in an environment and must learn a particular behavior through repeated trial-and-error interactions with it [18]; at each time step, the agent observes the state of the environment, chooses its action based on these observations and in return receives some kind of “reward”, in other words a *reinforcement signal*, from the environment as feedback. Usually, it is assumed that the decision of the agent depends only on the current state but not the previous ones, i.e. has the Markovian property. The aim of the agent is to find a policy, a way of choosing actions, that maximizes its overall gain. Here, the gain is defined as a function of rewards, such as the (discounted) sum or average over a time period. Unlike supervised learning problem, in RL correct input/output pairs, i.e. optimal action at a given situation, are not presented to the agent, nor sub-optimal actions explicitly corrected. One key aspect of RL is that the rewards can be *delayed* in the sense that immediate rewards received by the agent may not be reflecting the true values of the chosen actions. For example, in the game of chess a move which causes your opponent to capture a piece of yours can be regarded as a “bad” move. However, a series of such moves on purpose may be essential to win the game and consequently receive a higher reward in the future. In a simplified setting, RL as defined above is closely related with the kind of learning and decision making problems that human beings face in their daily lives.

There are two main approaches for solving RL problems. In the first approach, the agent maintains a function $V^\pi(s)$, called *value function*, that estimates the expected return when starting in state s and following policy π thereafter, and tries to converge to the value function of the optimal policy. The policy is inferred from the value function. Alternatively, in *policy space* approaches, policy is represented as a parameterized function from states to actions and optimal policy is searched directly in the space of such functions. There also exist methods that combine both approaches. Note that, in any case, the functions that we are learning (either value function, policy, or both) are naturally functions of the state (observation) variables. However, in a given problem (i) all these variables may not be relevant, which leads to *feature selection* problem, i.e. selecting a subset of state variables, or worse (ii) in their raw form they may be inadequate for successful and/or efficient learning and it may be essential to use some kind of *feature discovery*.

The most obvious situation where the second case emerges is when the number of states is large, or infinite, and each state variable by itself reflects limited and local information about the problem. Let us consider the popular game of Tetris. In Tetris, traditionally each state variable corresponds to the binary (occupied/empty) status of a particular cell of the grid. Not only the number of possible states increases exponentially with respect to the size of the grid, but also each state variable tells very little about the overall situation. A human player (most successful computer players

as well) instead takes into consideration more informative features that are computable from the state variables, such as the height of each column or the number of holes in the occupied regions of the grid, and decides on his actions accordingly. Similar reductions are also quite common in other domains, such as image processing applications where instead of the raw image various high level features derived from it are fed into learning algorithms. On the other end of the spectrum, in some cases, additional features can be useful to improve the performance of learning. An example of this situation is presented in Figure 1 for the classical cart-pole balancing problem. By adding sine and cosine of the pole’s angle as new features to existing state variables, optimal policy can be attained much faster. A related question is, of course, given a problem what these features are and how to find them. Note that feature discovery, which will also be our main objective, is a more general problem and includes feature selection as a special case.

Usually, the set of features that are to be used instead of or together with state variables are defined by the user based on extensive domain knowledge. They can either be fixed, or one can start from an initial subset of possible features and iteratively introduce remaining features based on the performance of the current set, so called *feature iteration approach* [3]. However, as the complexity of the problem increases it also gets progressively more difficult to come up with a good set of features. Therefore, given a problem, it is highly desirable to find such features automatically in an unsupervised fashion solely based on the observations of the agent. In this paper, we explored using a genetic programming based approach for that purpose.

The rest of the paper is organized as follows: In Section 2, we review related work on feature discovery in RL and GP based approaches in other fields. Section 3 describes our method in detail. Section 4 presents some empirical evaluations of our approach on some benchmark problems. We conclude in Section 5 with a discussion of results and future work.

2 Related Work

Feature discovery, in essence, is an information transformation problem; the input data is converted into another form that “better” describes the underlying concept and relationships, and “easier” to process by the agent. As such, it can be applied as a preprocessing step to a wide range of problems and it has attracted attention from other fields particularly for classification. In [10], Krawiec studies the change of representation of input data for machine learners and genetic programming based construction of features within the scope of classification. Each individual encodes a *fixed* number of new feature definitions expressed as S-expressions. In order to determine the fitness of an individual, first a new data set is generated by computing feature values for all training examples and then a classifier (decision tree learner) is trained on this data set. The resulting average accuracy of classification becomes the evaluation of the individual. He also proposes an




Figure 1: Policy gradient with Rprop update and optimal baseline on the cart-pole problem using a linear parameterized policy (see [13] for more detailed discussion). In the cart-pole problem, the objective is to hold a pole, which is attached to a cart moving along a track, in upright position by applying certain amount of force to the cart. The state variables are the pole’s angle and angular velocity and the cart’s horizontal position and velocity. (a) Learning performance when sine and cosine of the angle of the pole are added as new features. (b) Performance of the features found by GP.

extended method in which each feature of an individual is assigned a utility that measures how valuable that feature is and the most valuable features are not involved in evolutionary search process. The idea behind this extension is to protect valuable features from possible harmful modifications so that the probability of accidentally abandoning promising search directions would be reduced. It is possible to view this extension as an elitist scheme at the level of an individual. While Krawiec’s approach has some similarities with our approach presented in this paper, they differ in their focus of attention and furthermore we consider the case in which the number of features is not fixed but also determined by GP.

Smith and Bull [16] have also used GP for feature construction in classification. However, they follow a layered approach: in the first stage, a fixed number of new features (equal to the number of attributes in the data set subject to a minimum 7) is generated as in Krawiec (again using a decision tree classifier and without the protection extension), and then a genetic algorithm is used to select the most predictive ones from the union of new features and the original ones by trying different combinations. Although their method can automatically determine the number of features after the second stage, in problems with large number of attributes the first stage is likely to suffer as it will try to find a large number of features as well (which consequently affects the second stage).

In RL, Sanner [14] recently introduced a technique for online feature discovery in relational reinforcement learning, in which the value function is represented as a ground relational naive Bayes net and structure learning

is focused on frequently visited portions of the state space. The features are relations built from the problem attributes and the method uses a variant of *Apriori* data mining algorithm to identify features that co-occur with a high frequency and creates a new joint feature as necessary.

3 Feature Discovery in RL using GP

Genetic programming (GP) is an evolutionary algorithm based methodology in which each individual represents a computer program and a population of individuals is progressively evolved under the influence of several genetic operators (such as crossover and mutation) and evaluated over a series of generations [7]. Evolution and “survival of the fittest” dynamics results in individuals having higher fitness scores, i.e. programs that better solve the posed problem. GP is shown to be quite effective on various domains and in some cases produced results that are competitive with or surpass human performance [9].

When GP is being applied to a particular problem, there are three main issues that need to be addressed:

1. structure and building blocks (i.e. primitive functions and terminals) of the individuals,
2. set of genetic operators, and
3. fitness function.

In our case, due to the fact that we are interested in identifying useful features for a given RL problem, each individual must essentially be a program that generates a set of features based on the state variables. Therefore, state variables are the *independent variables* of the problem and are included in the set of terminals together with (ephemeral) constants and possible problem specific zero argument functions. An individual consists of a list of S-expressions, called *feature-functions*, such that each S-expression corresponds to a unique feature represented as a function of various arithmetic and logical operators and terminals. Given the values of state variables, the agent can calculate the value of features by evaluating the S-expressions. In our implementation, we linearized each S-expression in prefix-order and then concatenated together to obtain the final encoding of the individual. This compact form helps to reduce memory requirements and also simplifies operations. As we will describe later, each individual is dynamically compiled into executable binary form for evaluation and consequently this encoding is accessed/modified only when genetic operators are applied to the individual.

Note that the number of useful features is not known a priori (we are indeed searching for them) and has to be determined. Instead of fixing this number to an arbitrary value, we allowed the individuals to accommodate varying number of feature functions (S-expressions) within a range, typically less than a multiple of the number of raw state variables, and let




Figure 2: Single point cross-over on feature lists of two individuals. Cross-over point is shown by vertical dashed line. The number of features represented by the off-springs differ from that of their parents.

Cross-over One of the nodes in any feature function of an individual and the whole branch under it is switched with another node from another individual in the population.

Mutation (node) One of the nodes in any feature function of an individual is substituted with another nodes compatible one – a terminal or a zero argument function is replaced with a terminal or a zero argument function, and an n -ary operator is replaced with an n -ary operator. Note that the branch under the mutated node, if any, is not affected.

Mutation (tree) One of the nodes in any feature function of an individual and the whole branch under it is substituted with a new randomly generated sub-tree having a depth of 3 or less.

Shrinkage One of the operator nodes in any feature function of an individual is substituted with one of its children.

Feature-list cross-over See text and Figure 2.

Table 1: Genetic operators.

the evolutionary mechanism search for an optimal value. To facilitate the search, in addition to regular genetic operators presented in Table 1 we also defined a single-point crossover operator over the feature function lists of two individuals. Let n and m be the number of features of two individuals selected for cross-over, and $0 < i < n$ and $0 < j < m$ be two random numbers. The first i features of the first individual are merged with the last $m - j$ features of the second individual, and the first j features of the second individual are merged with the last $n - i$ features of the first individual to generate two off-springs (Figure 2). The generated off-springs contain a mixture of features from both parents and may have different number of features compared to them.




Figure 3: Speedup for the evaluation of functions in the form of complete binary trees having depth 2-8. Each function is executed 1000000 times and the results are averaged over 10 independent runs).

Since our overall goal is to improve the performance of learning, the obvious choice for the fitness of an individual is the expected performance level achieved by the agent when the corresponding feature functions are applied on a particular RL algorithm. In this work, we opted for two different algorithms, namely λ policy iteration and policy gradient method with RProp update; they are described in more detail in Section 4. In both RL algorithms, we represented the value function and the policy as a linear combination of feature functions, hence the parameters correspond to the coefficients of each feature function. The fitness scores of individuals are calculated by taking their average performance over a small number (around 4-10) of short learning trials using the corresponding RL algorithm. In the experiments, we observed that both algorithms converges quickly towards an approximately optimal policy when the basis feature functions capture the important aspects of the complicated nonlinear mapping between states and actions. We also penalized feature functions according to their size to avoid very large programs, but in practice we observed that this penalization had very little effect as feature sets consisting of simpler functions tend to perform better and receive higher scores.

Accelerating the Computations

It is well known that GP is computationally demanding. In our case, which also applies in general, two bottlenecks can be identified: (i) the time required to execute the program represented by an individual, and (ii) the need to evaluate many individuals in each generation.

During the evaluation of a single individual, feature functions are called repeatedly for different values of state variables in order to calculate the corresponding feature values. If at each call, the actual tree structure of each feature function (or its linear form) is interpreted directly by traversing the S-expression, much time is spent in auxiliary operations such as following nodes, pushing/popping values onto the stack, parsing node types

etc. This overhead can easily, and in fact eventually, become a bottleneck as the size of the individuals (i.e the number of nodes) and the number of calls (in the order of thousands or more) increase. Several approaches have been proposed to overcome this problem, such as directly manipulating machine language instructions as opposed to higher level expressions [12, 1] or compiling tree representation into machine code [5]. Following the work of Fukunaga et.al. [5], we dynamically generate machine code at runtime for each feature function using GNU Lightning library, and execute the compiled code at each call. GNU Lightning is a portable, fast and easily retargetable dynamic code generation library [6]. It defines a standardized RISC instruction set with general-purpose integer and floating point registers, and abstracts the user from the target CPU. As it translates code directly from a machine independent interface to that of the underlying architecture without creating intermediate data structures the compilation process is very efficient and requires only a single pass over the tree representation or linearized form of an individual¹. Furthermore, problem specific native operators or functions (mathematical functions etc.) can be easily called from within compiled code. Figure 3 shows the speedup of an optimized implementation of standard approach compared to the dynamically compiled code on randomly generated functions that have a complete binary tree form (i.e. contains $2^d - 1$ nodes where d is the depth of the tree). The speed-up increases with the size of the functions, reaching a high of about 75 fold performance improvement which is substantial.

In GP, at each generation the individuals are evaluated independently of each other, that is the evaluation process is highly parallelizable. As such, it can be implemented efficiently on parallel computers or distributed computing systems. By taking advantage of this important property, we developed a parallel GP system using MPICH2 library [11], which is an implementation of the *Message Passing Interface*, and run the experiments on a Grid platform. This also had a significant impact on the total execution time.

4 Experiments

We evaluated the proposed GP based feature discovery method on three different benchmark problems: Acrobot [17], multi-segment swimmer [4] and Tetris [3] (Figure 4). The first two problems, Acrobot and multi-segment swimmer, are dynamical systems where the state is defined by the position and velocity of the elements of the system, and action being an acceleration which, according to Newton’s law, defines the next state. These are non-linear control tasks with continuous state and action spaces; the number of state variables are respectively 4 and $2n + 2$ where n is the number of segments of the swimmer. Despite their seemingly easiness, these two tasks are difficult to learn (to say the least, far from obvious). In Acrobot, there is only a single control variable, whereas in swimmer the agent has to decide on torques applied to each of $n - 1$ joints. Although it is similar in

¹Time to compile a function of 64 nodes is around 100 microseconds on a 2.2Ghz PC.




Figure 4: (a) Acrobot (figure taken from [18]), (b) multi-segment swimmer, and (c) Tetris.

nature to Acrobot, swimmer problem has significantly more complex state and control spaces that can be varied by changing the number of segments. As the number of segments increase the problem also becomes harder. Our third benchmark problem, the game of Tetris, has discrete state and action spaces. The state variables are the following: (i) the heights of each column, (ii) the absolute difference between the heights of consecutive columns, (iii) the maximum wall height, (iv) the number of holes in the wall (i.e. the number of empty cells that have an occupied cell above them), and (v) the shape of the current object. We used a 12×8 grid and 7 different shapes that consist of 4 pieces, thus in our case the number of features is 18.

For evaluating the individuals and testing the performance of the discovered features, we employed two different RL algorithms: λ policy iteration for the Tetris problem, and policy gradient method with RProp update for the Acrobot and swimmer problems. λ policy iteration is a family of algorithms introduced by Ioffe and Bertsekas which generalizes standard value iteration and policy iteration algorithms [2]. Value iteration starts with an arbitrary value function and at each step updates it using the Bellman optimality equation (with one step backup); the resulting optimal policy is greedy with respect to the value function². On the other hand, policy iteration starts with an initial policy and generates a sequence of improving policies such that each policy is greedy with respect to the estimated value (calculated by policy evaluation) of its predecessor. λ policy iteration fuses both algorithms together with a parameter $\lambda \in (0, 1)$ by taking a λ -adjustable step toward the value of next greedy policy in the sequence [15]. We used the approximate version of λ policy iteration as defined in Section 8.3 of [3]. Policy gradient method also works on the policy space, but approximates a parameterized (stochastic) policy directly. Starting from an initial policy, the policy parameters are updated by taking small steps in the direction of the gradient of its performance and under certain conditions converge to a local optima in the performance measure [19]. The gradient is usually estimated using Monte Carlo roll-outs. With RProp update, instead of directly relying on the magnitude of the gradient for the updates (which may lead to slow convergence or oscillations depending on the learning rate), each parameter is updated in the direction of the corresponding

²For example, as the policy that selects in each state the action with highest estimated value.




Figure 5: Results for Acrobot. (a) Performance of discovered features, and (b) fitness values of best individuals in each generation.

partial derivative with an individual time-varying value. The update values are determined using an adaptive process that depends on the change in the sign of the partial derivatives.

In the experiments, a population consisting of 100 individuals evolved for 50 generations. We set crossover probability to 0.7 (with a ratio of 1/6 for the feature-list crossover), and the remaining 0.3 is distributed among mutation and shrinkage operators. Node mutation is given two times higher probability than the others. There is a certain level of elitism, 10% of best performing individuals of each generation are directly transferred to the next generation. The set of operators is $\{+, -, *, /, \sin, \cos, \sqrt{\cdot}\}$ for the Acrobot and swimmer problems and $\{+, -, *, /, \min, \max, \|\cdot\|\}$ for the Tetris problem where $\|\cdot\|$ denotes the absolute difference between two values. The terminals consists of the set of original state variables as given above and $\{1, 2, e\}$ where e denotes an ephemeral random constant $\in [0, 1]$. In the policy gradient method (Acrobot and swimmer problems), an optimal baseline is calculated to minimize the variance of the gradient estimate, and the policy is updated every 10 episodes. We tested with two different sets of parameters: $(\Delta_{min} = 0.01, \Delta_{ini} = 0.1, \Delta_{max} = 0.5)$ and $(\Delta_{min} = 0.002, \Delta_{ini} = 0.02, \Delta_{max} = 0.1)$. In λ policy iteration (Tetris), we run 30 iterations and sampled 100 trajectories per iteration using the greedy policy at that iteration. λ is taken as 0.6. The results presented here are obtained using a single GP run for each problem. The discovered features are then tested on the same RL algorithms but with a longer training period (50000 iterations for policy gradient, and 30 iterations for λ policy iteration) to verify how well they perform. We averaged over 20 such test trainings.

Figure 5a and Figure 6[a,c] show the testing results for the Acrobot and multiple-segment swimmer problems, respectively. In both cases, features found by GP show an improvement over original features and allow the agent to learn policies with larger return. The improvement is more evident in swimmer problem, where the agents utilizing the discovered features can learn policies that perform on average 50% better. Since candi-

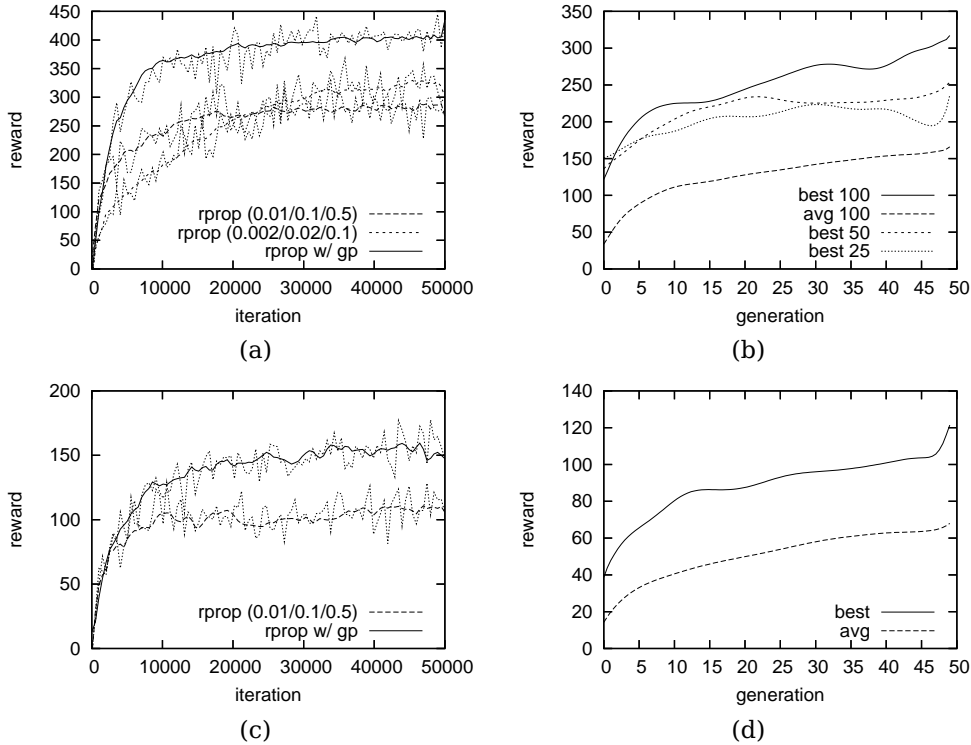


Figure 6: Results for 3 (figures a and b) and 5 (figures c and d) segment swimmers. (a, c) Performance of discovered features, and (b, d) fitness values of best individuals and average fitness of the population in each generation. Figure b also shows the fitness values of best individuals for the population sizes of 25 and 50.

date feature-functions are evaluated based on their average performances on short learning trials, learning speed is also faster in the initial stages as expected. The fitness values of best individuals and average fitness of the population in each generation during the feature discovery process indicate that the evolutionary search drives towards better solutions and further improvement may be possible with longer GP runs especially in swimmer problem (Figure 5b and Figure 6[b,d]). We obtained inferior results with smaller population sizes (Figure 6b).

Although we used a different RL algorithm, the results for Tetris are also similar to those of Acrobot and swimmer, and show considerable improvement in terms of the performance of the resulting policies (Figure 7).

5 Conclusion

In this paper, we explored a novel genetic programming based approach for discovering useful features in reinforcement learning problems. Empirical




Figure 7: Results for Tetris. (a) Performance of discovered features, and (b) fitness values of best individuals and average fitness of the population in each generation.

results show that evolutionary search is effective in generating functions of state variables that when fed into RL algorithms allow the agent to learn better policies. As supported by previous results in classification tasks, the approach may also be applicable in supervised settings by changing the learning algorithm used for evaluating the individuals. However, care must be taken to choose algorithms that converge quickly when supplied with a “good” state representation.

One important point of the proposed method is that it allows the user to guide the search and if needed incorporate domain knowledge simply by specifying the set of program primitives (i.e. ingredients of the feature functions). Furthermore, resulting feature functions are readable by humans (and not hard to comprehend) which makes it possible to fine-tune and also transfer knowledge to (feature extraction process of) similar problems. This can be done either manually, or by converting them into meta functions, as in *automatically defined functions* [8], leading to a hierarchical decomposition. We pursue future research in this direction.

Acknowledgment

Experiments presented in this paper were carried out using the Grid’5000 experimental testbed, an initiative from the French Ministry of Research through the ACI GRID incentive action, INRIA, CNRS and RENATER and other contributing partners (see [https:// www.grid5000.fr](https://www.grid5000.fr))

References

- [1] Wolfgang Banzhaf, Frank D. Francone, Robert E. Keller, and Peter Nordin. *Genetic programming: an introduction: on the automatic evo-*

- lution of computer programs and its applications*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1998.
- [2] D. Bertsekas and S. Ioffe. Temporal differences-based policy iteration and applications in neuro-dynamic programming. Technical Report LIDS-P-2349, MIT, 1996.
- [3] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.
- [4] Rémi Coulom. *Reinforcement Learning Using Neural Networks, with Applications to Motor Control*. PhD thesis, Institut National Polytechnique de Grenoble, 2002.
- [5] Alex Fukunaga, Andre Stechert, and Darren Mutz. A genome compiler for high performance genetic programming. In John R. Koza, Wolfgang Banzhaf, Kumar Chellapilla, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max H. Garzon, David E. Goldberg, Hitoshi Iba, and Rick Riolo, editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 86–94, University of Wisconsin, Madison, Wisconsin, USA, 22-25 1998. Morgan Kaufmann.
- [6] GNU. *Lightning*, 2007. Available from <http://www.gnu.org/software/lightning/>.
- [7] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [8] John R. Koza. *Genetic programming II: automatic discovery of reusable programs*. MIT Press, Cambridge, MA, USA, 1994.
- [9] John R. Koza, Martin A. Keane, Matthew J. Streeter, William Mydlowec, Jessen Yu, and Guido Lanza. *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers, Norwell, MA, USA, 2003.
- [10] Krzysztof Krawiec. Genetic programming-based construction of features for machine learning and knowledge discovery tasks. *Genetic Programming and Evolvable Machines*, 3(4):329–343, 2002.
- [11] Argonne National Laboratory. *Mpich2*, 2007. Available from <http://www-unix.mcs.anl.gov/mpi/mpich2/>.
- [12] Peter Nordin. A compiling genetic programming system that directly manipulates the machine code. In Kenneth E. Kinneer, Jr., editor, *Advances in Genetic Programming*, chapter 14, pages 311–331. MIT Press, 1994.
- [13] M. Riedmiller, J. Peters, and S. Schaal. Evaluation of policy gradient methods and variants on the cart-pole benchmark. pages 254–261, 04 2007.

-
- [14] Scott Sanner. Online feature discovery in relational reinforcement learning. In *Open Problems in Statistical Relational Learning Workshop (SRL-06)*, 2006.
 - [15] Bruno Scherrer. Performance bounds for lambda policy iteration. 2007.
 - [16] Matthew G. Smith and Larry Bull. Genetic programming with a genetic algorithm for feature construction and selection. *Genetic Programming and Evolvable Machines*, 6(3):265–281, September 2005. Published online: 17 August 2005.
 - [17] Mark W. Spong. Swing up control of the acrobot. In *ICRA*, pages 2356–2361, 1994.
 - [18] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998. A Bradford Book.
 - [19] Richard S. Sutton, David A. McAllester, Satinder P. Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *NIPS*, pages 1057–1063, 1999.

Contents

1 Introduction	3
2 Related Work	4
3 Feature Discovery in RL using GP	6
4 Experiments	9
5 Conclusion	12



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399