Predicting the metabolic energy costs of bipedalism using evolutionary robotics

W. I. Sellers[1,*], L. A. Dennis[2] and R. H. Crompton[3]

[1]*Department of Human Sciences, Loughborough University, Loughborough LE11 3TU, UK*, [2]*School of Computer Science and Information Technology, University of Nottingham, Jubilee Campus, Nottingham, NG8 1BB, UK* and [3]*Department of Human Anatomy and Cell Biology, University of Liverpool, The Sherrington Buildings, Ashton Street, Liverpool L69 3GE, UK*

*Author for correspondence (e-mail: W.I.Sellers@lboro.ac.uk)

Summary

To understand the evolution of bipedalism among the hominoids in an ecological context we need to be able to estimate the energetic cost of locomotion in fossil forms. Ideally such an estimate would be based entirely on morphology since, except for the rare instances where footprints are preserved, this is the only primary source of evidence available. In this paper we use evolutionary robotics techniques (genetic algorithms, pattern generators and mechanical modeling) to produce a biomimetic simulation of bipedalism based on human body dimensions. The mechanical simulation is a seven-segment, two-dimensional model with motive force provided by tension generators representing the major muscle groups acting around the lower-limb joints. Metabolic energy costs are calculated from the muscle model, and bipedal gait is generated using a finite-state pattern generator whose parameters are produced using a genetic algorithm with locomotor economy (maximum distance for a fixed energy cost) as the fitness criterion. The model is validated by comparing the values it generates with those for modern humans. The result (maximum efficiency of $200\,\mathrm{J\,m^{-1}}$) is within 15% of the experimentally derived value, which is very encouraging and suggests that this is a useful analytic technique for investigating the locomotor behaviour of fossil forms. Initial work suggests that in the future this technique could be used to estimate other locomotor parameters such as top speed. In addition, the animations produced by this technique are qualitatively very convincing, which suggests that this may also be a useful technique for visualizing bipedal locomotion.

Movies available on line.

Key words: bipedalism, biomechanics, locomotion, evolutionary computing, human.

Introduction

Humans are unique among living primates in that they are habitually bipedal. Other great apes walk bipedally on occasion (Alexander, 1991) but more often practice pronograde (African apes; Doran, 1996) or orthograde (orang-utan; Crompton et al., in press) forms of 'quadrupedal' locomotion). A switch from facultative bipedalism to habitual bipedalism thus appears to have occurred either very rapidly and during the last 6 million years (the common ancestor of humans and chimpanzees probably lived over 6 million years ago; Senut et al., 2001), or as part of a much more gradual process over the last 23 million years since the first appearance of the hominoids at the beginning of the Miocene (Martin, 1990), with bipedalism becoming progressively more favoured by natural selection. Many theories have been postulated to explain the adaptive pressures that led to the adoption of habitual bipedalism such as tool-use, display, load carrying, foraging and thermoregulation (for a review, see Fleagle, 1999), but to evaluate any of these theories in an ecological context we need to know as much as possible about the locomotor ecology of early hominids. Theories of foraging and resource exploitation in particular need to estimate the metabolic cost of locomotion as an element of energy budgets, while to evaluate hypotheses based on predator/prey interactions we also need to know the likely top speed and daily ranging capabilities of the fossil species. These values are not easy to obtain, even for modern species, but are impossible to obtain directly for fossil forms.

Traditionally the locomotor behaviour of fossil hominids has been reconstructed by detailed investigation of the morphology of elements of the post-cranial skeleton (see, for example, Rose, 1993). The dimensions of various anatomical features are compared with those of living primates, and when a match is found it is assumed that the locomotor capabilities of the fossil would have been similar to those of the modern species. Many of the fossils of interest, however, combine anatomical features that are not found together in any single living species, and alternative reconstructions of locomotor behaviour based on different features are often contradictory.

It is therefore preferable to assess locomotor behaviour on

whole-body or body-system attributes. The most predictively reliable measure in living primates is the intermembral index: the ratio of forelimb to hindlimb length. The IMI tends to be <1.0 in leapers, approximately 1.0 for quadrupeds, and >1.0 for arm-swinging primates (Napier and Napier, 1985), while the low value in humans appears to be associated with hand-carrying of loads (Wang et al., 2003). This sort of analysis works well for identifying animals that are highly specialised for a particular sort of locomotion (e.g. leapers or brachiators) but does not permit such reliable evaluation of frequency of a particular form of locomotion within the total repertoire of animals that are locomotor generalists (Oxnard et al., 1989).

More recently various forms of biomechanical analysis have been used to investigate locomotion where direct approaches are impossible, for example the top speed of dinosaurs (Alexander, 1989), the leaping capabilities of lorises (Sellers, 1996), and the likely gait of the Miocene hominoid *Proconsul hesoloni* (Li et al., 2002). Some of these approaches have concentrated on the underlying mechanics of bipedalism (Alexander, 1992a; Minetti and Alexander, 1997), whereas others have attempted to simulate the gait of extinct species based on body proportions (Crompton et al., 1998; Kramer, 1999). These latter models have used kinematics derived from motion-capture to drive the model's movements. Other models have used central control systems to drive the simulation (Yamazaki et al., 1996; Ogihara and Yamazaki, 2001), which has the advantage of allowing the model much greater freedom in the movements generated, and this is the approach adopted in this paper.

One major difficulty of all modelling approaches is the conversion of the mechanical work calculated from these models to metabolic energy consumption, which is the ecologically more important parameter. Fortunately Minetti and Alexander (1997) derived a formula from empirical data of Ma and Zahalak (1991) that allows this conversion.

Biological processes are often used as inspiration for developing computation techniques, and evolution is no exception. It has been observed that evolutionary processes have led to highly optimised solutions to complex problems, and since the 1950s evolutionary theory has been used by computer scientists as a source of inspiration for optimisation and machine learning algorithms. Evolutionary strategies (Rechenberg, 1965) are the most developed of these early techniques. They encode the problem under consideration as a sequence of real numbers and then randomly mutate these numbers. Each time a mutation is created it is compared with the previous solution and, if it performs better according to some metric, it replaces the original solution, otherwise it is discarded. Using a Gaussian mutation (adding a random value selected from a Gaussian probability distribution with a mean of zero) it is possible to generate any sequence, although the new sequences are more likely to be similar to the previous sequence than otherwise. This approach takes advantage of the fact that sequences close to the 'best' solution are likely to be similar to that solution.

Genetic algorithms were invented by John Holland in the 1960s (Holland, 1975). The genetic algorithm uses a population of solutions. Members of this population (called chromosomes) are allowed to contribute to the next generation by 'crossover', whereby two chromosomes exchange subsequences to create two new chromosomes. The selection of parent chromosomes is done randomly but is influenced by their fitness, which is calculated in some way by a fitness function. The genetic algorithm *sensu strictu* uses a fixed length sequence of bits (zeros or ones) as its chromosome. However, the genetic algorithm has had such a large impact on the field of evolutionary computation that concepts such as populations and crossover have been incorporated into other techniques and the term is now used to cover almost any population-based evolutionary search technique (for a more thorough introduction to genetic algorithms, see Davis, 1991). As our work uses a sequence of real numbers and our primary method for creating the next generation is Gaussian mutation, our technique should probably be considered an 'evolutionary strategy' approach, but we will adopt the more commonly used term, genetic algorithm (GA). These techniques have been widely used for difficult computational search problems and have become popular in the field of evolutionary robotics, where they are used to find the parameters that specify the control systems that drive the robots (Nolfi and Floreano, 2000). It is important to note that we are not modelling the evolutionary process itself. Rather, we are modelling the biomechanics of a pre-existing evolutionary sequence and using a search technique inspired by evolutionary theory to generate the control system needed to drive our biomechanical models.

The process of implementing a generic GA is shown in Fig. 1. The first step is to decide how the problem under investigation can be encoded onto a chromosome. In modeling situations the problem can usually be resolved into a number of floating point values that can be stored as a fixed-length list. In other circumstances this list may need to be of variable length, or a branching tree-like data structure may be more appropriate. The second step is to generate a starting population from a collection of chromosomes. This can be the same individual duplicated a number of times, a number of randomly generated individuals, or a population from a previous attempt. In the third step the individual chromosomes in the population are evaluated to see how well they are able to solve the problem. Once again the designer has to decide how the solution can be evaluated, and in general this is done by generating a numerical score for the individual. The higher the score, the better the solution, and the fitter the genome. In step four, individuals are chosen to reproduce. Their chance of reproduction is governed by their fitness score so that higher fitness individuals are likely to produce more offspring. In step five, reproduction occurs, and is achieved by taking copies of the parental chromosomes and altering them so that they are not identical to their parents. This alteration is often due to a combination of random changes applied to the values in the chromosome (mutation) and merging values from two parent chromosomes (crossover). The result of steps four and five
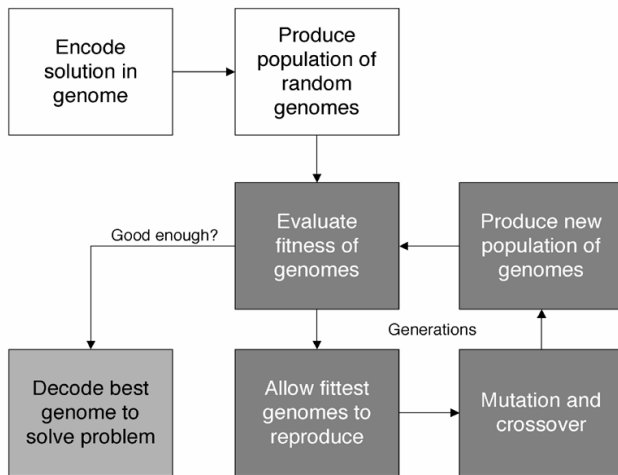
What is a genetic algorithm?

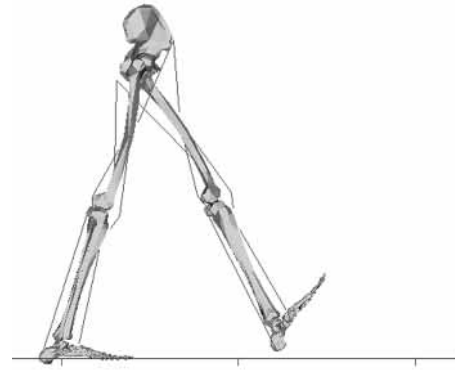Fig. 1. Diagram illustrating the process of encoding and solving a problem using a genetic algorithm.

Fig. 2. Diagram illustrating the 7-segment model used (Delp et al., 1990). The segments are: 1, head, arms and torso (HAT), origin at the centre of the hip joint; 2, 3 right and left upper leg, extending from the centre of the hip joint to the centre of the knee joint; 4, 5 right and left lower leg, extending from the centre of the knee joint to the centre of the tibio-talar joint; 6, 7 right and left foot, extending from the tibio-talar joint to the head of the 1st metatarsal.

is a new population where, hopefully, the values in the chromosomes that led to good fitness results have been preserved and the values that led to poor results have been lost. In addition there is also a chance that mutation has produced better values than in the parent generation and that crossover has led to the aggregation of high-fitness values. This new population is now returned to step three, and steps three, four and five are repeated a large number of times. Eventually the result is considered good enough, either by evaluating against some goal criteria or simply by limiting the number of generations, the process is terminated and the best individual chromosome is decoded to provide the solution to the problem.

Materials and methods

Multi-body dynamic analysis is a technique where the body is treated as a set of rigid segments acted upon by various forces and constrained by the joints between the segments. The movement of the segments can then be calculated by solving the equations of motion – the so-called forward solution (for a review, see Winter, 1990). Setting up these equations by hand is an error-prone process but there are a number of computer packages that can be used (for a list of a variety of packages, see the International Society of Biomechanics web site at http://www.isbweb.org). These packages vary in the facilities they offer, their ease of use, their flexibility and their prices. We chose Dynamechs (McMillan, 1994; McMillan et al., 1995; http://dynamechs.sourceforge.net), which is an open source C++ library that allows extensive customization, so that we could develop a simulator that would work well with our in-house genetic algorithm optimization package.

We produced a mechanical simulation of human walking. As always, there is a trade-off between biofidelity and computational cost, so a fairly simple 7-segment model consisting of 3 segment legs and a torso was used (Fig. 2). As a first approximation,

human bipedalism can be considered as an essentially two-dimensional movement, most work being done in the parasagittal plane (see, for example, Winter, 1990), so the model was constrained to move in the vertical plane by defining the lower limb joints as parallel hinge joints. The bone outlines and muscle attachment points were based on data from Delp et al. (1990) made available on the ISB website (http://www.isbweb.org), and segment inertial properties were obtained from Zatiorski and Seluyanov (1983) (see Tables 1 and 2).

Muscles were modeled as multipoint tension elements and were implemented as custom-written extensions to Dynamechs. The muscle line of action runs from its origin to insertion through a number of intermediate points that guide the muscle's path. For example, in the case of the knee an intermediate point is used to model the movement of the patella. The positions of these points can be recalculated at each step in the simulation to allow them to move, and the total length and its rate of change are stored so they can be used for force and energy calculations. The force on the end-points is simply the tension in the muscle, directed towards the neighboring intermediate points. The force on the intermediate points is the vector sum of the upstream and downstream tension force vectors (Fig. 3).

Table 1. *Joint parameters used in the model**

Joint	X (m)	Y (m)	Extension (degrees)	Flexion (degrees)
Hip	0	0	10	90
Knee	0.396	0	0	90
Ankle	0.826	0	30	60

*The model was taken from (Delp et al., 1990).

Note values of extension and flexion are relative to the human anatomical position.

X and Y are the world coordinates of the joint centres.

Table 2. *Mass distribution properties used in model**

Segment	Local origin	Mass (kg)	Ixx (kg m^2)	Iyy (kg m^2)	Izz (kg m^2)	CM X offset (m)	CM Y offset (m)
Torso	Hip joint	39.00	0.34	6.00	5.84	–0.328	0
Thigh	Hip joint	10.36	0.041	0.54	0.54	0.180	0
Calf	Knee joint	3.16	0.06	0.13	0.13	0.174	0
Foot	Ankle joint	1.00	0.018	0.001	0.018	0	0.116

*The model was taken from (Zatiorski and Seluyanov, 1983).

Ixx, *Iyy* and *Izz* are the principle moments of inertia of the segment about the local origin. CM *X* and CM *Y* are the local coordinates of the segment's centre of mass.

The force vectors $\mathbf{F}_1$ to $\mathbf{F}_n$ generated by a muscle with tension T at the points with position vectors $\mathbf{P}_1$ to $\mathbf{P}_2$ are given by:

$$\mathbf{F}_1 = T(\mathbf{P}_2 - \mathbf{P}_1)/|\mathbf{P}_2 - \mathbf{P}_1|,$$
$$\mathbf{F}_n = T(\mathbf{P}_{n-1} - \mathbf{P}_n)/|\mathbf{P}_{n-1} - \mathbf{P}_n|, \qquad (1)$$
$$\mathbf{F}_i = [T(\mathbf{P}_{i-1} - \mathbf{P}_i)/|\mathbf{P}_{i-1} - \mathbf{P}_i|] + [T(\mathbf{P}_{i+1} - \mathbf{P}_i)/|\mathbf{P}_{i+1} - \mathbf{P}_i|],$$

where $(\mathbf{P}_2 - \mathbf{P}_1)$ represents the vector from $\mathbf{P}_1$ to $\mathbf{P}_2$, $|\mathbf{P}_2 - \mathbf{P}_1|$ represents the magnitude of this vector, and i is the contact point number.

The controller produces a fractional muscle activation from 0 to 1. We converted this to tension using the theoretical model of Minetti and Alexander (1997), which is based on empirical data from Ma and Zahalak (1991). The model formulae are written in terms of torques and joint angular velocities, but it can be shown by rearrangement that it is the relative moment (current torque/maximum torque) that is dependent on the relative angular velocity (current angular velocity/maximum angular velocity); the power also depends on the relative angular velocity. This means that these angular terms can be replaced by their equivalent linear terms, relative tensions and relative rates of shortening. Therefore for a muscle with a maximum tension of T_0 (in N), maximum rate of shortening of v_{max} (in m s^{-1}), current rate of shortening of v (in m s^{-1}), and activation of α, the force is given by:

$$v \leq 0 \quad T = \alpha T_0 \left[1.8 - 0.8 \left(\frac{v_{max} + v}{v_{max} - 7.56v/k} \right) \right],$$

$$0 < v < v_{max} \quad T = \alpha T_0 \left(\frac{v_{max} - v}{v_{max} + v/k} \right). \qquad (2)$$

k is defined below.

The model also allows the metabolic power of each muscle, P (in W), to be calculated (using the same symbols as in Equation 2):

$$P = \alpha T_0 v_{max} \Phi(v/v_{max}), \qquad (3)$$

where

$$\Phi(v/v_{max}) =$$

$$\frac{0.054 + 0.506(v/v_{max}) + 2.46(v/v_{max})^2}{1 - 1.13(v/v_{max}) + 12.8(v/v_{max})^2 - 1.64(v/v_{max})^3}. \qquad (4)$$

The value of k used by Minetti and Alexander is 0.17. v_{max} and T_0 varied from muscle group to muscle group. We used a
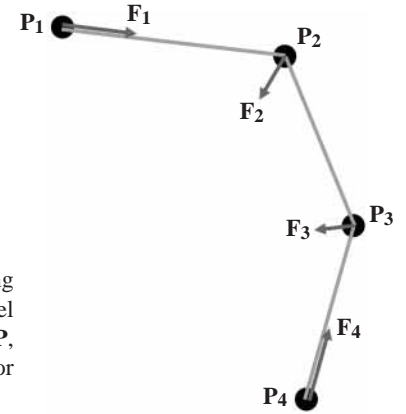

Fig. 3. Diagram illustrating the generic muscle model used. **F**, force vector; **P**, position vector. See text for explanation.

value of v_{max} equal to 6 times the resting length of the muscle, as recommended by Winter (1990), and T_0 was calculated from the physiological cross-sectional areas (PCAs) given in Pierrynowski (1995), multiplied by the intrinsic muscular strength value of 20 000 N m^{-2} given in Winter (1990). The muscles were aggregated as a set of flexors and extensors around the three joints in each limb and the values used for total PCA and mean resting length are given in Table 3. Since the model calculates the metabolic power of each muscle we can sum these to find the total metabolic power of the simulation without requiring any internal or external work calculations.

To produce the walking pattern we used a finite state machine with five states. Each state contained seven floating point values between –1 and +1: the first value was the duration of the state in seconds (the absolute value was used); the subsequent six values represented the activation level of the muscle groups (the value was assigned to the flexors if positive and to the extensors if negative). The first two states were used to initialize the walking pattern since the model started from an upright, stationary position. The first state could be equated to lifting the leading foot off the ground by extending the hip and flexing the knee, and the second state would equate to the hip being further extended with the knee extending to place the lead foot on the ground anterior to the torso. The subsequent three states were then cycled with the values applied to the left-hand side swapped with those on the right-hand side after each cycle. These states would equate to the foot being pushed against the ground by flexing the ankle to toe-off; the forward swing with the hip extending and the knee flexing; and the
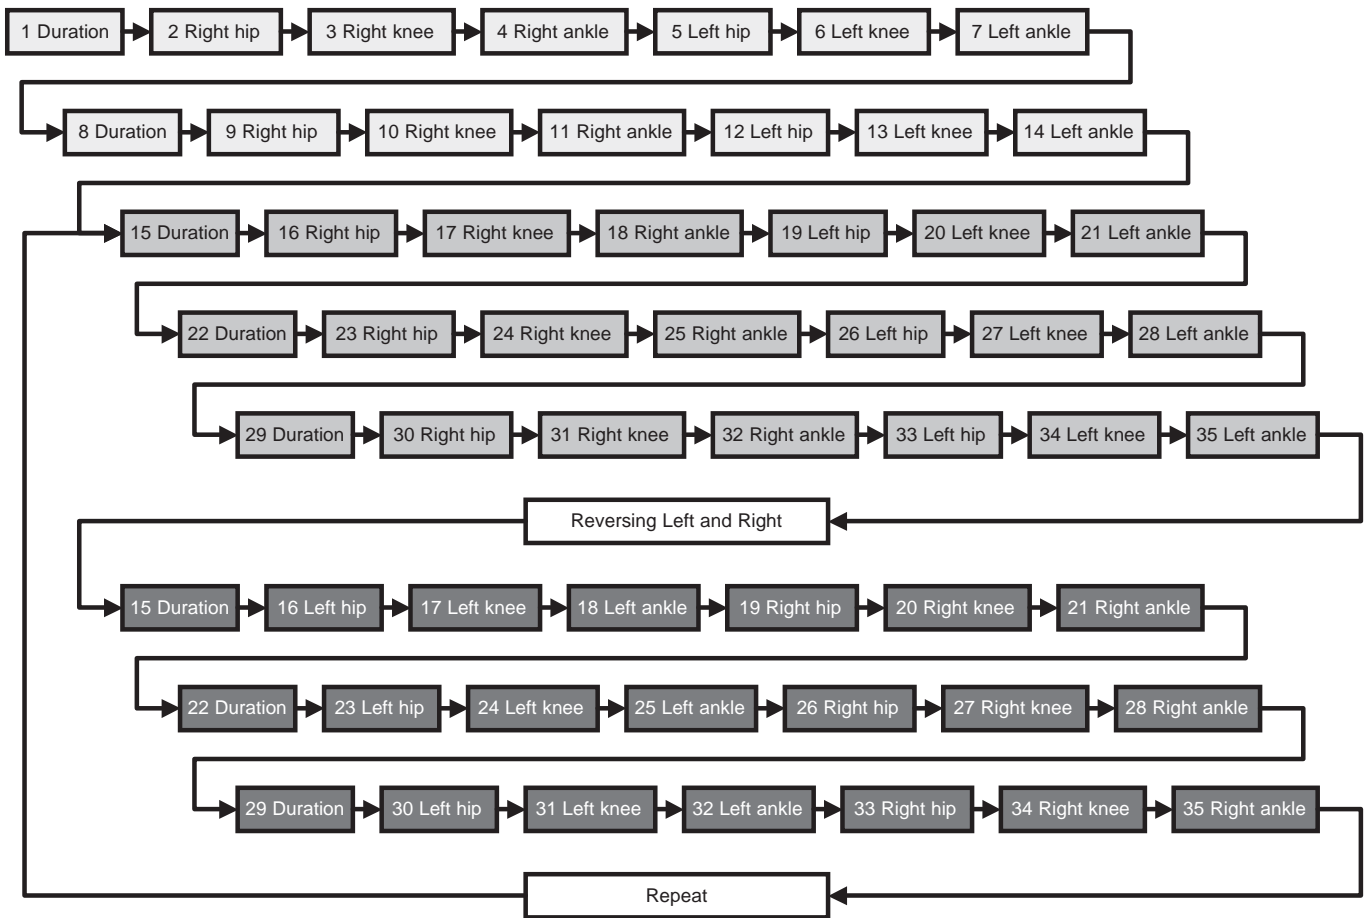
Fig. 4. Diagram illustrating the encoding used for the genome. Each value in the list of 35 values was a floating point number between –1 and +1.

Table 3. *Physiological cross-sectional areas and resting length values used in the model*

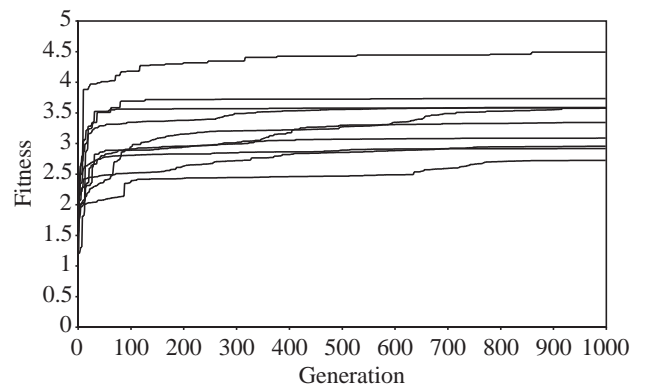| Muscle group	PCA (m^2)	Resting length (m)
Hip extensors	0.014	0.26
Hip flexors	0.007	0.24
Knee extensors	0.024	0.23
Knee flexors	0.012	0.23
Ankle extensors	0.005	0.29
Ankle flexors	0.026	0.41

PCA, physiological cross-sectional area.

Fig. 5. Graph showing the fitness of the best genome when starting with a randomly generated population of 100 individuals and a metabolic energy cutoff of 5000 J.

forward swing with the knee and hip both extending on to heel-strike. Thus the walking pattern was controlled by a list of 35 values, which was used directly as the genome for the genetic algorithm. The encoding is shown diagrammatically in Fig. 4.

Our intention was that the genetic algorithm process would optimise the muscle activation values used in the chromosome. In particular we hoped to investigate metabolic energy, so we wanted the fitness score to favour the most efficient set of muscle activation values. The fitness score for the genome was calculated by running the Dynamechs simulation until either the

model fell over or until a fixed amount of metabolic energy was used. The score was the distance that the torso had moved forward before the model terminated. In any given experiment a population of genomes was optimized using a steady state genetic algorithm, where the population for the next generation consisted of a combination of the best unique parents and the best offspring. Individuals were chosen to reproduce using a
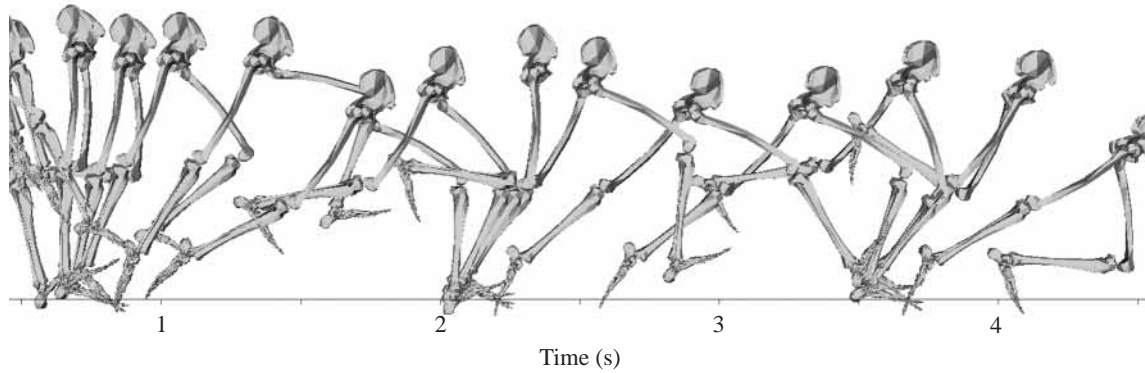
Fig. 6. Overlay image of the animation generated by the best genome when starting with a randomly generated population (see movie1.mov). The images are 0.2 s apart.
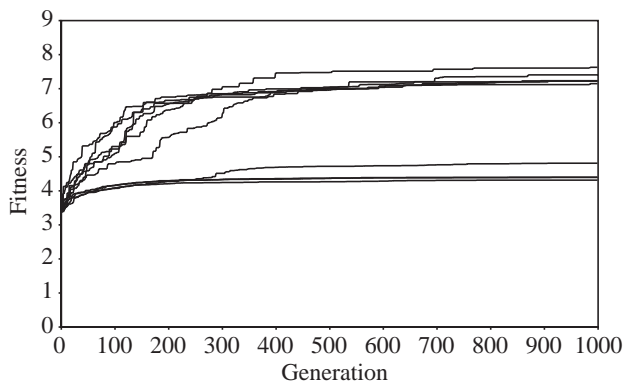

Fig. 7. Graph showing the fitness of the best genome when starting with a screened population of 100 individuals and a metabolic energy cutoff of 5000 J.

roulette-wheel selection process, where the chance of being selected was proportional to an individual's fitness score. Single-point crossover was used and mutation was effected by adding a Gaussian-distributed random value at random locations along the genome. Considerable experimentation was required to choose the various parameters and in the end a population size of 50 or 100 was used with half of the population being retained between generations. The chance of mutation was 0.05 per allele and the crossover chance was 0.01. The standard deviation (s.d.) of the Gaussian noise value was 0.1 with a mean of zero. The use of a relatively low rate of crossover agrees with the findings of evolutionary roboticists who have used similar techniques to generate gait controllers (Nolfi and Floreano, 2000). Experimentation also showed that the rate of improvement after 1000 generations tended to be relatively small, and since this generally represented a week of computer time it was used as a stop point.

Results

The first experiment used a population of 100, a randomly generated starting population, and a metabolic energy limit of 5000 J. The best fitness obtained in 1000 generations repeated over 10 trials, was a distance of 4.5 m, as shown in Fig. 5. However, inspection of the animations produced showed that stable walking had not been achieved, although it is quite possible that it would have been if more generations had been allowed (Fig. 6). The difficulty seems to be that the search space is very sparse – most randomly generated genomes have
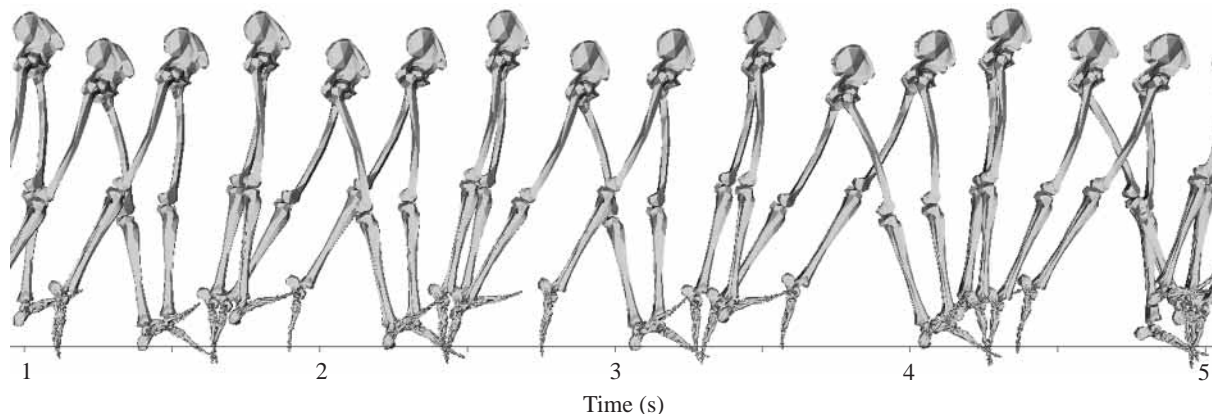

Fig. 8. Overlay image of the animation generated by the best genome when starting with a screened random population (see movie2.mov). The images are 0.2 s apart.
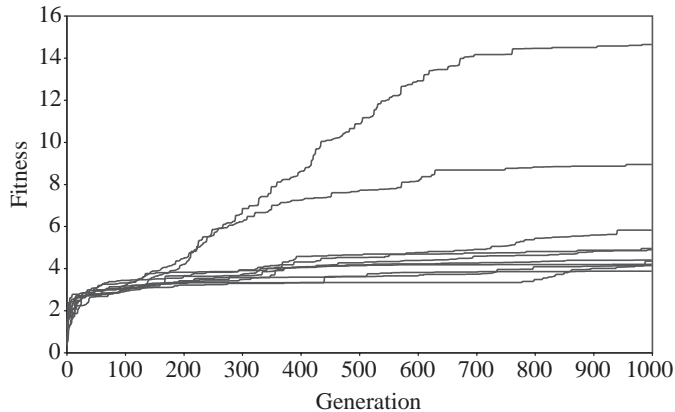
Fig. 9. Graph showing the fitness of the best genome when starting with a population of 50 individuals, derived from a hand-crafted genome and a metabolic energy cutoff of 3000 J.

a very low fitness and random mutation is unlikely to be able to transform the genome sufficiently to move it into a higher fitness region of the search space in a reasonable time.

To overcome this we ran a screening program where 3 000 000 randomly generated genomes were assessed for fitness and the top 100 were chosen as the starting population. Screening alone produced a maximum fitness of 3.4 m. This population was then evolved exactly as before. The best fitness obtained is 7.6 m in this case and the results of all trials are shown in Fig. 7.

In this case stable walking was obtained although the gait was novel: it was digitigrade with the knee kept fully extended and the ankle joints alternately flexed and extended to allow the forward swinging limb to clear the ground. This can be seen in the overlay image (Fig. 8).

The screened population did produce stable walking but it was clear from the animations that it was not taking advantage of the startup phase, so was constrained to produce a walking pattern that was able to cope with both startup and continuous locomotion. In an attempt to overcome this, we hand-crafted a genome that took advantage of the startup phase and was able to produce a single step before falling over. This involved manually setting the values in the genome and watching the

animation produced. Thus for phase 1, the left hip extensor and the left knee flexor were activated and the duration set so the model raised its left leg off the ground. For phase 2, the left knee extensor was activated and the duration set so that by the end of phase 2 the left heel had just touched the ground. For phase 3, the right ankle was strongly flexed to push the right leg off the ground and the duration set so that the right ankle had just reached full extension by the end of the phase. For phase 4, the right hip extensor and right knee flexor were activated to further raise the right leg and swing it forward with the duration set so that the foot had reached the midpoint of the forward swing by the end of the phase. For phase 5, the right knee extensor was activated to further swing forward the right leg and the duration set so that by the end of this phase the right heel had just touched the ground with the knee fully extended. Finding these values by hand was a very slow and laborious process of trial and error. The starting population was then created by incrementally mutating this genome. The original hand-crafted genome had a low fitness but we hoped it would direct the search to a profitable area of the search space (a process known as 'shaping'). This approach proved to be the most successful by far and to reduce the running time the metabolic energy cutoff was changed to 3000 J and the population reduced to 50 individuals; and even so, the best fitness produced was 14.7 m. The results for all ten trials are shown in Fig. 9.

The animation produced with this genome is also impressive (see Fig. 10). The kinematics approximate compass gait, an extreme form of 'stiff' walking (Cavagna et al., 1977; Alexander and Jayes, 1978). The compass gait is in part due to the model's two-dimensionality: in human walking the pelvis is tilted sideways during the support part of the gait cycle (this action is emphasised during speed walking) to reduce the maximum vertical movement of the centre of gravity. This obviously cannot occur in our simulation. In normal human comfortable walking there are also transitory knee flexions at heelstrike, midstance and toe-off, which have similar effects. Together these contribute to a certain amount of 'springiness' in the support limb that smooths walking and reduces the peak vertical ground reaction forces. (This normal gait should not be confused with compliant walking, where knee flexions are
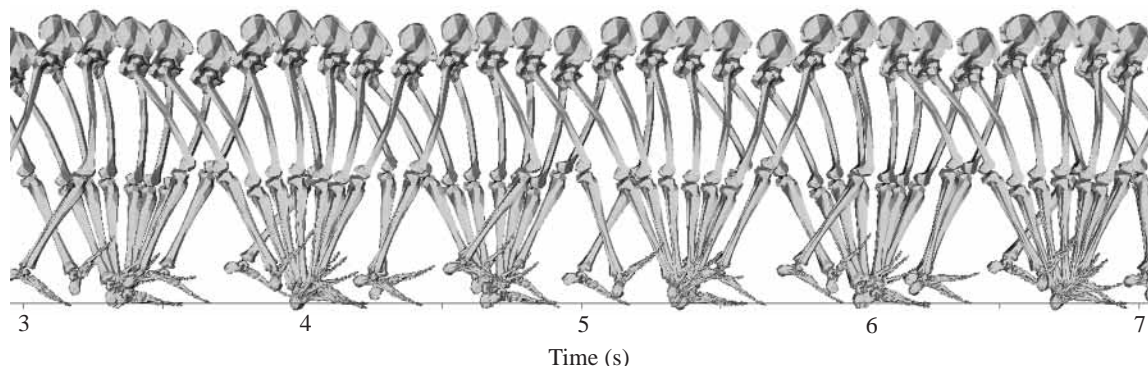

Fig. 10. Overlay image of the animation generated by the best genome when starting with a population derived from a hand-crafted genome (see movie3.mov). The images are 0.2 s apart.
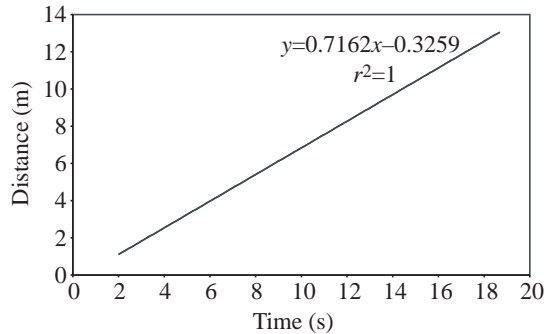
Fig. 11. Graph showing the forward displacement of the torso with respect to time.
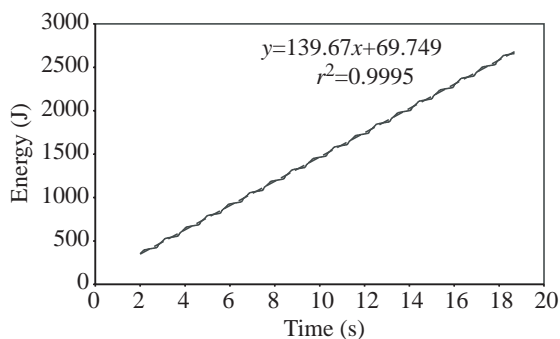

Fig. 13. Graph showing the total metabolic cost of the muscles with respect to the forward distance traveled.

Fig. 12. Graph showing the total metabolic energy cost of the muscles with respect to time.

Fig. 14. Experimentally derived data on the metabolic energy cost of human locomotion. Based on data from Alexander (1992b).

sufficient and long-lasting enough to bring the oscillations of kinetic and potential energies of the body centre of mass into phase, eliminating energy exchange and thus increasing energy costs.)

To further investigate the simulated gait we plotted the characteristics of the gait once it reached a steady cyclic phase (approximately 2 s into the simulation). These data are shown in Figs 11, 12 and 13. Fig. 11 shows the forward displacement of the torso over time. The forward velocity of the torso is very nearly constant at $0.72\,\mathrm{m\,s^{-1}}$. Fig. 12 shows the metabolic energy costs over time. There is a small amount of fluctuation, depending on the phase of the gait cycle, but the mean power is 140 W. Fig. 13 shows the metabolic energy cost against forward distance traveled, and the mean value for the specific energy costs is $200\,\mathrm{J\,m^{-1}}$.

Discussion

The primary aim of the model was to estimate the optimum energetic cost of human bipedal locomotion almost entirely from the proportions of the skeleton. The best figure calculated was $200\,\mathrm{J\,m^{-1}}$. Fig. 14 shows the experimentally derived metabolic cost of human locomotion and the best value obtained is $230\,\mathrm{J\,m^{-1}}$ (Alexander, 1992b). Thus the model's prediction is correct to within 15% of empirical values, which is extremely encouraging and suggests that this approach could
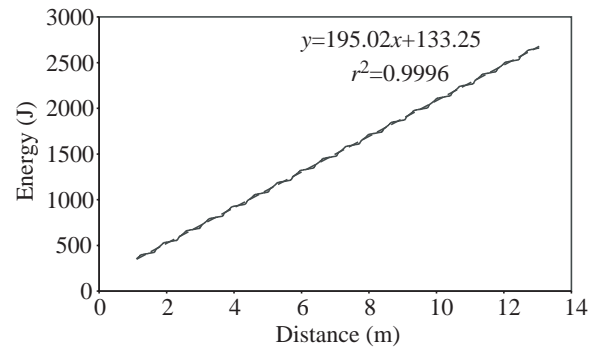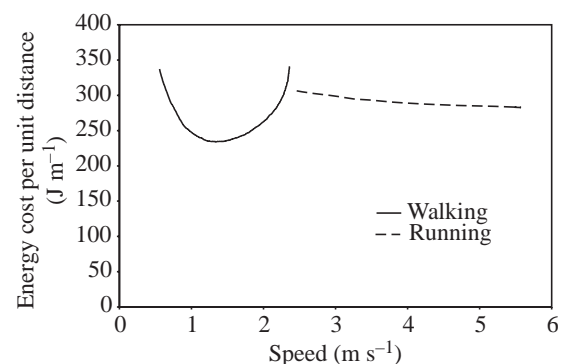
reliably be used to investigate costs in other bipeds. However, the details of the gait produced suggest that the model needs improvement before it could be used for wider investigations of locomotor performance. The speed at which the most efficient bipedalism was obtained in the model was $0.72\,\mathrm{m\,s^{-1}}$, which is appreciably lower than the $1.3\,\mathrm{m\,s^{-1}}$ obtained experimentally. We would suggest two reasons for this discrepancy. Firstly there are no spring elements in our model. Spring elements would tend to accelerate limbs back towards their resting positions, which would tend to increase the pendular frequency of the limbs, in turn leading to a faster optimal gait. Secondly our model has only a very short start-up phase, so the eventual walking pattern has to be stable at low speeds that will limit the speed the model can sustain later in the simulation. Indeed, our preliminary investigations into using this technique for top-speed analysis support this suggestion since, although we can evolve a running gait, the current top speed is $2.2\,\mathrm{m\,s^{-1}}$. Also, since the genetic algorithm is a stochastic process we may simply have not run enough repeats to find a better solution. Our best solution was still improving when we terminated the program and, given enough time, we would have found a better answer. Inspection of the fitness curves suggests that our evolutionary algorithm tends to generate periods of fitness stasis interspersed with short periods of rapid fitness improvement, corresponding to the

model jumping from a local maxima to an ultimately fitter region of the search space.

From an evolutionary perspective, it would also be desirable to substitute into the model not only alternative body proportions, but also alternative sets of segment mass distributions and muscle parameters. Some muscle parameters (such as maximum tension and length/tension relationships) can only be derived by types of experimentation that cannot ethically be performed on non-human primates, and hence we have no choice but to use values for humans (and phylogenetic considerations, of course, suggest that human values are the most likely to be appropriate for extinct hominids). But others, such as physiological cross-sectional areas and lever arms, can be derived from cadavers, as can segment inertial properties. We have now gathered the latter data for all the living apes, and future research will therefore examine the consequences of different muscle geometries, and can be expected to identify which adaptive features of living apes would be most likely to favour the acquisition of habitual bipedality. Other studies using a similar approach can and will examine the performance effectiveness of given morphologies in activities such as climbing, and together these studies will help identify the range of behaviours that extinct species were best adapted to perform.

The model would also benefit from greater biofidelity. For example, it could be fully three-dimensional; it could incorporate sensory feedback to modulate the performance and improve balance; the muscle groups could be subdivided into smaller functional units including two-joint as well as single-joint muscles; the finite-state control system could employ a greater number of states in the control system to provide more precise control; and spring elements could be incorporated into the muscle model. This poses two problems. Firstly, the simulation itself would then take longer to compute, which would increase the duration of the fitness evaluation. Secondly, a more complex model would require a larger genome to specify the parameters, and this would very greatly increase the size of the search space. The simulation problem is not completely intractable, since an advantage of the genetic algorithm approach is that it can be efficiently implemented to take advantage of massively parallel computer systems. The current implementation runs on a 4-processor computer and would scale to run on the next generation of 1000-processor clusters with very little overhead. This would reduce the current experimental time of weeks to hours, so that considerable increase in model complexity would be possible. The search space problem is more difficult to solve since the size of the space grows far too rapidly to be overcome by simply increasing the computational power available. However, an incremental evolutionary approach may be the answer, since we can evolve a set of good parameters from the current model and use these as the starting point for a series of more complex models. In this way we should be able to restrict the search to a region of the search space that is likely to be profitable.

One interesting outcome is the spontaneously derived walking pattern shown in Fig. 8. The knee is kept fully extended at all times and movement at the ankle joint is used to allow the swinging leg to clear the ground. This walking pattern is completely different from anything we might have expected and, although not as efficient as normal walking, it is mechanically stable and shows an alternative solution to the fundamental mechanical problem. When evaluating the locomotion of fossil forms we tend to limit ourselves to patterns normally seen in extant animals, but given the limited amount of direct evidence available from fossil footprints it is important to consider alternative locomotor strategies. This will be even more important when looking at older fossils since their morphology can be very different from modern animals.

Another important area for further work is investigating the sensitivity of the model predictions to the fixed modeling parameters. The dimensions of the skeleton can be accurately measured, but there is appreciable variation between individuals. Mass distribution parameters need to be estimated and there is considerable variation in the values reported in the literature for humans, although it has been suggested that this variation may not be important (Yoko et al., 1998). Muscle mechanical parameters vary from muscle to muscle and from individual to individual. The current study used generic values, but there is a need to use subject-specific values to allow direct experimental validation – especially in the case of 'non-standard' groups such as children and non-human primates. This would also allow us to identify the modeling parameters that have a disproportionate effect on the outcome and should therefore experience the greatest evolutionary pressure. We appreciate that soft tissue parameters will always be difficult to predict for fossil forms and would encourage the further development of techniques to aid such predictions.

The animation generated by the model was qualitatively impressive. Since it was based on a simulation, the ground and foot interaction was convincing with minimal intersection and, unlike motion-capture driven animation, the animation was not limited to a few gait cycles that then need to be reused to generate longer walking sequences. Our sequences could cover over 50 m, and although the gait looks smooth there are small differences between each step, which adds to the naturalness of the effect. For use in animation we would need to incorporate a fully articulated upper body to allow arm-swinging, since this is visually important. This would also increase biofidelity, since arm-swing in human walking seems to assist free vertical moments under the foot in balancing trunk torques (Li et al., 2001). An added bonus is that each run of the simulation will produce a slightly different optimal gait so that each character in an animation can easily have a realistic gait that is unique for that character.

References

Alexander, R. McN. (1989). *Dynamics of Dinosaurs and Other Extinct Giants*. New York: Columbia University Press.

Alexander, R. McN. (1991). Characteristics and advantages of human bipedalism. In *Biomechanics in Evolution* (ed. J. M. V. Rayner and R. J. Wooton), pp. 255-266. Cambridge: Cambridge University Press.

Alexander, R. McN. (1992a). A model of locomotion on compliant legs. *Phil. Trans. R. Soc. Lond. B* **338**, 189-198.

Alexander, R. McN. (1992b). Human Locomotion. In *The Cambridge Encyclopedia of Human Evolution* (ed. S. Jones, R. D. Martin and D. Pilbeam), pp. 80-85. Cambridge: Cambridge University Press.

Alexander, R. McN. and Jayes, A. S. (1978). Vertical movements in running and walking. *J. Zool. Lond.* **185**, 27-40.

Cavagna, G. A., Heglund, N. C. and Taylor, C. R. (1977). Mechanical work in terrestrial locomotion, two basic mechanisms for minimizing energy expenditure. *Am. J. Physiol.* **233**, R243-R261.

Crompton, R. H., Li, Y., Günther, M. M. and Savage, R. (1998). The mechanical effectiveness of erect and 'bent-hip, bent-knee' bipedal walking in *Australopithecus afarensis*. *J. Human Evol.* **35**, 55-74.

Crompton, R. H., Li, Y., Thorpe, S. K., Wang, W. J., Savage, R. Payne, R., Carey, T. C., Aerts, P., Van Elsacker, L., Hofstetter, A., Günther, M. M., D'Aout, K. and DeClerq, D. (in press). The Biomechanical Evolution of Erect Bipedality. *Courier Forschuungsinstitut Senckenberg.*

Davis, L. (1991). *Handbook of Genetic Algorithms*. Van Nostrand Reinhold: New York.

Delp, S. L., Loan, J. P., Hoy, M. G., Zajac, F. E., Topp, E. L. and Rosen, J. M. (1990). An interactive graphics-based model of the lower extremity to study orthopaedic surgical procedures. *IEEE Trans. Biomed. Eng.* **37**, 757-767.

Doran, D. M. (1996). Comparative positional behavior of the African apes. In *Great Ape Societies* (ed. W. C. McGrew, L. F. Marchant, L. F. and T. Nishida), pp. 213-225. Cambridge: Cambridge University Press.

Fleagle, J. G. (1999). *Primate Adaptation and Evolution*. Academic Press: San Diego.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press: Mitchigan.

Kramer, P. A. (1999). Modelling the locomotor energetics of extinct hominids. *J. Exp. Biol.* **202**, 2807-2818.

Li, Y., Crompton, R. H., Wang, W.-J. and Günther, M. M. (2002). Reconstructing the mechanics of quadrupedalism in an extinct hominoid. *Zeits. Morphol. Anthropol.* **83**, 265-274.

Li, Y., Wang, W.-J., Crompton, R. H. and Günther, M. M. (2001). Free vertical moments and transverse forces in human walking and their role in relation to arm-swing. *J. Exp. Biol.* **204**, 47-58.

Ma, S. and Zahalak, G. I. (1991). A distribution-moment model of energetics in skeletal muscle. *J. Biomech.* **24**, 21-35.

Martin, R. D. (1990). *Primate Origins and Evolution: A Phylogenetic Reconstruction*. Princetown University Press: Princeton, New Jersey.

McMillan, S. (1994). Computational dynamics for robotic systems on land and under water. PhD thesis, Ohio State University, Columbus, OH, USA.

McMillan, S., Orin, D. E. and McGhee, R. B. (1995). DynaMechs: An object oriented software package for efficient dynamic simulation of underwater robotic vehicles. In *Underwater Robotic Vehicles: Design and Control* (ed. J. Yuh), pp. 73-98. Albuquerque, New Mexico: TSI Press.

Minetti, A. E. and Alexander, R. McN. (1997). A theory of metabolic costs for bipedal gaits. *J. theor. Biol.* **186**, 467-476.

Napier, J. R. and Napier, P. H. (1985). *The Natural History of The Primates*. London: British Museum (Natural History).

Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics*. Cambridge, Mass.: MIT Press.

Ogihara, N. and Yamazaki, N. (2001). Generation of human bipedal locomotion by a bio-mimetic neuro-musculo-skeletal model. *Biol. Cybern.* **84**, 1-11.

Oxnard, C. E., Crompton, R. H. and Liebermann, S. S. (1989). *Animal Lifestyles and Anatomies: The Case of the Prosimian Primates*. Seattle: Washington University Press.

Pierrynowski, M. R. (1995). Analytical representation of muscle line of action and geometry. In *Three-dimensional Analysis of Human Movement* (ed. P. Allard, I. A. F. Stokes and J.-P. Blanch), pp. 215-256. Champaign, Illinois: Human Kinetic Publishers.

Rechenberg, I. (1965). *Cybernetic Solution Path of an Experimental Problem. Ministry of Aviation*. Royal Aircraft Establishment (UK): Ministry of Aviation.

Rose, M. D. (1993). Locomotor anatomy of Miocene hominoids. In *Postcranial Adaptation in Nonhuman Primates* (ed. D. L. Gebo), pp. 252-272. DeKalb, Illinois: Northern Illinois University Press.

Sellers, W. I. (1996). A biomechanical investigation into the absence of leaping in the locomotor repertoire of the Slender Loris (*Loris tardigradus*). *Folia Primatol.* **67**, 1-14.

Senut, B., Pickford, M., Gommery, D., Mein, P., Cheboi, K. and Coppens, Y. (2001). First hominid from the Miocene (Lukeino Formation, Kenya). *C. R. Acad. Sci. II* **332**, 137-144.

Wang, W.-J., Crompton, R. H. and Günther, M. M. (2003). Optimum ratio of upper to lower limb lengths in hand-carrying of a load under the assumption of frequency coordination. *J. Biomech.* **36**, 249-252.

Winter, D. A. (1990). *Biomechanics and Motor Control of Human Movement*. New York: John Wiley and Sons.

Yamazaki, N., Hase, K., Ogihara, N. and Hayamizu, N. (1996). Biomechanical analysis of the development of human bipedal walking by a neuro-musculo-skeletal model. *Folia Primatol.* **66**, 253-271.

Yoko, T., Takahashi, A., Okada, H., Ohyama, K. B. and Muraoka, M. (1998). Is the selection of body segment inertia parameters critical to the results of kinematic and kinetic analysis of human movement? *Anthrop. Sci.* **106**, 371-383.

Zatiorski, V. and Seluyanov, V. (1983). The mass and inertia characteristics of the main segments of the human body. In *Biomechanics VIII-B* (ed. H. Matsui and K. Koayashi), pp. 1154-1159. Champaign, Illinois: Human Kinetic Publishers.