

Automated Gait Adaptation for Legged Robots

Joel D. Weingarten† Gabriel A. D. Lopes† Martin Buehler§ Richard E. Groff‡ Daniel E. Koditschek†
jweingar@umich.edu glopes@umich.edu buehler@BostonDynamics.com regroff@eecs.berkeley.edu kod@umich.edu

†Department of Electrical Engineering and Computer Science, The University of Michigan

§Boston Dynamics, Cambridge, MA

‡Department of Electrical Engineering and Computer Science, University of California, Berkeley

Abstract—Gait parameter adaptation on a physical robot is an error-prone, tedious and time-consuming process. In this paper we present a system for gait adaptation in our RHex series of hexapedal robots that renders this arduous process nearly autonomous. The robot adapts its gait parameters by recourse to a modified version of Nelder-Mead descent, while managing its self-experiments and measuring the outcome by visual servoing within a partially engineered environment. The resulting performance gains extend considerably beyond what we have managed with hand tuning. For example, the best hand tuned alternating tripod gaits never exceeded 0.8 m/s nor achieved specific resistance below 2.0. In contrast, Nelder-Mead based tuning has yielded alternating tripod gaits at 2.7 m/s (well over 5 body lengths per second) and reduced specific resistance to 0.6 while requiring little human intervention at low and moderate speeds. Comparable gains have been achieved on the much larger ruggedized version of this machine.

I. INTRODUCTION

In this paper we document the performance improvements in a hexapedal robot achieved by a nearly autonomous gait adaptation system. Appropriately designed gait variant¹ parameter optimization has improved top speed and energy efficiency by a factor of three beyond what any prior hand tuned settings could achieve. Significantly, the new parameter settings drive the robot into a qualitatively different operational regime with a pronounced aerial phase — typically more than 35% of the complete gait cycle, as documented in Fig. 1. In this regime, forward speed exceeds that of a motor's output shaft angular velocity scaled by leg length — the speed of an equivalent wheeled vehicle with the same motor gear assemblies powering wheels of the same radius. Thus, well in advance of our much desired but still imperfect analytical understanding, empirical gait adaptation in RHex begins to suggest the advantages of springy legs that can store energy at the motor's power limits and then return it far more quickly, at just the right time, and in just the right direction to produce the faster, more efficient aerial phase.

RHex (see Figure 1) is a power- and computation-autonomous hexapod robot [1]. Inspired by cockroach locomotion [2], RHex features compliant legs and a simple

¹The term "gait", seems to encompass both the discrete notion of a pattern of footfalls and the continuous notion of their relative timing and magnitude. In this paper, we will use the term "gait pattern" to denote the discrete notion. Formally, the homotopy class of a simple closed curve (the embedding of a circle) in the appropriate torus. We will use the term "gait variant" to denote the continuous notion. Formally, the particular choice of embedding within specified pattern. It is this latter aspect of gait whose adaptation we discuss in the present paper.

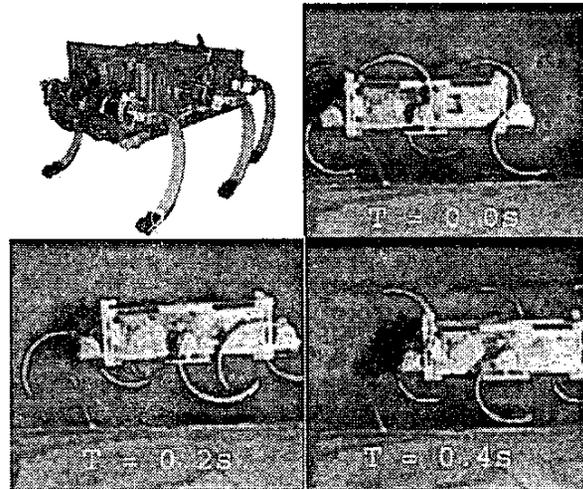


Fig. 1. RHex is a power- and computation-autonomous robotic hexapod, featuring compliant legs and a simple mechanical design. The chassis measures 48cm×22cm×12.5cm, and the distance from hip to ground in normal standing posture is 15.5cm. This time sequence from a typical stride of an efficiency optimized gait exemplifies the repeated maximal-compression, flight-apex, maximal-compression phase cycle. In this instance, liftoff occurred at approx. $t=0.16s$ and touchdown at $t=0.3s$, resulting in an aerial phase of $0.14s/0.4s = 35\%$.

mechanical design. Each leg has a single actuated mechanical degree of freedom and can rotate fully about the hip joint [3]. A growing body of evidence suggests that high speed cockroach runners employ open loop feedforward style gait control, since the lag due to neural signal propagation from brain to leg is large relative to the speed of the gait [4]. Further inspired by this principle of cockroach locomotion, the original control design for RHex employs an essentially open loop control strategy incorporating hand-tuned reference trajectories for the leg joint angles, the "clock" signal depicted in Figure 2. In section IV.A we offer a brief physical interpretation of the four gait parameters that are denoted by the knot points q_1 and q_2 in Figure 2.

More sophisticated closed loop controllers for RHex are under active development [5]. However, notwithstanding its simplicity, the open loop clock driven scheme lent RHex a degree of mobility unprecedented in autonomous legged machines at the time of its initial communication [1] and deserves study and improvement in its own right. While recent

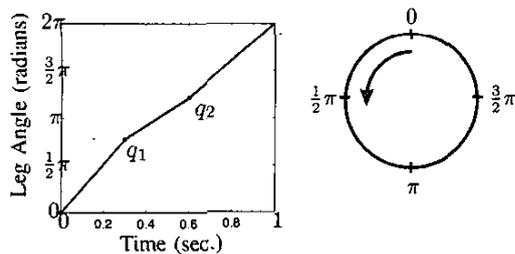


Fig. 2. *Left* - The "clock" signal that drives the legs for walking *Right* - angle naming convention

analysis of reduced degree of freedom models of this "simple" scheme has begun to reveal the underlying basis of gait stabilization [6], we are very far from understanding the factors of performance in different environments. Empirically, it is clear that for each gait pattern and fixed variant, performance varies considerably with details of the terrain type — e.g. on linoleum, concrete, pavement, gravel, grass, etc. Conversely, for a fixed surface, we have found that slight variations in the parameters that select these variants can have a significant impact on performance. Moreover, the admissible region of the parameter space is quite large. These observations motivate the central focus of this paper — the development of an automated method for tuning the gait variant parameters.

The organization of this paper is as follows: Section II describes how we implement the gait adaptation process as an offline parameter optimization problem and discusses our choice of descent algorithm. Section III presents our vision-based automation system with emphasis on the state machine that governs its autonomy. Finally, Section IV presents the results of a series of gait optimization experiments performed on the hexapedal platforms.

II. GAIT ADAPTATION

Using intuition, an experienced designer can often conceive of a gait pattern for a given task, however finding an appropriate operating point in the associated gait variant parameter space is typically less amenable to intuition. Fortunately, the designer will often have an idea in mind of the desirable performance attributes that would distinguish a better variant from a worse, and it is quite natural to encode these desired properties in the form of a scalar valued cost function. Hence, tuning can generally be reduced to an empirically formulated optimization problem. [7]. In a legged robotic system, especially one featuring compliant legs such as RHex, it is difficult to obtain accurate models of, for example, actuators, nonlinear springs and damping in the legs, varying friction coefficients, and complex ground-body interactions. The lack of a good model necessitates that all experimentation is done on the physical robot.

Learning and optimization to improve behavior has previously been successfully implemented in a growing number of robotics settings. Among the most successful, Atkeson and Schaal have used reinforcement learning (R.L.) to improve

or learn robot behaviors, for example, with a devil-sticking robot [8] while Ng has used his R.L. based PEGASUS [9] algorithm to autonomously control helicopter robots. Both approaches require an accurate model to run experiments in simulation, precluding their use in the present setting. Porta et al. [10] use reinforcement learning to generate a free gait for a simulation of the Genghis II robot and Davidor [11] among others have used genetic algorithm techniques to optimize over robot trajectories. Again, in our system, the burden of experimentation due to the absence of a viable model make these techniques ill-suited. Most recently, Kohl and Stone [12] report significant performance gains in the Aibo ERS-210A robot using a form of policy gradient R.L.

NASA's Ambler robot led by Thorpe [13] takes a deliberative planning approach to generate static gaits. This planning approach is used to determine carefully footfalls for the robot over varied terrain. I-Ming Chen et al. [14] have explored gait generation for an inchworm robot, modeling the segments of the robot as a finite state automaton and searching through the resulting state transition graph to generate gaits. It is not clear how these kinematic approaches to learning in legged locomotion might be adapted to the present dynamical setting.

Given the significant cost of experiments and measurement noise in their assessment, we have chosen the Nelder-Mead algorithm [15] [16], a derivative free simplex method for scalar function optimization, to implement gait optimization. Convergence of Nelder Mead has been established only for convex functions in one and two dimensions [17]. Even in two dimensions, convergence results are weak (the simplex volume vanishes asymptotically and vertex values converge but not necessarily over the same points) and there are established cautionary examples (i.e. cases where the algorithm converges to a non-critical point) of a seemingly benign (i.e., smooth and strictly convex) nature [18]. Nevertheless, Nelder Mead incurs in principle the least experimental cost per step of any of the other "direct search" (derivative free) methods and despite some published accounts of its breakdown in very regular application settings, it has been empirically observed to perform well on a wide range of optimization problems [19]. Note that a derivative free approach to hill climbing is desired in this setting because experimental variability makes the approximation of gradients difficult and untrustworthy.

III. SYSTEM AUTOMATION

The effort and difficulty of executing a trial and collecting the associated data required for parameter tuning make a compelling case for its automation. A single descent generally requires hours of robot time and the inevitable operator fatigue introduces errors. Over this lengthy period, additional uncontrolled variation inevitably arises through the natural aging of the physical system: changes in leg stiffness as its constituent materials degrade and varying estimates of power usage as battery levels change. Automating the descent decreases the operator induced noise, thereby avoiding unnecessary trial repetition, shortening the total length of the descent and diminishing as well these effects of natural aging.

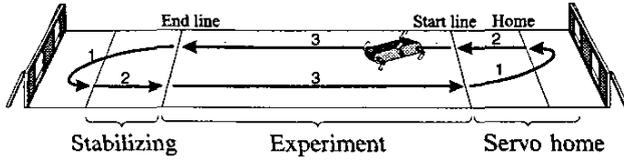


Fig. 3. Illustration of a typical setup used for automated gait adaptation. A set of 3 beacons is placed at each side of a corridor. The robot moves back and forward registering against the beacons. The lines perpendicular to the corridor represent the start location for the stabilizing phase and the start/stop for the experiment phase.

The robot receives visual data at 30hz from an onboard Sony DFW300 firewire camera. We use the visual registration algorithm described in [20] implemented by recourse to engineered beacons (bright red vertically striped panels as depicted in figure. 3. In order to better describe the implementation we distinguish three main components: the finite state-machine acting as a high level supervisor, the controllers associated with each supervisor state and finally the camera map.

A. Sequential Composition of Controllers

Transition events between discrete supervisor states occur when the robot reaches (or, via surrogate means, supervisor states "believes" itself to have reached, in the cases noted below wherein it lacks the sensory modality to measure the relevant aspects of its state directly) its goal inside the domain. These concepts may be formalized [21] as follows. Let Φ_i be a controller with domain of attraction $\mathcal{D}(\Phi_i)$ and goal $\mathcal{G}(\Phi_i)$. We say that controller Φ_i prepares controller Φ_{i+1} , denoted by $\Phi_i \succeq \Phi_{i+1}$, if the goal of the first lies in the domain of attraction of the second: $\mathcal{G}(\Phi_i) \subset \mathcal{D}(\Phi_{i+1})$. By construction the set of controllers $\mathcal{U} = \{\Phi_1, \dots, \Phi_4\}$ associated with each state represented in figure 3 induce a directed cyclic graph, ie $\Phi_i \succeq \Phi_{i+1}$ and $\Phi_4 \succeq \Phi_1$. To guarantee that the robot can handle any situation the robot's workspace \mathcal{W} should be covered by the domains of attraction of the set of controllers: $\mathcal{W} \subset \bigcup_{\Phi_i \in \mathcal{U}} \mathcal{D}(\Phi_i)$.

B. State-Machine Model

The sequential composition of the constituent continuous controllers is implemented by a supervisor defined by the standard finite state machine illustrated in figure 4. The standard "prepares" events that label the transition arrows in fig 4, as defined above, are triggered by vision, the Nelder-Mead algorithm and (in critical situations only) the user. The three primary supervisor states in an optimization trial are: servo-home, stabilize and experiment. Additional states are added to deal with undesired events. The numbered states illustrated in figure 3 are described next:

- **Servo home Domain of attraction:** Entire workspace. The controller assumes that the robot is in any upright configuration inside the optimization area (in all experiments reported here we have used a 15×2 m corridor). If no beacons appear on the robot's FOV then it rotates in place until it finds an

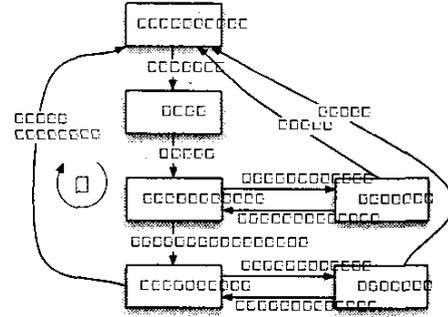


Fig. 4. Illustration of state machine used for automated gait optimization. A trial is considered successful if loop A occurs. Other loops occur if the robot loses sight of the beacons for longer than a predefined time or a critical situation occurs. Only critical situations require human intervention, in general the robot is able to recover by rotating in place until the beacons appear in the field of view of the onboard camera.

appropriate constellation of 3 beacons.²

Goal: Move the robot into a predefined home location preparing itself to start a new trial. A navigation function [22], [20] drives the robot to the home position while guaranteeing that the beacons stay in the FOV at all times.

- **Stabilizing phase Domain of attraction:** Locations in which the robot is behind the home line illustrated in figure 3 and a set of beacons is centered in the FOV.

Goal: Cross start line illustrated in figure 3. This stage of the composition is introduced to eliminate the transient response of the gait being tested. The controller used is the same as in the experiment phase described next. Since we have no sensor capable of measuring directly when the transient response has ended, the goal in this state is triggered by distance.³

- **Experiment phase Domain of attraction:** The robot must be over the start line.

Goal: Cross end line illustrated in figure 3. The experiment phase drives the robot in a straight line for a fixed length. The controller maintains a constant forward velocity and steers the robot through the corridor so that it stays on a line as much as possible. In order to eliminate disturbances introduced by the steering leg offsets a dead zone is added to the yaw controller resulting in a 90% no steering motion on slow gaits.

The recovery states illustrated in figure 4 are activated when the robot temporarily loses the beacons during a trial. Heuristically, the robot turns in the direction in which the beacons are spotted last. If the recovery does not bring the robot back into track within a couple of seconds then the trial is aborted and the robot returns home using the previously described servo

²The corridor is so engineered with beacons that for every location therein, some interval of heading angles is guaranteed to afford a clear view of an appropriate constellation. It is for this reason that the domain of attraction arising from the servo home controller includes the entire workspace.

³Thus while we adhere to the formal definition of sequential composition [21] with respect to a surrogate projection of greatly reduced dimension (a projection of the robot's three degree of freedom configuration in the horizontal plane), this is only a coarse substitute for the more refined goal that would need to be defined in the underlying state space of the robot's full 12 dimensional rigid body position and velocity.

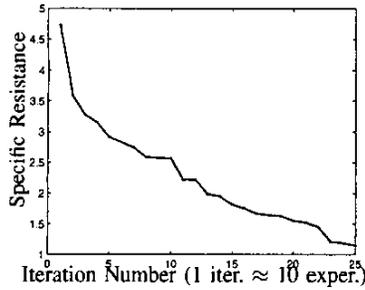


Fig. 5. A typical instance of human-driven (no vision system) Nelder-Mead optimization of specific resistance.

home controller.

IV. EXPERIMENTAL RESULTS

Experiments described in this section take the form of repeated runs over a fixed 8 m track depicted in Fig. 3. A cost function is computed from the average speed and average power recorded over each run, and the gait parameters listed below are adjusted before each subsequent run according to the Nelder-Mead variant described in section II. Two different cost functions — specific resistance (1), and speed-weighted specific resistance (2) — are used to achieve, respectively, highly efficient, and fast stable gaits. We discuss outcomes for three different physical settings. In section B we present the results of hand measured and human driven runs with RHex [1] and with Rugged RHex [23]. In Section C, we discuss a set of autonomously generated runs driven and measured by the visually servoed state machine described in Section III.

A. Gait Variant Parameter Space and Cost Functions

Our parameterization of the walking gait yielded an eight dimensional space which allows affordance over the slopes and timing of the piece-wise linear function graphed in figure 2 (i.e. moving knot points q_1 and q_2 in figure 2), the PD gains at the hip joints of the robot, a trajectory smoothing factor, and the period of the gait. In this “alternating tripod” scheme, the same reference trajectory is applied to each leg, but the signal seen by the left tripod is 180 degrees out of phase from the signal seen by the right tripod [1]. Intuitively, the reference trajectory imposes a slower rotational velocity on the legs while putatively on ground, and faster while recirculating through the air. Moving the knot points changes the timing, relative speed and the length of the two phases of the reference trajectory. At each hip, a single actuator applies torque to a leg shaft through a local PD controller that regulates the difference between the reference signal and the motor shaft angle and velocity. While the period regulates the average speed of the motors the relationship between period and forward velocity is strongly non-monotonic. Indeed, there is very little in the way of an intuitively compelling relationship between these parameters and the robot’s physical motion. Nor does our best present mathematical understanding, outlined in [6], yet provide anything close to an approximation of the mapping.

In the absence of further understanding, we resort to purely empirical tuning of intuitively prescribed fitness measures.

The cost function we use to encode efficiency is the average specific resistance [24] [1],

$$f_{sr} = P_{av}/mgv_{av}, \quad (1)$$

a dimensionless quantity which has become a standard measure of vehicle efficiency. Here, P_{av} is the average power⁴, and v_{av} is average velocity, measured over the course of the 8m run. Constants, m and g are the mass of the robot and acceleration of gravity respectively. To encode speed, the inverse of velocity was tested and rejected as a performance criterion, because it led to gaits that were fast but extremely sensitive to perturbations from the environment to the point of instability. Instead, we chose a speed weighted version of Specific Resistance which combines the desirable properties of specific resistance with the desire to find faster gaits.

$$f_v = f_{sr}/v^2 = P_{av}/mgv_{av}^3, \quad (2)$$

It is our feeling that specific resistance and stability are strongly correlated as unstable gaits tend to “waste” energy.

TABLE I
PERFORMANCE IMPROVEMENTS: HUMAN DRIVEN OPTIMIZATION

Robot	Cost Fn.	Measure	Pre-tune Measure	Post-tune Measure
RHex	f_{sr}	Spec. Res.	2.0	0.72
RHex	f_v	Spec. Res.	4.0	0.84
RHex	f_{sr}	Speed	0.5m/s	1.2m/s
RHex	f_v	Speed	1.2m/s	2.7m/s
Rugged	f_{sr}	Spec. Res.	2.2	0.80
Rugged	f_v	Spec. Res.	2.2	0.85
Rugged	f_{sr}	Speed	0.4m/s	0.9m/s
Rugged	f_v	Speed	0.4/s	1.2m/s

B. Adapting Gait Variant Parameters: Human Driven

We first tested the optimization as applied to our gait parameterization without the vision system enabled. Instead a highly experienced driver was used to run experiments. We show the results using two different hexapedal robots, RHex and Rugged RHex

1) RHex:: For each cost function we performed approximately 10 descents each typically involving 300-500 trials. Figure 5 shows how the current best specific resistance decreases over a sample descent. Table I shows that maximum velocity of the speed gait increased threefold, up to 2.7m/s, and specific resistance was lowered to 0.6. As mentioned in the introduction, with both the speed and endurance gaits RHex achieves a true aerial phase and is thus running rather than walking. Using the optimized endurance gait, RHex can travel over 3.3km on a single set of batteries, up from 750m.

⁴In this work, the total power (which includes power for the on-board computation and inefficiencies in the electronics) is used to compute specific resistance. Some other studies consider only mechanical power, which yields a lower specific resistance.

2) *Rugged*: While the parameterization of Rugged RHex has the same control architecture as does RHex, but at almost twice the mass, its higher torque and, hence, lower maximum speed motors add additional constraints to the robot's locomotion speed and efficiency. Nevertheless, applying our parameter optimization scheme to Rugged RHex yielded similar results. Table I shows nearly a factor of three improvement in both top speed and specific resistance. The similar forward velocities for each of the different cost functions can be attributed to the reduced maximum angular rate of the motor shafts.

TABLE II
ACCURACY AND RELIABILITY OF AUTOMATED SYSTEM

Test	Speed	Automated	Human	Trials(N)
Trial Timing (std. dev.)	0.5m/s	~ 0.2 sec.	~ 0.5 sec.	50
	1.0m/s	~ 0.3 sec.	~ 0.5 sec.	50
	2.0m/s	N.A.	~ 0.5 sec.	20
Steering Rate (percentage)	0.5m/s	~ 10%	~ 70%	10
	1.0m/s	~ 40%	~ 85%	10
	2.0m/s	N.A.	~ 90%	10
Run Redo Rate (percentage)	0.5m/s	~ 2%	~ 10%	50
	1.0m/s	~ 10%	~ 10%	50
	2.0m/s	100%	~ 20%	20
Human Assisted (percentage)	0.5m/s	0%	100%	50
	1.0m/s	~ 3%	100%	30
	2.0m/s	100%	100%	5

C. Autonomous Gait Variant Parameter Adaptation

In this section we present the results of experiments to tune both the speed and specific resistance of the alternating tripod gait pattern using the autonomous vision guided system introduced in Section III. Once again, we report the results of two sets of (roughly 300 - 500) 8m runs on linoleum, although similar results were obtained operating outdoors on concrete.

1) *Level of Automation*: Judging the efficacy of any automation system entails an assessment of the extent to which it reduces the need for human intervention. While the state machine in our system is formally complete in the sense that its constituent basins cover the entire set of legal configurations in the horizontal plane of the robot's rigid body placements (in other words, every contingency is in principle accounted for), this is a mere projection of the robot's true physical state (at the very least, at 48 dimensional quantity [1]) and there are a number of situations where human intervention is still necessary. In particular the automated system is presently unable to recover when the robot has flipped on it's back, nor is it equipped with thermal sensors permitting the detection of motor temperatures near or at the point of incurring motor damage. For these reasons, we never run the automated system without a human assistant to watch the robot's progress and resolve collisions with these unmodeled and fatal obstacles. Thus, while not entirely displaced, the burden on the human operator is substantially reduced, allowing useful attention to other work while tuning progresses, thereby allowing for longer and more accurate tuning sessions.

2) *Accuracy and Reliability of the Automated System*: Besides making it significantly easier on the operator the

automation system is both more reliable and accurate than the human operated version at speeds less than 1.3m/s. To test the attributes of the vision system we ran trials at three constant speeds over our 8m linoleum course. Table II shows how the vision system achieved more than a factor of 2 reduction in timing variance while significantly reducing the percentage of the run where steering inputs are used to keep the robot on course. Furthermore, the percentage of experiments that need to be re-evaluated (redo rate = successfully completed runs/total runs) is greatly reduced (with the vision system on, re-evaluation is triggered when the beacons are lost or the robot flips. In the human operated case these can be attributed to operator error or flipping). At lowest speeds (approx. 0.5m/s) our vision system proved to work entirely without human assistance as opposed to every run without vision. As the velocity of the robot increases, it becomes more prone to flipping and thus the experimenter had to intervene to right the robot. As can be seen the automation system fails

TABLE III
PERFORMANCE IMPROVEMENT: INEXPERIENCED HUMAN VS.
AUTOMATED SYSTEM

Cost fn.	Measure	Pre-tune	Human	Automated
f_{sr}	Spec. Res.	1.8	1.1	0.63
f_v	Spec. Res.	1.8	1.3	0.8
f_{sr}	Speed	0.4m/s	0.9m/s	1.0m/s
f_v	Speed	0.4 m/s	1.4m/s	1.3m/s

at high speed. We attribute this failure to the low frame rate returned by our vision system and image blur due to a long exposure time. Currently we have a dedicated vision processor on RHex equivalent to a Pentium II 300Mhz which yields 15 frames/sec when running our vision algorithms. We feel the a faster frame rate coupled with a smaller exposure time will allow our system to be successful at the higher speed.

3) *Tuning Walking with Vision*: Table III shows the results of tuning using the vision system. To give a sense of how the difficulty of driving a hexapedal robot affects the results of the optimization we have also compared the vision system to results obtained by an inexperienced driver. In both cases the initial conditions were chosen via the same method and several descents performed. The automated system matched the inexperienced human's final speed and trounced him with respect to efficiency. As can be seen from the second line of Table I, the automated system beat even the experienced driver (over 500 hours driving time) in the final efficiency of its speed targeted optimization by about 10%. descents performed. We attribute this improved performance to the increased steering and timing ability of the automated system documented in Table II. In contrast, neither the inexperienced driver nor the automated system were able to operate at the high speeds of the experienced human. It is quite difficult for a human to get a feel for this and we have already remarked upon the limitations of the vision system that the automated tuner relies upon.

D. Discussion

Although we have presented evidence of effective adaptation only over a simple test course on level ground, we have in fact successfully tuned up RHex's gait over many different surfaces and terrains (hard packed dirt, grass, concrete, small rock beds etc.) using this framework. The resulting open loop controller consistently exhibits a rapid return to its steady state gait pattern even in the presence of significant ground perturbations during runs with long aerial phases (albeit the specific resistance may no longer be as favorable).

While hard to characterize in quantitative terms, these fitness landscapes seem to be extremely complex over the available gait variant parameter space. For example, tuning outcome seems to depend heavily upon the "quality" of the initial condition given to Nelder-Mead. In our experience, a "good" initial condition entails the generation of an initial simplex whose vertices (i.e., nine different parameter settings) are not too close to each other and also yield reasonably "high quality" gaits. Absent the intuition of an experienced experimenter to create "good" initial conditions both the inexperienced human and the automated system were typically led to "dead-ends" — conditions of continued anomalous measurement that ended a run, or simply poor quality local minima. Nonetheless, as our tables show, the successful descents, from good initial conditions, yielded significant performance increases.

V. CONCLUSION

In this paper we have presented a system for automated gait tuning that achieves dramatic improvements in performance by recourse to an appropriate descent technique, the Nelder-Mead direct search algorithm, applied to an intuitively designed cost function. We have used a family of existing visual servo algorithms together with carefully engineered beacons to allow the robot to implement the tuning procedure in a nearly autonomous manner.

We have applied this gait adaptation system to two members of the RHex family of hexapods with performance improvements by a factor of three in specific resistance and by a factor of five beyond what had earlier been documented as the speed record for autonomous legged robots. Testing with the vision system showed increased accuracy and improved steering over human-controlled experiments. Moreover, these experiments demonstrated the system's ability to run without human intervention at low and moderate speeds. The speed- and endurance-optimized gaits yield strongly stable dynamical running with aerial phases of up to 35% of the gait cycle.

Further work will be required to improve the robustness of the vision automation system at high speeds and to improve the error handling to remove the need for a human observer. Work now in progress focuses on adding feedback to our gait controllers to improve performance over less favorable terrain and to adapt to changing terrains - both in an online manner.

ACKNOWLEDGMENTS

This work could not have been accomplished without the prior RHex development efforts of Uluc Saranli, Eric Klavins,

Katherine Scott and Haldun Komsuoglu. We are grateful for many helpful discussions with Satinder Baveja. The work was supported in part by DARPA/ONR Grant N00014-98-0747.⁵

REFERENCES

- [1] U. Saranli, M. Buehler, and D.E. Koditschek, "Rhex: A simple and highly mobile hexapod robot," *The International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.
- [2] H. Komsuoglu D.E. Koditschek H.B. Brown Jr. M. Buehler N. Moore D. McMordie R. Altendorfer, U. Saranli and R. Full, "Rhex: A biologically inspired hexapod runner," *Aut. Rob.*, vol. 11, 2001.
- [3] M. Buehler, U. Saranli, and D.E. Koditschek, "Single actuator per leg robotic hexapod," McGill University, The Regents of the University of Michigan: USA, 2002, US Patent 6,481,513.
- [4] D.L. Jindrich and R.J. Full, "Dynamic stabilization of rapid hexapedal locomotion," *J Exp Biol*, vol. 205, pp. 2803–23, 2002.
- [5] U. Saranli and D.E. Koditschek, "Template based control of hexapedal running," in *International Conference in Robotics and Automation*, Taipei, Taiwan, 2003, IEEE.
- [6] R. Altendorfer, D.E. Koditschek, and P. Holmes, "Towards a factored analysis of legged locomotion models," in *International Conference in Robotics and Automation*, Taipei, Taiwan, 2003, IEEE.
- [7] U. Saranli, "Simsect hybrid dynamical simulation environment," Tech. Rep. CSE-TR-436-00, The University of Michigan Department of Computer Science Technical Report, 2000.
- [8] S. Schaal and C.G. Atkeson, "Robot juggling: An implementation of memory-based learning," *Con. Sys. Mag.*, vol. 14, no. 1.
- [9] Andrew Y. Ng and Michael Jordan, "Pegasus: A policy search method for large mdps and pomdps," in *Uncertainty in Artificial Intelligence, Proceedings of the Sixteenth Conference*, 2000.
- [10] Josep M Porta and Enric Celaya, "Efficient gait generation using reinforcement learning," in *CLAWAR*, Germany, 2001.
- [11] Yuval Davidor, *Genetic Algorithms and Robotics: A Heuristic Strategy for Optimization*, World Scientific, 1991.
- [12] Nate Kohl and Peter Stone, "Policy gradient reinforcement learning for fast quadrupedal locomotion," in *Submitted to: Int. Conf. Robotics and Automation*, IEEE, 2004.
- [13] David Watergreen and Chuck Thorpe, "Developing planning and reactive control for a hexapod robot," in *Int. Conf. Rob. and Aut.* IEEE, 1996.
- [14] I-Ming Chen, Song Hual Yeo, and Yan Gao, "Locomotive gait generation for inchwork-like robots using finite state approach," *Robotica*, vol. 19, pp. 535–542, 2001.
- [15] J. A. Nelder and R. Mead, "A simplex method for function minimization," *Computer Journal*, vol. 7, pp. 308–313, 1965.
- [16] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge (UK) and N.Y., 1st edition, 1986.
- [17] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the nelder-mead simplex algorithm in low dimensions," *SIAM Journal on Optimization*, vol. 9, pp. 112–147, 1998.
- [18] K. I. M. McKinnon, "Convergence of the nelder mead simplex method to a nonstationary point," *SIAM Journal on Optimization*, vol. 9, pp. 148–158, 1998.
- [19] R. Lewis, V. Torczón, and M. Trosset, "Direct search methods: Then and now," *Journal of Computational and Applied Mathematics*, vol. 124, pp. 191–207, 2000.
- [20] Gabriel A. D. Lopes and Daniel E. Koditschek, "Visual registration and navigation using planar features," in *International Conference in Robotics and Automation*, Taipei, Taiwan, 2003, IEEE.
- [21] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *The International Journal of Robotics Research*, vol. 18, no. 6, pp. 534–555, June 1999.
- [22] Elon Rimon and D. E. Koditschek, "Exact robot navigation using artificial potential fields," *IEEE Trans. on Rob. and Aut.*, vol. 8, no. 5, pp. 501–518, Oct 1992.
- [23] Martin Buehler, "www.bostondynamics.com," 2003.
- [24] G. Gabrielli and T. H. von Karman, "What price speed?," *Mech. Eng.*, vol. 72, no. 10, 1950.

⁵This work was performed while Buehler was an Associate Professor at McGill, before moving to Boston Dynamics in September 2003.