AN ADAPTIVE APPROACH TO HUMANOID LOCOMOTION

José Pontes, Stephane Doncieux, Cristina Peixoto Santos, Vincent Padois

► To cite this version:

José Pontes, Stephane Doncieux, Cristina Peixoto Santos, Vincent Padois. AN ADAPTIVE APPROACH TO HUMANOID LOCOMOTION. *Advances in Cooperative Robotics – Proceedings of the 19th International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines*, World Scientific Publishing, pp.437-444, 2016, 978-981-3149-12-0. 10.1142/9789813149137_0052 . hal-01398513

HAL Id: hal-01398513

<https://hal.science/hal-01398513>

Submitted on 17 Nov 2016

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.

An Adaptive Approach to Humanoid Locomotion

J. PONTES^{1/†/‡}, S. DONCIEUX^{2/‡}, C. SANTOS^{3/†} and V. PADOIS^{4/‡}

[†] *CMEMs, University of Minho,
Campus de Azurém, 4800-058, Guimarães, Portugal*

[‡] *Sorbonne Universités, UPMC Univ Paris 06, CNRS
Institut des Systèmes Intelligents et de Robotique (ISIR),
75252 PARIS cedex 05, France*

¹*pontesjpc@gmail.com* ²*doncieux@isir.upmc.fr*
³*cristina@dei.uminho.pt* ⁴*padois@isir.upmc.fr*

Controllers for the locomotion of legged robots often face challenges regarding their optimization towards different objectives and different terrains. We propose an architecture that uses the information gathered in an exploration phase to adapt to a terrain with unknown characteristics. In the exploration phase virtual simulations are used to optimize the parameters of the controller in different terrains. The results of this optimization are used to identify the unknown terrain characteristics, and these values are used to select the best parameters for this particular terrain. The approach was tested in simulation, on terrains with variable friction, on an iCub robot, against a naive approach, and another where the friction was identified at random, and it clearly outperformed in both cases.

Keywords: Humanoid Locomotion, Multi-objective Optimization, Evolutionary Computation

1. Introduction

Locomotion is an essential skill for humanoids, and it has nonlinear and discontinuous dynamics. The task of walking also does not uniquely specify how the limbs must be coordinated in order to achieve the desired displacement of the robot, making it a problem with many solutions. This two factors mean that designing and tuning a controller that is “good”, even if it is not optimal, may be difficult. On top of this, the user may want the locomotion optimized towards objectives like speed or energy consumption, and the terrain in which the robot walks can present certain characteristics (*e.g.* roughness, slopes) that make it harder to walk on.¹ In this context,

the controller’s parameters that result in the desired solutions are sometimes hard or impossible to find analytically, and tuning them manually, through experimentation, results in a lengthy and cumbersome trial and error process, and sometimes in sub-optimal controllers.

We propose a method that uses information from several simulations to adapt the locomotion task of a robot on a terrain with unknown characteristics, as well as to different goals prioritized by the user (*e.g.* speed). The method is separated in two distinct phases: an exploration phase, which is longer and done in preparation to the adaptation phase, which should be quick, taking no more than a couple of minutes. In the exploration phase thousands of simulations are performed, resulting in a large amount of information that can be used in the next phase. The simulations are guided by an optimization algorithm that does not rely on an analytical analysis of the locomotion. In the adaptation phase the robot walks in an unknown terrain it has to adapt to. The behavior observed in that terrain is compared to the results from the exploration phase, and a set of parameters that results in an optimal or near-optimal behavior is selected. In the continuity of this paper this method is presented, along with some experiments that were conducted to test it.

2. Related Work

Naive, manual hand-tuning of locomotion controllers tends to lead to non-robust or sub-optimal controllers. An alternative approach is doing an analytical study of the dynamics of the robot’s locomotion.^{2,3} These studies can be very complex and computationally expensive to conduct, and there is no general framework that is applicable to every case. When an analytic approach is impractical, the information needed for the optimization of the controllers’ parameters may be gathered by experimentation that provides information for stochastic and derivative-free methods, usually requiring thousands of evaluations. Some examples of these methods employed in robotics control optimization are Genetic Algorithms (GAs),^{4,5} Covariance Matrix Adaptation (CMA),⁶ and Particle Swarm Optimization (PSO).⁷

Like the optimization towards different objectives, solutions for the adaptation required by terrain changes vary in their complexity and approach. Some examples are: a linear interpolation between control tables;² designing the control system from the start to comply with expected necessary conditions for stable dynamic walking on uneven terrain;⁸ and using a reflex mechanism as feedback that will change the step length while walking in a slope.⁹ More recently, Cully *et al.*¹⁰ introduced an intelligent trial and

error algorithm that uses the information from a map of high-performance locomotion behaviors to discover a compensatory behavior that works in spite of damage to the robot.

Because it optimizes the control parameters without sensory feedback, or internal feedback from the controller, this last approach is more general than the previous ones, allowing it to be implemented in different controllers and robots, without having to be fundamentally changed. Our proposed method is similar in these general aspects.

3. Methods

We propose a method that uses information from several simulations to adapt the locomotion task of a robot to a specific terrain. The simulations are done in virtual terrains with different values of modeled variable(s), in what is called the exploration phase.

Different sets of parameters for a given controller are a tentative solution $\mathbf{x}$ for the task of walking. The solutions are evaluated terrains with variable parameters (*e.g.* friction, slope) $\boldsymbol{\theta}$, resulting in observed behavioral features of the locomotion (*e.g.* speed, applied torque) $f(\mathbf{x}, \boldsymbol{\theta})$. In the exploration phase the control of the robot is simulated with different solutions and in different terrains, saving the results in a training data set $\mathbb{D}_t = \{\mathbf{x}, \boldsymbol{\theta}, f(\mathbf{x}, \boldsymbol{\theta})\}$. The solutions are selected by a multi-objective, evolutionary, optimization algorithm that takes into account the resulting features from previous tests, and selects new solutions in order to optimize some of them.

The adaptation phase (see Figure 1) involves identifying the variable of a terrain with unknown characteristics. From the exploration phase, various data sets are obtained $\mathbb{D}_{t_1}, \mathbb{D}_{t_2} \dots \mathbb{D}_{t_n}$, one for each n values of terrain parameters tested ($\mathbb{D}_\theta = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \dots \boldsymbol{\theta}_n\}$). From each data set a list of solutions optimal towards different locomotion optimization objectives is selected, and are then all combined into one single list $\mathbb{D}_x$. The next step is to test all the solutions from $\mathbb{D}_x$ in the different terrain parameters from $\mathbb{D}_\theta$, in order to obtain a data set that contains the optimal solutions across all terrains tested. This is defined as $\mathbb{D}_{t_{\text{opt}}} = \{\mathbb{D}_x, \mathbb{D}_\theta, f(\mathbf{x}, \boldsymbol{\theta}) \mid \forall \mathbf{x} \in \mathbb{D}_x, \forall \boldsymbol{\theta} \in \mathbb{D}_\theta\}$.

The most important step of the adaptation phase is the identification of the unknown variable of the terrain the locomotion is adapted to. To this end, a solution is picked from $\mathbb{D}_{t_{\text{opt}}}$, one that maximizes the variation in resulting locomotion features across the n values of $\boldsymbol{\theta}$ tested for previously, making it more likely to correctly identify the terrain in the next step,

Fig. 1. Architecture of the identification phase.

$$\mathbf{x}_s = \arg\max_x \{ \text{Var}(\{f(\mathbf{x}, \theta) \mid \theta \in \mathbb{D}_\theta\}) \}. \quad (1)$$

This solution, defined as $\mathbf{x}_s$, is then tested in the terrain with unknown parameter θ_{new} . The resulting features $f(\mathbf{x}_s, \theta_{new})$ are compared with the ones resulting from the tests of $\mathbf{x}_s$ with each of the values from $\mathbb{D}_\theta$, by calculating the Euclidean norm of these vectors

$$\theta_{id} = \arg\min_{\theta} \{ \|f(\mathbf{x}_s, \theta) - f(\mathbf{x}_s, \theta_{new})\| \mid \theta \in \mathbb{D}_\theta \}. \quad (2)$$

The value of the terrain parameter associated with the most similar features is selected as the estimated value for the new terrain, θ_{id} . With the value of the parameter that defines the terrain estimated, from the data set of optimal solutions $\mathbb{D}_{t_{opt}}$ a solution $\mathbf{x}_f$ is chosen, so that it is optimal in regards to the performance indicators with higher priority,

$$\mathbf{x}_f = \arg\max_x \{ f(\mathbf{x}, \theta_{id}), \forall \mathbf{x} \in \mathbb{D}_x \}. \quad (3)$$

4. Experiments and Results

The method was tested in simulations on an iCub robot,¹¹ in a simulation framework called XDE, developed by CEA-LIST.¹² The controller used for the locomotion control is one that organizes tasks in a hierarchy, with the main task for walking based on a Zero Moment Point (ZMP) predictive control.¹³ For the exploration a multi-objective Genetic Algorithm called NSGA-II was used, implemented in the EA framework SFERESv2.¹⁴ The locomotion controller was optimized towards higher speeds and a lower total of torque applied by the joints, as well as two stability measures.

The controller was optimized to walk forward for 20 seconds, by changing the values of seven parameters: length, width, and height of the steps, time taken for each step, the ratio between the swing and support phases of the steps, the time horizon of the ZMP controller, and the friction coefficient estimated for the controller. The locomotion features optimized for were the speed of the robot (maximized), the total torque applied by the joints, squared (minimized)^a, the maximum friction coefficient observed between the feet and the ground (minimized) and the average error of the ZMP trajectory tracking. The optimization was conducted for different values of floor friction (0.05, 0.10, 0.25, 0.50, 0.75, 1.00, 1.25 and 1.50), simulated using the Coulomb model. NSGA-II ran, for each friction value tested, 100 generations, for a population size of 100. Runs which ended with a controller failure, resulting in a robot fall, were considered as having fitness penalties for the distance and torque objectives. The objective of this phase was to have a set of controllers that were optimal, or near-optimal, for the four different objectives. This set is built by selecting the pareto optimal solutions, *i.e.* the set of the best trade-off solutions.

For the adaptation phase, our methods was tested for different friction values, comparing the results with a naive approach and a random one. In the naive approach, the solution that is best suited for our objectives was selected ignoring the terrain parameter. In the random approach, a value of friction coefficient was randomly selected, instead of being identified.

The three approaches were tested for different hierarchies of priorities of locomotion objectives, that define how $\mathbf{x}_f$ is chosen in the last step of the adaptation. The solutions from $\mathbb{D}_x$ are filtered at four stages, according to their performance in the terrain defined by θ_{id} . At each stage a number

^a $\int_0^{t_f} \tau^\top \tau dt$. τ is the vector of the joints's torques and t_f the final time stamp of the simulation.

of solutions are kept according to their performance regarding one of the objectives. The hierarchy is defined as an order of the four objectives that defines which performances are looked at in each stage, and a percentile assigned to each objective that defines the percentage of solutions kept in each stage. Two hierarchies were tested for: one referred to as balanced, which prioritizes first the torque with the 70th percentile, then the speed with the 70th, and then the controller’s friction estimation with the 50th; the second one prioritizes speed to the 100th percentile, which results in ignoring the other indicators.

Each approach was tested by having them adapt to the robot walking in a terrain with a given value of friction. The values of friction tested for were from 0.05 to 2.00, with intervals of 0.05 (extremes included). The results for each approach were compared with a Wilcoxon test, which can take results from two different approaches and assesses if their performances differs. It uses ranks to compare the difference between each matched pair (in our case pairs of results for the test of the same friction value in different approaches). The lowest differences having a rank of 1, and higher differences result in higher ranks and lower p-values.

Table 1 shows the Wilcoxon test applied to the speed and torque features of the locomotion. The pairs compared were the performances that two approaches presented for a given friction value. Better performance is expressed as higher sums of rank for the speed feature, and lower for the torque feature. ”Winning“ statistically significant values are in bold. A p-value lower than 0.05 means that the null hypothesis that the two approaches performed on a comparative level can be rejected. Table 2 shows information regarding the success rate (a failure is defined as when the robot falls) of each approach for each priority hierarchy, as well as the average speed and torque squared per second. The best values for each hierarchy group are in bold font.

From Table 1, and regarding speed, our approach showed higher performance for both priority hierarchies, against both the naive and the random approach. Its torque values are generally worse (higher) but the only statistically significant instances are one in which they are better and other in which they are worse. Looking at Table 2, our identification approach had a higher success rate in both priority hierarchies, as well as higher average speeds and lower average total torque squared per second.

Table 1. Wilcoxon test results for our approach (ID) against random and naive approaches, speed and torque features.

Objectives	Feature	Approach	Ranks sum ID	Ranks sum other	p-value
Balanced	Speed	Random	589	221	0.0143
Balanced	Speed	Naive	820	0	0.0000
Speed	Speed	Random	595	224	0.0123
Speed	Speed	Naive	571	243	0.0261
Balanced	Torque	Random	558	252	0.0573
Balanced	Torque	Naive	667	153	0.0006
Speed	Torque	Random	430	389	0.8125
Speed	Torque	Naive	146	668	0.0006

Table 2. Success rates and average values of performance indicators for the adaptation phase tests.

Approach	Objectives	Success %	Average speed (m/s)	Average torque squared per second (N ² .m ² /s)
ID	Balanced	95.0	0.1020	422.44
Random	Balanced	92.5	0.0936	567.28
Naive	Balanced	85.0	0.0672	453.46
ID	Speed	80.0	0.2676	691.36
Random	Speed	72.5	0.2563	716.65
Naive	Speed	55.0	0.2531	1102.63

5. Conclusions

Given these results, we conclude that the exploration, terrain parameter modeling, and identification we proposed shows success in adaptation to terrains with unknown characteristics, when compared with usual approaches akin to manual optimization. The terrain parameter has to be modeled in a specific way, and we only tested the approach for friction. In the future we plan to use other variables. The approach also lacks the capability to adapt to failed identifications, which is correlated with the fact that it does not include feedback loops, and ignores a lot of the information gathered in the exploration phase. We also plan to expand the method regarding this areas, as well as testing it in a physical robot.

Acknowledgments

This work was supported by FCT with the reference project UID/EEA/04436/2013, by FEDER funds through the COMPETE 2020 - Programa Operacional Competitividade e Internacionalizao (POCI) with the reference project POCI-01-0145-FEDER-006941, and by the Programa de Acções Universitárias Integradas Luso-francesas (PAUILF 2013/14/15). José Pontes is supported by FCT with a grant with reference UMINHO/BI/78/2012.

References

1. E. Westervelt, J. Grizzle, C. Chevallereau, J. H. Choi and B. Morris, *Feedback Control of Dynamic Bipedal Robot Locomotion.*, first edn. (CRC Press, 2007).
2. M. van de Panne, E. Fiume and Z. Vranesic, A controller for the dynamic walk of a biped across variable terrain, in *Proc. of IEEE CDC 1992*, (IEEE, 1992).
3. Y. Zheng and K. Yamane, Optimization and control of cyclic biped locomotion on a rolling ball, in *Proc. of IEEE-RAS Humanoids 2011*, (IEEE-RAS, oct 2011).
4. L. Rodrigues, M. Prado, P. Tavares, K. Da Silva and A. Rosa, Simulation and control of biped locomotion-GA optimization, in *Proc. of IEEE CEC 1996*, (IEEE, 1996).
5. C. Paul and J. Bongard, The road less travelled: morphology in the optimization of biped robot locomotion, in *IROS 2001 - IEEE/RSJ*, (IEEE, 2001).
6. J.-c. Wu and Z. Popović, Terrain-adaptive bipedal locomotion control, in *ACM SIGGRAPH 2010*, (ACM Press, jul 2010).
7. J. Kennedy and R. Eberhart, Particle swarm optimization, in *Proc. of IEEE ICNN 1995*, (IEEE, 1995).
8. H. Kimura, Y. Fukuoka and A. H. Cohen, *The International Journal of Robotics Research* **26**, 475(may 2007).
9. C. Liu, D. Wang and Q. Chen, *IEEE SMC 2013* **43**, 1206(sep 2013).
10. A. Cully, J. Clune, D. Tarapore and J.-B. Mouret, *Nature* **521**, 503(may 2015).
11. G. Sandini, G. Metta and D. Vernon, 50 Years of Artificial Intelligence, Lecture Notes in Computer Science Vol. 4850 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007)
12. X. Merlhiot, J. L. Garrec, G. Saupin and C. Andriot, The XDE mechanical kernel: Efficient and robust simulation of multibody dynamics with intermittent nonsmooth contacts., in *Proc. of the 2nd Joint International Conference on Multibody System Dynamics*, 2012.
13. J. Salini, V. Padois and P. Bidaud, Synthesis of complex humanoid whole-body behavior: A focus on sequencing and tasks transitions, in *IProc. of IEEE ICRA 2011*, (IEEE, may 2011).
14. J.-B. Mouret and S. Doncieux, {SFERES}v2: Evolvin' in the Multi-Core World, in *Proc. of IEEE CEC 2010*, 2010.