Efficient bipedal robots based on passive-dynamic walkers

Steve Collins,[1] Andy Ruina,[2*] Russ Tedrake,[3] Martijn Wisse[4]

[1]Mechanical Engineering, University of Michigan, Ann Arbor, MI 48104, USA
[2]Theoretical & Applied Mechanics, Cornell University, Ithaca, NY 14853, USA
[3]Brain & Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[4]Mechanical Engineering, Delft University of Technology, NL-2628 CD Delft, NETHERLANDS

[*]To whom correspondence should be addressed; E-mail: ruina@cornell.edu.

November 19, 2004

Passive-dynamic walkers are simple mechanical devices, comprised of solid parts connected with joints, that can walk stably down a slope. They have no motors or controllers, yet can have remarkably human-like motions. Here we present three new robots which extend passive-dynamic walking principles to walking on level ground by using an active power source. These robots use less control and less energy than other powered robots, yet walk more naturally. These results highlight the importance of the coupling between form and function in human and animal locomotion.

Most researchers study human locomotion by observing people as they walk, measuring joint angles and ground reaction forces (*1*). Our approach is different - we study human locomotion by designing and testing walking machines which we compare to humans in terms of morphology, gait appearance, energetic efficiency, and control.

Previous bipedal robots (*2–4*) with human-like forms have demonstrated smooth versatile motions. These impressive robots are based on the mainstream control paradigm, namely precise joint-angle control. For the study of human walking, this control paradigm is unsatisfactory because it requires actuators with higher precision and frequency response than human muscles have (*5*) and requires an order of magnitude more energy. To address these issues, *passive-dynamic walkers* (*6*) (Fig. 1) were proposed as a new design and control paradigm. In contrast to mainstream robots which actively control every joint angle at all times, passive-dynamic walkers do not control *any* joint angle at *any* time. Although these walkers are made only of solid parts connected with joints and have no control or actuation, they can walk downhill with startlingly human-like gaits (*7*).

Our research demonstrates that the human-like properties of passive-dynamic machines are not restricted to their gravitational power source, but rather that these principles remain important in level-ground walking. We built three powered walking robots (Fig. 2) at three different institutions by substituting gravitational power with simple actuation, and each walks stably on a level surface.

The Cornell biped (Fig. 2a) is based on the passive device in Fig. 1d, and is electrically powered at ankle push-off. It has five internal degrees of freedom (two ankles, two knees, and a hip), the arms are mechanically linked to the opposite leg, and the small body is kinematically constrained so that its midline bisects the hip angle. The Delft biped (Fig. 2b) has a similar morphology, but is powered by pneumatic hip actuation, and has a passive ankle. The MIT learning biped (Fig. 2c) is based on the simpler ramp-walkers in Fig. 1a-b with a passive hip, is powered by two servo motors in each ankle, and uses reinforcement learning to automatically acquire the controller. It has six internal degrees of freedom (two in each ankle and two hips), the arms are mechanically linked to the opposite leg, and the body hangs passively. The supplementary videos show these robots walking and the supplementary text describes their construction details.

The Cornell biped is specifically designed for minimal energy use. For walking at constant speeds, the primary energy losses are due to dissipation when a foot hits the ground and to active braking by the actuators (negative work). The Cornell design demonstrates that it is possible to completely avoid this negative actuator work. The only work done by the actuators is positive: the left ankle actively extends when triggered by the right foot hitting the ground, and *vice versa*. The hip joint is not powered and the knee joints only have latches. The average mechanical power (*8*) of the two ankle joints is about 3 watts, almost identical to the scaled gravitational power consumed by the passive-dynamic robot on which it is based (*7*). Including electronics, micro-controller, and actuators, the Cornell biped consumes about 11 watts total (*9*).

To compare energy usage between humans and robots of different sizes, it is convenient to use the dimensionless *specific cost of transport*, c_t = (energy used)/(weight$\times$distance traveled). In order to isolate the effectiveness of the mechanical design and control system from the actuator efficiency, we distinguish between the specific *energetic* cost of transport, c_{et}, and the specific *mechanical* cost of transport, c_{mt}. Whereas c_{et} uses the total energy consumed by the system (11 watts for the Cornell biped), c_{mt} only considers the positive mechanical work of the actuators (3 watts for the Cornell biped). The 13 kg Cornell biped walking at 0.4 m/s has $c_{et} \approx 0.2$ and $c_{mt} \approx 0.055$. Humans are similarly energy effective in walking with $c_{et} \approx 0.2$, as estimated by the oxygen they consume (V02), and $c_{mt} \approx 0.04$ (*10–12*). Measurement of actuator work on the Delft-biped yields $c_{mt} \approx 0.08$. Based on the small slopes that it descends when passive, we estimate the MIT biped to have $c_{mt} \approx 0.02$. Although the MIT and Delft bipeds here were not specifically designed for low energy use, both inherit energetic features from the passive-dynamic walkers upon which they are based. By contrast, we estimate the state-of-the-art Honda humanoid Asimo to have $c_{et} \approx 3.2$ and $c_{mt} \approx 1.6$ (*13*). Thus Asimo, perhaps representative of joint-angle controlled robots, uses about 20 times the energy (scaled) of a typical human.

Controllers for state-of-the-art, level-ground walking robots are typically complex, requiring substantial real-time computation. In contrast, the Delft and Cornell bipeds walk with primitive controllers. Their only sensors detect ground contact and their only motor commands are on/off signals issued once per step. In addition to powering the motion, hip actuation in the Delft biped also improves fore-aft stability by swiftly placing the swing leg in front of the robot before it falls forward (*14, 15*).

The MIT biped (Fig. 2c) is designed to test the utility of motor learning on a passive-dynamic mechanical design. The goal of the learning is to find a controller which stabilizes the robot's trajectory on level terrain, using the passive ramp-walking trajectory as the target. The robot acquires a feedback control policy which maps sensors to actions using a function approximator with 35 parameters. With every step that the robot takes, it makes small, random changes to the parameters, and correlates those changes to changes in walking performance to improve the stability of the step-to-step dynamics (Fig. 3). Using an actor-critic reinforcement learning algorithm (*16*), this correlation is estimated efficiently on the real robot despite sensor noise, imperfect actuators, and uncertainty in the environment. The robot's actuators are positioned so that when they are commanded to their zero position the robot imitates its passive counterpart. Starting from this zero-controller, the learning system quickly and reliably acquires an effective controller for walking, using only data taken from the actual robot (no simulations), typically converging in ten minutes or approximately 600 steps. Figure 3 illustrates that the learned controller not only achieves the desired trajectory, but is also robust to disturbances. The robot can start, stop, steer, and walk forward and backward at a small range of speeds. This optimized learning system works quickly enough that the robot is able to continually adapt to the terrain as it walks.

Each of the robots presented in this paper has some design features that are intended to mimic humans. The Cornell and Delft bipeds utilize anthropomorphic geometry and mass distribution in their legs and demonstrate ankle push-off and powered leg swinging, both present in human walking (*12*). They do not use high-power nor high-frequency actuation, which are unavailable to humans. These robots walk with human-like efficiency and human-like motions (Fig. 4 and supplementary videos). The motor learning system on the MIT biped is also inspired by biology, both in the mechanism of the learning rule (*17*), and in the formulation of motor learning as an optimal feedback control problem (*18*).

The Cornell and Delft bipeds demonstrate that walking can be accomplished with extremely simple control. These robots do not rely upon sophisticated real-time calculations nor on significant sensory feedback such as from continuous sensing of torques, angles, or attitudes. This suggests that steady-state human walking might require only simple control, as well. The sequencing of human joint-angles in time could determined as much by morphology as by motor control. Note that no other robots have done particularly better at generating human-like gaits even when using high-performance motors, a plethora of sensors, and sophisticated control. The conclusion that natural dynamics may largely govern locomotion patterns was already suggested by passive-dynamic machines. A common misconception has been that gravity power is essential to passive-dynamic walking, making it irrelevant to understanding human walking. These machines demonstrate that there is nothing special about gravity as a power source. We achieve equally successful walking using gravity power, ankle power, or hip power.

The learning results of the MIT biped also suggest that the mechanical design of walking robots, or of the human musculo-skeletal system, can have a major impact on the efficiency of motor learning. Previous attempts at learning control for bipedal robots have required a huge number of learning trials in simulation (*19*), or an initial hand-designed controller on the robot (*20*). By exploiting the natural stability of walking trajectories on the passive-dynamic walker, our robot was able to learn in just a few minutes without requiring any initial control knowledge.

All three of the robots presented here suggest the value of using passive dynamics to understand human walking. We expect that robot designs will improve dramatically in the coming years, becoming significantly more versatile and more efficient. This might come to pass through further improving control of passive-based robots, or it may come through paying closer attention to energy efficiency in joint-controlled robots (*21*). Whatever the future of humanoid robots, the success of human mimicry demonstrated here, without using joint-angle control, strongly suggests an intimate relationship between body architecture and control in human walking.

References and Notes

1. D. A. Winter (Univ. of Waterloo Press, 1987).

2. Y. Sakagami, *et al.*, *Proc. Int. Conf. Intelligent Robots and Systems* (2002), pp. 2478–2483.

3. E. R. Westervelt, J. W. Grizzle, D. E. Koditschek, *IEEE Trans. Automatic Control* **48**, 42 (2003).

4. F. Pfeiffer, K. Löffler, M. Gienger, *Proc., IEEE Int. Conf. on Robotics and Automation* (Washington DC, 2002), pp. 3129 – 3135.

5. F. Zajac, *Crit Rev Biomed Eng. 1989* **17**, 359 (1989).

6. T. McGeer, *Int. J. Robotics Res.* **9**, 62 (1990).

7. S. H. Collins, M. Wisse, A. Ruina, *Int. J. Robotics Res.* **20**, 607 (2001).

8. Mechanical power is defined here as net positive mechanical work at the joints= $\int_0^T \sum [\omega_i M_i]^+ \, dt/T$ where T is the period of one step, ω_i is the relative angular veloc-

ity at one joint, M_i is the torque across that joint, $[x]^+ = x$ if $x > 0$ and 0 otherwise, and the sum is over all the joints. Because only the ankle does positive work on the Cornell biped, this can be measured by measuring the foot force as the ankle extends in a synthetic push-off.

9. For the Cornell biped total power was measured by averaging the voltage across a 1Ω resistor put in series with the battery.

10. E. Atzler, R. Herbst, *Pflg. Arch. ges. Physiol.* **215**, 291 (1927).

11. N. H. Molen, R. H. Rozendal, W. Boon, *Proc. Kon. Ned. Akad* **75**, 305 (1972).

12. J. M. Donelan, R. Kram, A. D. Kuo, *J. Exp. Biol.* **205**, 3717 (2002).

13. Honda's ASIMO can walk at a variety of speeds, kick balls, and even climb stairs. It weighs 510 N, can walk at speeds up to 1.6 km/hr, and drains a 38.4V, 10 amp-hour battery in about 30 minutes (http://world.honda.com/ASIMO/). Using these numbers, we estimate $c_{et} \approx 3.2$ and, assuming a 50% drive train efficiency, $c_{mt} \approx 1.6$.

14. M. Wisse, J. van Frankenhuyzen, *Proc. Conf. Adaptive Motion of Animals and Machines* (AMAM, Kyoto, Japan, 2003).

15. M. Wisse, A. L. Schwab, R. Q. van der Linde, F. C. T. van der Helm, *Submitted to IEEE Trans. on Robotics and Automation* (2004).

16. R. Tedrake, T. W. Zhang, M. Fong, H. S. Seung, *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems* (IEEE/RSJ, Sendai, Japan, 2004).

17. H. Seung, *Neuron* **40**, 1063 (2003).

18. E. Todorov, *Nature Neuroscience* **5**, 1226 (2002).

19. H. Benbrahim, J. A. Franklin, *Robotics and Autonomous Systems* **22**, 283 (1997).

20. W. T. Miller III, *IEEE Control Systems Mag.* **14**, 41 (1994).

21. J. Pratt, Exploiting inherent robustness and natural dynamics in the control of bipedal walking robots, Ph.D. thesis, MIT (2000).

22. J. E. Wilson, Walking toy, *Patent*, U.S. Patent Office (1938). Patent # 2,140,275 available at http://www.tam.cornell.edu/˜ruina/hplab/.

23. R. Tedrake, T. W. Zhang, M. Fong, H. S. Seung, *Proc. IEEE Int. Conf. on Robotics and Automation* (IEEE, New Orleans, LA, 2004).

24. M. Garcia, A. Chatterjee, A. Ruina, *Dynamics and Stability of Systems* **15**, 75 (2000).

25. G. T. Fallis, Walking toy, *Patent*, U.S. Patent Office (1888). Available at http://www.tam.cornell.edu/˜ruina/hplab.

26. A. D. Kuo, *J. Biomech. Eng.* **124**, 113 (2002).

27. A. D. Kuo, *Int. J. of Robotics Res.* **18**, 917 (1999).

28. M. Hubbard, *J. of Applied Mechanics* **46**, 931 (1979).

29. **Acknowledgments.** The Cornell robot was developed by S.C. with suggestions from A.R.; the Delft robot was developed by M.W. and Jan van Frankenhuyzen on an STW grant, with help from Arend Schwab; the MIT robot was developed by R.T. and Teresa Weirui Zhang with help from Ming-fai Fong and Derrick Tan in the lab of H. Sebastian Seung. A.R. & S.C. were funded by an NSF Biomechanics grant. R.T. was funded by the Packard Foundation and the NSF. The text was improved by comments from Naveen Agnihotri, Chris Atkeson, Joseph Burns, Anindya Chatterjee, Michael Coleman, Phil Holmes, Art Kuo, Yonatan

Loewenstein, Saskya van Nouhuys, Dan Paluska, Andrew Richardson, Sebastian Seung, Manoj Srinivasan, Steve Strogatz, and Nyk Sydor.

Supporting Online Material

www.sciencemag.org

Materials and Methods

Movies S1, S2, S3

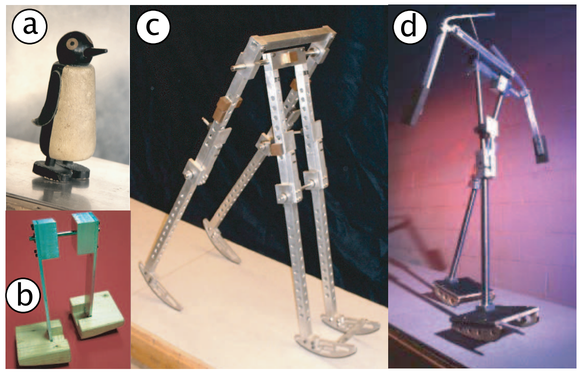

Figure 1: ‘Ramp-walking’, ‘downhill’, ‘unpowered’ or ‘passive-dynamic’ robots. Our powered bipeds are based on these passive designs. (a) The Wilson “Walkie” (*22*). (b) Our improved version (*23*). Both (a) and (b) walk down a slight ramp with the “comical, awkward, waddling gait of the penguin” (*22*). (b) Cornell copy (*24*) of McGeer’s capstone design (*6*). This four-legged ‘biped’ has two pairs of legs, an inner and outer pair, to prevent falling sideways. (c) The Cornell passive biped with arms (photo by Hank Morgan). This walker has knees and arms and is perhaps the most human-like passive-dynamic walker to date (*7*).

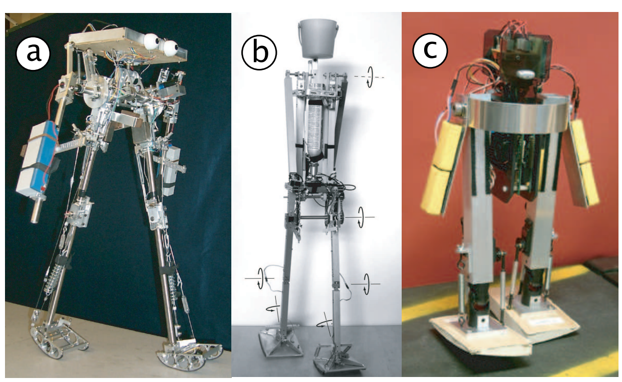

Figure 2: Three level-ground powered walking robots based on the ramp-walking designs of Fig. 1. (a) The Cornell biped. (b) The Delft biped. (c) The MIT learning biped. These powered robots have motions close to their ramp-walking counterparts as seen in the supplementary movies. Information on their construction is in the supplementary text.

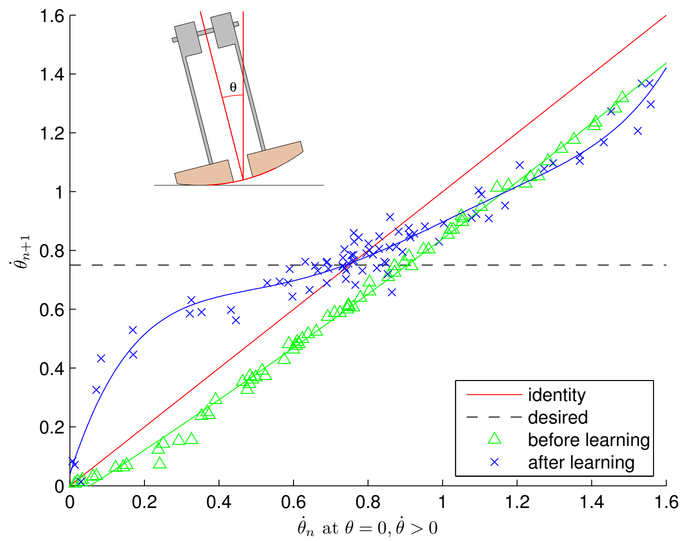

Figure 3: Step-to-step dynamics of the MIT biped walking in place on a level surface, before (triangles) and after (crosses) learning. Shown is the roll angular velocity when the right foot collides with the ground ($\theta = 0, \dot{\theta} > 0$) at step $n + 1$ versus step n. Intersections of the plots with the solid identity line are fixed points. The horizontal dashed line is the theoretical ideal, the robot would reach $\dot{\theta} = 0.75\, s^{-1}$ in one step. This ideal cannot be achieved due to limitations in the controllability of the actuation system. On a level surface, before learning the robot loses energy on every step ($\dot{\theta}_{n+1} < \dot{\theta}_n$), eventually coming to rest at $\dot{\theta} = 0$. After learning the robot quickly converges near $\dot{\theta} = 0.75\, s^{-1}$ for $0 \leq \dot{\theta}_0 \leq 1.7\, s^{-1}$.

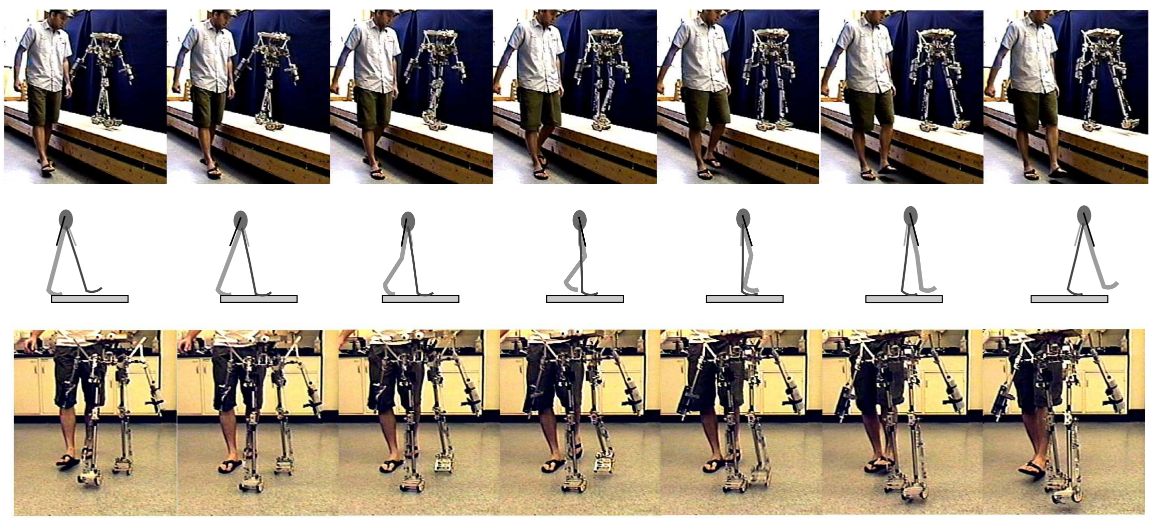

Figure 4: Two sets of video stills of the Cornell ankle-powered biped walking on a level surface next to a person. A little less than one step is shown at 7.5 frames/s. Both the robot and the person are walking at about 1 step/s. The stick figure indicates the leg angles for the corresponding video stills; the right arm and leg are darker than the left.

Supporting Online Material

This supporting online material for the paper "Efficient bipedal robots based on passive-dynamic walkers" by S. Collins, A. Ruina, R. Tedrake, and M. Wisse is also available at `http://tam.cornell.edu/~ruina/powerwalk.html` . We hope that the reviewers will watch the videos on that site.

There is one video for each of the three robots described in this paper:

- **S1** - Cornell powered biped movie. This movie shows videos of the robot walking on flat terrain, and a slow-motion segment which illustrates the ankle push-off actuation.
- **S2** - Delft powered biped movie. This movie shows the robot walking down a hall with views from the front, side, and back.
- **S3** - MIT learning biped movie. This movie begins with the powered robot imitating passive walking down a ramp. Then it shows the robot learning to walk on flat terrain with foam protective pads. After the robot learns to walk in place, its stability was tested by applying disturbances. We show the robot walking down the hall, on tiles, and outside.

Videos of the passive-dynamic walkers that inspired these designs are available on the authors' websites, and are linked off the URL above.

Materials and Methods

Cornell powered biped (Fig. 2a). This fully autonomous robot consists of two 0.8 m long legs attached at a hip joint, each having knees, curved-bottom feet, arms, and a small torso which is kept upright by connection to the legs with an angle-dividing mechanism. Each arm carries a battery. The right arm is rigidly attached to the left leg and *vice versa*, reducing yaw oscillations (*7,25*). The machine weighs 12.7 kg and has 5 internal degrees of freedom (one hip, two knees, and two ankles). The thigh to shank length and mass ratios are 0.91 to 1 and 3.3 to 1, respectively, which is important to the passive dynamics of the system. The hip joint is fully passive. A latch at each knee passively locks the shank to be parallel with its proximal thigh throughout stance. This latch is released by a solenoid at the completion of ankle push-off, at which point the knee is fully passive until knee-strike. Ankle extension is fully controlled such that the push-off restores energy lost to collisions. To minimize the needed motor size, energy for ankle push-off is stored in a compression spring between steps.

Electronics are located in the hip/torso/head visible in Fig. 2a. A finite-state machine with eight binary inputs and outputs is implemented in 68 lines of code on an Atmel AT90S8515 chip running on an ATSTK500 standard development board. A second board with relays and passive conditioning components connects the board to the electromechanical and sensory parts. During the first state, Left Leg Swing, all actuators are unpowered and the left knee latch passively locks at knee strike. When switches below the left foot detect impending heel strike, the state switches to Right Ankle Push-Off. This begins a timed activation of the solenoids that release the plantar-flexor spring of the right foot. When switches detect full foot extension, the state switches to Right Toe Return. During this state, a 9.5 Watt, 6.4 oz MicroMo® motor is activated, slowly retracting the foot and restoring spring energy. A timed activation of the solenoids simultaneously unlocks the right knee. When a switch on the motor indicates full foot retraction, the state switches to Right Leg Swing, and the motor is deactivated. The machine then swaps left and right legs and goes to the initial state. Taking all sensing, including the sensing of internal degrees of freedom which could in principle be made open loop, about 20 bits of information per step flows to the processor. Environmental sensing, i.e., the instant of foot contact, is about one third of that.

This machine is designed for minimal energy use and has only one capability: walking forward. Its speed, path and joint motions are not shaped or controlled but follow from its mechanical design and primitive ankle push-off actuation. Ankle extension occurs mostly after the opposite leg has completed heel-strike collision, so in principle the machine could be made to consume about four times less energy by having ankle push-off before, rather than after, the opposing leg's foot-to-ground collision (*26*).

The Cornell powered biped first walked successfully in August 2003.

Delft powered biped (Fig. 2b). The powered robot weighs 8 kg with 5 internal degrees of freedom (one hip, two knees, two skateboard-like ankles), and is 1.5 m tall. It is also entirely autonomous. The robot consists of two legs, each with knees and ankles, and an upper body. The knees have mechanical stops to avoid hyper-extension, and are locked with a controllable latch. Two antagonist pairs of air-actuated artificial muscles (McKibben muscles) provide a torque across the hip joint to power the walking motion. The muscles are fed with CO2 from a 58 atm cannister, pressure-reduced in two steps to 6 atm through locally developed miniature pneumatics. Low-power, two-state valves from SMC Pneumatics® connect the "muscles" either to the 6 atm supply pressure or to 0 atm.

McKibben muscles have a low stiffness when unactuated, leaving the joints to behave almost passively at zero pressure. At higher pressures, the McKibben muscles behave as progressively stiffer springs. By activating opposing muscles in different proportions, the relaxed angle of a joint can be controlled. This is applied at the hip where the artificial muscles alternate in action. At the start of each step, determined by a foot switch, one muscle is set to 6 atm and the other to 0 atm. The swing leg is thus accelerated forward until the relaxed angle of the hip is reached, where it (approximately) stays due to damping in the muscles and in the joint. If sufficient hip joint stiffness is obtained from the hip muscles, stable walking similar to that of McGeer's four-legged machine can be obtained. The upper body is kept upright not by active control, but via a mechanism at the hip which confines the upper body to the bisection angle of the two legs.

Lateral stability in two-legged robots can be obtained in a number of ways (*27*), and one solution was tested in the Delft robot. The feet are attached to the lower leg via special ankle joints which have a joint axis that runs from above the heel down through the middle of the foot, quite unlike the human ankle but much like skateboard trucks. The mechanism creates a nonholonomic constraint, enabling non-dissipative stability, as found in skateboards (*28*). If the robot starts to lean sideways as a result of a disturbance, the ankle does allow the foot to remain flat on the floor. Due to the tilted joint orientation, the leaning is accompanied by steering. So, if the walker has sufficient forward velocity, it will not fall sideways but instead change its heading.

A Universal Processor Board from Multi Motions® (based on the Microchip® PIC16F877 micro-controller) uses foot contact-switch signals to open or close the pneumatic valves. The control program is a state machine with two states: either the left or the right leg is in swing phase. At the beginning of the swing phase, the swing knee is bent. Four hundred milliseconds after the start of the swing phase, the knee latch is closed, waiting for the lower leg to reach full extension through its passive swing motion. Programmed in assembly, this amounts to about 30 lines of code. The only sensing is the time of foot contact, used once per step. Taking account of the implicit rounding from the processor loop time, we estimate the sensor information flow rate is about six bits per second.

The Delft powered biped first walked successfully in July 2004.

MIT learning biped (Fig. 2c). First we duplicated the Wilson design (Fig. 1a) using two rigid bodies connected by a simple hinge. The kneeless morphology was chosen to reduce the number of joints and actuators on the robot, minimizing the combinatorial explosion of states and control strategies that the learning algorithm needed to consider. The gait was iteratively improved in simulation by changing the foot shape for a given leg length, hip width, and mass distribution. The resulting ramp-walker (Fig. 1b) walks smoothly down a variety of slopes. The powered version (Fig. 2c) uses tilt sensors, rate gyros, and potentiometers at each joint to sense the configuration, and servo motors to actuate the ankles. The completed robot weighs 2.75 kg, is 43cm tall, and has 6 internal degrees of freedom (each leg has one at the hip and two at the ankle). Before adding power or control, we verified that this robot could walk stably downhill with the ankle joints locked.

The robot's control code runs at 200Hz on an embedded PC-104 Linux computer. The robot runs autonomously; the computer and motors are powered by lithium-polymer battery packs, and communication is provided by wireless ethernet. The learning controller, represented using a linear combination of local nonlinear basis functions, takes the body angle and angular velocity as inputs and generates target angles for the ankle servo motors as outputs. The learning cost function quadratically penalizes deviation from the desired state on the return map of the system, taken around the point where the robot transfers support from the left foot to the right foot. Before learning, outputs were zero everywhere regardless of the inputs, and the robot was able to walk stably down a ramp; because it lacks actuation, it would run out of energy when walking on a level surface. The robot kicks itself into a random starting position using a hand-designed control script to initialize the learning trials. The learning algorithm quickly and reliably finds a controller to stabilize this gait on level terrain. The resulting controller outputs ankle commands that are a simple, time-independent function of the state of the robot, and does not require any dynamic models, nor high-precision, high-frequency actuation. All learning trials were carried out on the physical biped with no offline simulations. The learned controller is quantifiably (using the eigenvalues of the return map) more stable than any controller we were able to design by hand, and recovers from most perturbations in as little as one step. The robot continually learns and adapts to the terrain as it walks.

The version of the MIT powered biped shown here first walked successfully in January 2004. The earliest powered prototype of this type at MIT first walked successfully in June 2003.