

Automated Optimization of Walking Parameters for the Nao Humanoid Robot

Nicolò Girardi ^a Chiel Kooijman ^a Auke J. Wiggers ^a Arnoud Visser ^a

^a *University of Amsterdam, Science Park 904, 1098 XH Amsterdam*

Abstract

This paper describes a framework for optimizing walking parameters for a Nao humanoid robot. In this case an omnidirectional walk is learned. The parameters are learned in simulation with an evolutionary approach. The best performance was obtained for a combination of a low mutation rate and a high cross-over rate.

1 Introduction

The main focus of the study is improving the walking behavior of Nao robots by optimizing a set of parameters for an existing walking engine. Our research is highly relevant for the soccer competition; the RoboCup Standard Platform League (SPL), where a fast walk is essential for winning the World Cup. Yet, a fast walk has also a broader value. Walking allows robots to move through difficult terrain (without the force a caterpillar exercises on the environment), and enables them to reach and function in places designed for humans. When interacting with humans they expect from a robot that it can move at the same speed. Currently most humanoid robots can only interact with elderly people.

This work can also be interpreted as a test-case for the artificial intelligence: To study the effectiveness of different machine learning methods in real life environments, which means taking into consideration factors such as field imperfections, sensor error, error in joint angles, and overheating of servos. To facilitate this research, a specific framework was created, which allows researchers to manage the learning process and obtain data for evaluation.

The RoboCup's aim is to promote Robotics and AI research in a playful and entertaining fashion that would appeal to non-expert people and potential future generations of researchers [5]. The event is also an occasion for insiders to share and expand their knowledge. The competition factor plays a role by creating an environment which stimulates improvement. The fact that contending in the SPL consists mainly of writing software, helps this sharing process, as code from the winning team has to be published after the end of the competition.

2 Related work

Multiple methods that combine elements of learning in real life and learning in simulation have been proposed. Haasdijk remarks that in online learning, the average performance of the population is more important than the best performing individual, as there is always a population, so each individual should perform well [3]. Bongard et al. suggest combining simulation and real-life evaluation to find an approximate evaluation in the simulator before assessment in the real world. This way the poorest performers may be prevented from influencing the average performance [1]. Nordin and Banzhaf note that genetic algorithms can be used




Figure 1: Overview of the 25 joints of the Nao robot (Courtesy Aldebaran).

in an on-line setting, where the evaluation will be different for each individual, as the averaging tendencies even out the effects of probabilistic sampling, resulting in a set of good solutions [8].

MacAlpine and Stone have been successful in the use of a genetic algorithm to teach a Nao robot to walk. Their approach uses the CMA-ES policy search, and did an extensive search of a policy space consisting of 14 out of over 40 parameters [6]. However, for this method to show any result, a huge number of rollouts (simulation runs) is needed. This is not necessarily true for some other gradient methods and genetic algorithms, which can show improvement early on (even before convergence). Both have been applied by Meijer, who tested performance of methods he refers to as Finite Difference, Vanilla, Episodic Natural Actor Critic and a genetic algorithm, in learning an optimal kicking motion for a Nao [7]. It should be noted that these methods were applied to improve performance of an already existing kicking motion.

2.1 Genetic algorithms

Genetic algorithms are a name for machine learning methods that aim to optimize through means similar to biological evolution. A basic genetic algorithm contains at least the following:

- Properties that define an organism: The ‘genes’ that can be altered to create new organisms.
- A population of organisms that is to be evaluated.
- A fitness function and corresponding evaluation tasks that specifies how these organisms are evaluated.
- A way to generate new generations based on old ones.
- Learning parameters that influence how new generations are formed.

The algorithm evaluates the population, and based on each organisms performance, selects ‘parents’ out of which new generations can be created, similar to the biological process of natural selection. Through the fitness function and corresponding evaluations, the search can be guided. A disadvantage of this method is that the fitness function must be user-specified, which may be hard to do in tasks with multiple objectives or in cases where the criteria for performance are unclear. The ways in which new generations are generated are similar to biological generation of offspring: Genes of either parent are transmitted directly, new genes are derived through crossovers by taking a (weighted) mean of genes of the two, or entirely new genes are created through mutation by adding random elements to the organism.

For each generation there can also be a part of individuals that is directly copied to the next one (the organisms survive). It’s important to set a proper number of survivors and parents in order to have the highest quality features of fittest individuals pass to the next generation without however limiting the search

space to a small area. If too many organisms survive, it could be possible that these and their offspring dominate the concurrent generations in a few timesteps, causing a local optimum to be reached.

3 Approach

3.1 Platform

The robot used in SPL is Nao (version H25) by Aldebaran [2], a humanoid, bipedal robot. It has inertia sensors that enable fall-detection, and a total of eight resistance sensors in the two feet that allow for early detection of imbalance. Both of these are used in the evaluation of parameter settings. Figure 1 shows an overview of the Nao's joints. It should be noted that the training and evaluation were both performed in simulation using models of this particular robot, due to the high number of iterations needed for these to be effective.

3.2 Walking engine

In the RoboCup 2013 we used the NaoTH framework for Nao by Berlin United [4] for various reasons: it features a modular code which makes it easy to expand, and provides a walking engine from which we can select a set of parameters for learning.

The walking system is closed loop, and uses a set of more than thirty parameters. A subset of this set was selected manually for learning, as these were likely to contribute to a stable and fast walk than others. Some examples are:

singleSupportTime Defines how long the Nao stands on a foot for during the walking. This determines the walking frequency.

stepHeight Defines the maximum distance between the ground and the base of the feet. If not too high improves stability, but given the low weight of the robots, if set over a certain limit, makes the robot bounce and actually lose stability. It also requires more work by the joints which translates in overheating and joints' efficiency decreasing.

maxStep _{x,y,θ} Defines the maximum length (for walking forward), width (for strafing left and right) and angle (for turning left and right). Given a certain singleSupportTime, those parameters determine the speed of the movements. If set too high the result is a loss in stability.

bodyOffset _{x} The body's offset in x -direction (forward). Setting this to a positive number will result in positioning the body forward during walking, which may contribute to faster walking in the x -direction.

bodyPitchOffset Indicates how much the body leans forward during walking with respect to the hips and legs. Similar to bodyOffset in that it can contribute to fast walks in x -direction.

CoMHeight Height of the center of mass of the robot during walking. The lower the body will be, the more stable the walk. However, the load on the leg joints will be higher and steps must be much smaller than for high center of mass.

The problem of selecting values for these parameters is to be solved using a genetic algorithm, as setting and testing these parameters manually is a time-consuming task and does not guarantee selection of an optimal parameter set. The differences between robots (for example actuator and sensor error) cause parameter settings to have a different effect on each robot. Also, the use of machine learning (and, more specifically, reinforcement learning) to teach robots to perform complex tasks is a relevant topic in AI research, and is a problem that is not limited to Nao robots, but applies to many other situations.

4 Implementation

For the use of machine learning methods described in the previous section, the NaoTH framework was extended with a module ‘MachineLearning’, which serves as a generic input for different machine learning methods. It consists of multiple classes:

MachineLearning The module class itself handles debug requests from users.

LearnToWalk The evaluation program used by the module to test performance of specific parameter sets.

MachineLearningMethod An abstract class that contains pure virtual functions, to be implemented by classes corresponding to specific machine learning methods.

This module can be controlled through debug requests. Alongside the module, an interface was created for easy parameter selection and the display of information during learning.

4.1 Learning Parameters

To set the genetic algorithm we use specific parameters:

Crossover rate Defines the weight for crossover.

Mutation rate Defines the weight for mutation.

Transmission rate Defines the weight for transmission.

maxGeneration The maximum number of generations that is be evaluated before the algorithm ends.

ResettingTime Time before the Nao touches the ground after a previous evaluation has ended.

RunningTime Duration of an evaluation.

StandingTime Time during the Nao stays still after touching the ground before starting a new evaluation.

It’s used to guarantee stability for the body when starting an evaluation.

Survivors The number of individuals that are copied from a generation to the next one.

Parents The number of individuals from a generation that are used to obtain a new generation.

PopulationSize The number of individuals for each generation.

4.2 Evaluation Tasks

Four different evaluation tasks were used. Walking (forward, then backward) is used to emphasize speed, strafing (left, then right) is incorporated because this is a task that occurs quite frequently during matches. Turning (clockwise, then counterclockwise) is also used quite often during matches. Walk diagonally (walk forward/backward + strafe left/right) was added to the evaluation tasks to ensure that the robot will be able to compensate for overshoot when walking straight towards an object. These four tasks can be combined to obtain more realistic and challenging action sequences.

These tasks can be combined to obtain more realistic and challenging action sequences.

4.3 Fitness Function

The walking distance is at the base of the fitness function used in the experiments. To emphasize that the robot should be able to walk in any direction, each evaluation consists of a set of tests, which can include walking forward, backward, sideways (strafing), rotating, and walking diagonally in any direction, we use the weighted harmonic mean (equation 1). It rewards increases closer to 0 more than at higher values, which should result in a behavior that should at least be able to perform each task to a certain minimum level.

$$\frac{\sum_{i=1}^n w_i}{\sum_{i=1}^n \frac{w_i}{x_i}} \quad (1)$$

The weights allow for focusing more on specific tasks such as walking forwards. The framework allows for giving a negative reward when the robot falls over. The individual can be killed to save time and reduce the chance of damage in the case of testing in the real world.

4.4 Graphical User Interface

The learning process is managed and controlled through a GUI implemented in Java (figure 2). Its main functions are:

Learn Starts the genetic algorithm

List Shows the list of learning parameters currently in use

GetInfo Shows relevant information such as the current generation number, the current individual number and the walking parameters for the fittest individual in the current generation.




Figure 2: Example of the graphical user interface (right) and SimSpark (left).

4.5 Simulator

The used simulator is SimSpark¹, which is the official 3D-simulator for the RoboCup Simulation League. It was designed to be able to simulate a large number of robots, which proves useful in machine learning tasks. A disadvantage of this simulator is that the models are slightly different from the real robots, so that learned parameters cannot be used directly without some tweaking. A second disadvantage is that the simulator does not guarantee identical results for repeated simulation runs, as it is influenced by network delays on the machine that disrupts the server.

5 Results

To test the performance of the implemented framework, two separate experiments were conducted. Both will be described below. For both of these, the performance on each task is measured by taking the traversed distance as an evaluation score. Falling down is penalized, and results in subtracting 1000 from fitness (empirically chosen), so that organisms that fall early in the episode are punished more severely than those that fall at the end of their task.

¹http://simspark.sourceforge.net/wiki/index.php/Main_Page

5.1 Experiment: Walking forward

To test the effectiveness of the method, it was first applied to improve the walk in the x -direction. The effect of different parameter settings can be seen in figure 3. The settings for these tests were as follows:

Figure	Transmission	Crossover	Mutation	# Parents	# Survivors	PopulationSize
Mutation	0.1	0.1	0.8	10	10	50
NoMutation	0.5	0.5	0.0	10	10	50
Crossover	0.1	0.8	0.1	10	10	50
Equal	0.1	0.1	0.1	10	10	50

It should be noted that the first generation of each test contains only copies of the first individual, making the performance for this generation prone to error.




Figure 3: Average fitness for different parameter settings.

5.2 Experiment: Holonomic walk

In this experiment, machine learning was applied to improve the omnidirectional walk of the Nao. Evaluation was performed by forcing Nao to execute multiple different subtasks, which were to be performed in quick succession - with no pause in between tasks. The tasks were chosen so that they would closely resemble actions during a football match. There was no bias toward any of the evaluations, with each phase lasting 7 seconds, although the order of the tasks would always be the same. Falling down was penalized, and caused the episode to finish immediately. A parameter set was evaluated based on performance in each of these tasks:

- Walking forward/backward
- Turning
- Walking towards a static target
- Quick stopping/starting and direction changes

The left graph in figure 4 shows how the average fitness of the population changes over time. The graph suggests that the fitness is generally increasing in the first 25 generations. The graph on the right of figure 4 shows the fitness of the fittest individual in the population over time, which is monotonically increasing.

Reason for this is that the surviving group of each population, of which the fittest individual is part of, is not re-evaluated, as suggested by [3].




Figure 4: Plot showing fitness of populations for a total of 40 generations.

5.3 Evaluation

Results for both experiments indicate an increase in performance, indicating that by learning the walk is optimized. As can be seen in figure 3 (Mutation versus Crossover), it proved useful to limit mutation, as random trials usually ended up in organisms with lower fitness than generations created using different settings. The test NoMutation is used as a benchmark, needed because of noise in the evaluation, by setting the Mutation rate to zero. This ensures that each generation contained only exact copies of the very first organism. Its low performance at generation 0 can be explained by this noise as well, as the first generation contained only copies of the initial organism. Although setting Transmission, Crossover and Mutation rate equal resulted in nearly the same performance, setting the Crossover rate relatively high resulted in higher fitness.

The performance in the second experiment cannot be directly compared to the first experiment; however, the relative increases in performance can be compared. The second experiment shows a more steady increase in performance. This can be explained by the complexity of the task: The chance that a mutation improves performance in just the task of walking forward is higher than this mutation improves performance in multiple subtasks simultaneously. Results indicate that the selected method and chosen parameter settings improve mean performance of the generation over time and will continue to do so.

6 Conclusion

In this study, the effectiveness of evolutionary algorithms is tested by applying one of these to optimization of walking parameters for a simulated bipedal robot. At the start of the learning process, a working set of walking parameters is available in an existing framework, meaning that not the joint values but the parameters for the framework must be optimized.

The algorithm and the system controlling the experiment are discussed shortly. Different parameter settings are tried to learn the optimal parameter-set for a single task, namely walking forward. Based on this first experiment, the best found settings are used to learn an optimal parameter-set for a holonomic walk. Results for both experiments indicate that there is indeed an improvement in performance, but that starting with a suboptimal set makes it so that the resulting increase in performance is small.

7 Discussion and Future Work

Since old individuals are not re-evaluated and the evaluation function is noisy, it is likely that the population's evolution is inhibited after an individual has gained a high fitness by chance. Re-evaluating should even out irregularities in the fitness and increase the overall performance [8].

The framework is eventually to be extended to facilitate on online learning, specifically learning during matches. This would allow for more accurate evaluation of fitness, as the tests in the lab setting are not the same as a real match. It would also allow for more optimal usage of time during an event, as the robots must be allowed to cool down before matches, and differences in carpet make evaluation before arriving at a RoboCup event inaccurate.

Another suggestion is to classify different kind of actions with different parameter sets, as parameters that work well for walking forward seem to be different from parameters that perform well on other tasks. This would, however require changes to the framework in order to transition smoothly and without falling.

References

- [1] Josh Bongard, Victor Zykov, and Hod Lipson. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006.
- [2] David Gouaillier, Vincent Hugel, Pierre Blazevic, Chris Kilner, Jérôme Monceaux, Pascal Lafourcade, Brice Marnier, Julien Serre, and Bruno Maisonnier. Mechatronic design of nao humanoid. In *IEEE International Conference on Robotics and Automation (ICRA'09)*, pages 769–774, 2009.
- [3] Evert Willem Haasdijk. *Never Too Old To Learn: On-line Evolution of Controllers in Swarm-and Modular Robotics*. PhD thesis, Vrije Universiteit Amsterdam, 2012.
- [4] Verena Hafner, Hans-Dieter Burkhard, Heinrich Mellmann, Thomas Krause, Marcus Scheunemann, Claas-Norman Ritter, and Paul Schütte. Berlin united - naoth 2013 - the robocup team of humboldt-universität zu berlin. Proceedings of the Symposium and Team Descriptions of the 17th RoboCup, Eindhoven, June 2013.
- [5] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawa, and Hitoshi Matsubara. Robocup: A challenge problem for ai. *AI magazine*, 18(1):73, 1997.
- [6] Patrick MacAlpine and Peter Stone. Using dynamic rewards to learn a fully holonomic bipedal walk. In *AAMAS Adaptive Learning Agents (ALA) Workshop*, June 2012.
- [7] Christiaan Meijer. Getting a kick out of humanoid robotics. Master's thesis, Universiteit van Amsterdam, 2012.
- [8] Peter Nordin and Wolfgang Banzhaf. An on-line method to evolve behavior and to control a miniature robot in real time with genetic programming. *Adaptive Behavior*, 5(2):107–140, 1997.