

Grammatical Evolution for Gait Retargeting

J. E. Murphy, H. Carr and M. O'Neill

University College Dublin, Ireland

Abstract

Artists and scientists require tools to construct physics-based animal models. However, animating these models requires motion data for realistic movement. Motion data may either be measured from real-life animals-in-motion or generated using an optimisation approach. We propose a solution for retargeting gait data from one animal to another. The retargeted gait cycles are generated using a Grammatical Evolution optimisation approach and the search space is constrained based on dynamic similarity principles.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism

1. Introduction

Physics-based animation uses the laws of physics to govern the movement of computer-constructed models. For animals, some motion data must be input for realism. Artists can generate motion data manually, but this is prohibitively expensive, as is capturing data from an animal. Optimisation approaches can prove computationally prohibitive.

To reduce the expense involved in generating motion data, we propose to retarget gait data from one animal to another. The target animal's gait cycle is generated via a series of hybrid models. We treat an animal's muscle groups as terms in a Fourier analysis and use an evolutionary algorithm to optimise them. Each optimisation is seeded with available gait data and constrained by dynamic similarity principles.

2. Problem

A physics-based animal model is constructed from rigid bodies (bones) connected by joints. A physics engine makes the bones react in a realistic manner to gravity, friction, and other forces. A small number of parameters are used to compute the joint torque forces, which generate desired motion.

The motions thus generated are discrete approximations of a gait cycle: the pattern of a limb's motion. Data for a single cycle is then repeated to sustain locomotion. Data can be acquired from photographs, videos or publications. Extracting and standardising this data is expensive. Data for many

animals is unavailable. This lack of data motivates the need for an inexpensive gait retargeting solution.

3. Related work

Gait retargeting modifies one animal's gait cycle data to another. Popović & Witkin [PW99] use an intermediate model to map animation from one biped to another. Monzani et al. retarget biped motion using an intermediate skeleton in [MBBT00]. Marsland & Lapeer retarget a kinematic animation of a trotting horse to a dynamic model in [ML05]. This work, however, retargets between animals of the same species, rather than between species.

Gait retargeting is possible through use of gait generating optimisation approaches, such as those presented in [vdPL95, HvS02]. The optimisation process is performed using computer simulations of animal locomotion, such as those presented in [HM01, vdB89]. Use of evolutionary algorithms for gait generation is gaining popularity, especially in the robotics field [KKW*02, GH06].

Evolutionary algorithms involve the simulation of evolving populations of solutions, guided by a fitness function towards an optimal solution. We propose to apply a relatively new type of evolutionary algorithm called Grammatical Evolution to gait retargeting. It has been successfully employed for financial prediction [OBRC01], but has not been applied to gait generation to date.

3.1. Grammatical Evolution

Grammatical Evolution (GE) adapts principles from molecular biology [OR03] to optimisation problems by evolving solutions from one generation to the next. In particular, GE applies evolutionary computation to a genotype rather than a phenotype, as in nature. Here, the phenotype is the solution in the problem domain, while the genotype is an abstraction into a more compact form. Thus, unlike other evolutionary algorithms, GE separates the search space and solution space for higher efficiency.

GE genotypes are typically either binary or integer strings that describe solutions, as DNA describes proteins. A grammar expands the genotype into a phenotype in the same way as DNA is expanded to a protein. GE then applies operators to each generation of genotypes to produce the next, evaluating the fitness of each solution in the problem domain - i.e. at the phenotypic level. High fitness solutions are preferentially transferred to the next generation, ahead of low fitness solutions. Because GE adopts a grammar-guided genotype-phenotype mapping, constraints may be applied to the fitness function to increase optimisation efficiency.

3.2. Dynamic similarity

We propose to constrain GE with observations from natural evolution and principles of dynamic similarity. These principles, demonstrated by Alexander & Jayes in [AJ83], observe that animals moving at equal values of the dimensionless Froude number have similar gaits.

While the Froude number does not predict properties of a gait pattern, if gait characteristics for one animal travelling at a particular Froude number are known, then those characteristics should hold for another animal travelling at the same number. Characteristics of interest include limb phase relationship, relative stride lengths and duty factors.

4. Solution components

An animal's skeletal dimensions and musculature have gradually adapted to its environment over millions of years through evolution [Fut79]. As such, cursorial quadrupeds have highly similar skeletal structures, most obviously differing in bone proportions.

To accomplish gait retargeting, we propose a constrained gait generation optimisation based on the naturally incremental nature of evolution. By retargeting gaits via a series of hybrid models, whose bone lengths change gradually, the optimisation remains close to a global minimum at all times.

To demonstrate this approach, we will retarget a horse's motion to a dog model through a series of horse-dog hybrid models (Figure 1). Each discrete hybrid model is a linear interpolation between the horse and dog. As the gait retargeting process proceeds, the bone proportions of subsequent hybrid models tend towards those of the target animal.

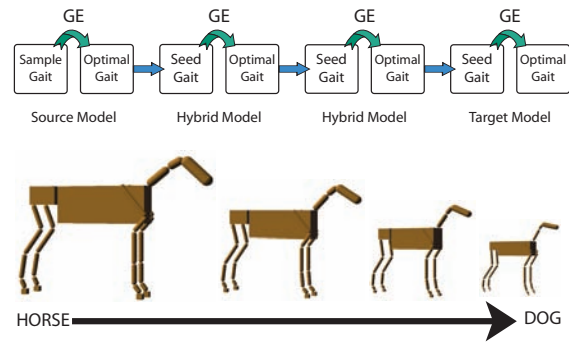


Figure 1: Incremental gait retargeting with Grammatical Evolution (GE). Optimisations are performed on a series of intermediate horse-dog hybrid models.

GE is first used to generate an optimal gait cycle for a pure horse model from measured data. This gait cycle then becomes a seed for a hybrid model, and GE generates a new optimal gait cycle. This process is repeated until a pure dog model is reached and optimised.

Optimality is judged on energy efficiency, since an animal's musculoskeletal system has evolved to minimise energy expenditure. Differences in musculature and bone proportions between species result in different gaits. As the incremental optimisation process progresses, each newly generated gait will further resemble the target model's real-life equivalent gait, due to the energy efficiency scoring system.

This proposed retargeting solution has two major components: the application in which a physics-based quadruped model can be simulated, and the optimisation approach.

4.1. Simulation environment

The core work to date is a horse simulation and gait development application (Figure 2). This takes skeletal data as input and constructs a 3D model using the Open Dynamics Engine [Smi06]. Skeletal data is acquired from the biomechanics literature [BSSB97] and joint inclination values from Muybridge [Muy85].

To produce movement, motion controllers apply torques about the joints. Torques are based on joint-angle data extracted from plots in the veterinary literature [BC01]. Motion curves are stored as discretised smooth approximations, where each value is a target angle at a given time. The torque required is then calculated using Hooke's law. These torques are not analogous to the forces applied by real-life muscle groups. Rather, successive "push" and "pull" torques move each bone towards its target orientation at each time step.

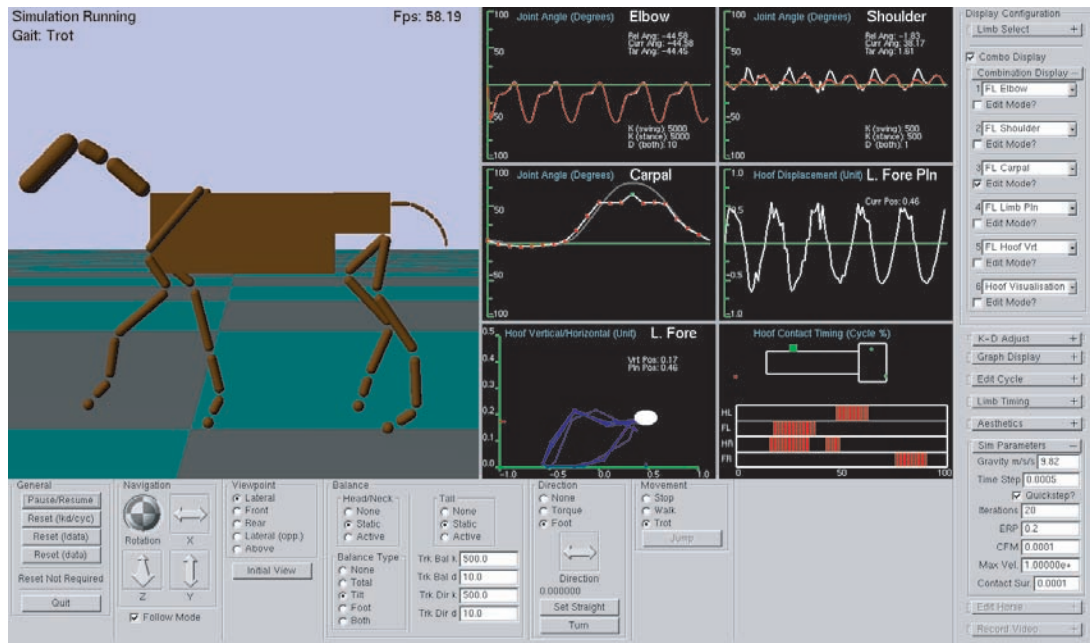


Figure 2: A screenshot of the simulation application shows an interactive physics-based horse animation alongside motion data visualisations. The user configures visualisation displays to show a combination of the following: current vs. target joint-angle data, target data in edit mode, limb extension plots, hoof-lift plot, hoof contact and timing visualisation.

4.2. Gait cycle representation

Gait data is stored as samples from a time varying function. Each sample is an optimisation parameter, resulting in a huge search space, which we reduce by applying Fourier analysis to the gait cycle to reduce the number of parameters. We can do so because locomotion involves muscles contracting and relaxing in a sinusoidal fashion. Thus, the gait cycle naturally lends itself to Fourier analysis.

Fourier analysis decomposes motion data into sinusoidal functions for a range of frequencies (Figure 3). The most influential peaks are recombined to form the minimal gait representation. Smaller peaks are discarded, as they are assumed to be unimportant. Representing a joint’s motion as a sum of phase and amplitude values reduces the search space and works well with GE.

4.3. Gait generation with Grammatical Evolution

To generate gaits, we apply GE to Fourier gait genotypes, initialising with random variations of a seed gait cycle. These random variations may be more or less efficient than the seed, but GE will cause efficient solutions to dominate over time, leading to an optimal solution.

A grammar translates the genotypic strings into phenotypic Fourier terms, which vary by frequency, amplitude and phase. Thus, compact genotypes are converted to complex

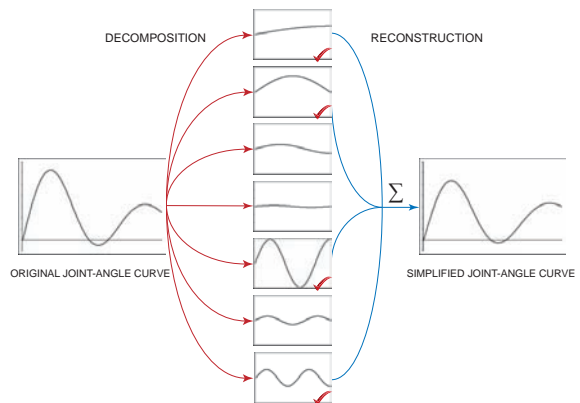


Figure 3: Fourier analysis decomposes gait data into a number of sinusoidal functions whose parameters are used to represent the gait.

phenotypic gait cycles which can be used for locomotion, as illustrated in Figure 4.

Each gait is tested in the simulation application and evaluated by a fitness function, based on energy efficiency of the gait. Energy efficiency is calculated in a manner similar to [KKW*02]. A valid solution must result in the animal travelling a set distance, in the required time whilst not violating any motion constraints.

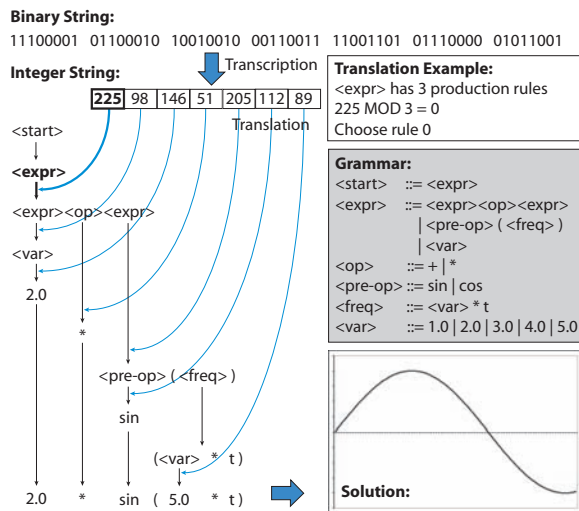


Figure 4: Generating a sinusoidal function from a binary string using a simple grammar.

4.4. Constraints for gait generation

An animal model's numerous joints present a huge search space of potential solutions. By exploiting knowledge of the problem domain, the optimisation can be constrained and the search space reduced. The dynamic similarity principles introduced in Section 3.2 can provide such a constraint.

As we retarget a gait cycle from one animal to another, the dynamic similarity hypothesis predicts that the target's optimal gait will share characteristics with the source gait. The limb phase differences of the source gait will be imposed on all target model simulations via the motion controllers. During the optimisation process, any potential solution that does not exhibit duty factors and relative stride lengths comparable to the source gait is rejected. The use of intermediate hybrid models minimise the variance between source gait and the generated gaits. In addition to this, the dynamic similarity constraints ensure that resources are not wasted on the exploration of unacceptable solutions.

5. Future work

In future, we intend to adapt the horse simulation application to arbitrary quadrupedal and hybrid skeletal data files. The Fourier gait representation will then be incorporated, as well as fitness calculations based on energy efficiency. This application will then be used as the fitness function for the evolution of populations of animations, using GE.

Dynamic similarity constraints will be implemented as a simple acceptance test associated with the GE fitness function. Although some constraint data is available through published sources, constraints for a full range of motion may require some video or motion capture data.

References

- [AJ83] ALEXANDER R. M., JAYES A. S.: A dynamic similarity hypothesis for the gaits of quadrupedal mammals. *Journal of Zoology (London)* 201 (1983), 135–152.
- [BC01] BACK W., CLAYTON H. M.: *Equine Locomotion*. Harcourt Publishers, 2001.
- [BSSB97] BUCHNER H. H. F., SAVELBERG H. H. C. M., SCHAMHARDT H. C., BARNEVELD A.: Inertial properties of dutch warmblood horses. *Journal of Biomechanics* 30, 6 (June 1997), 653–658.
- [Fut79] FUTUYMA D. J.: *Evolutionary Biology*. Sunderland, 1979.
- [GH06] GARDER L. M., HØVIN M. E.: Robot gaits evolved by combining genetic algorithms and binary hill climbing. In *Proc. GECCO* (July 2006), pp. 1165–1170.
- [HM01] HERR H. M., MCMAHON T. A.: A galloping horse model. *Journal of Robotics Research* 20, 1 (Jan. 2001), 26–37.
- [HvS02] HARDT M., VON STRYK O.: Increasing stability in dynamic gaits using numerical optimization. In *Proc. IFAC* (July 2002), pp. 1636–1641.
- [KKW*02] KIGUCHI K., KUSUMOTO Y., WATANABE K., IZUMI K., FUKUDA T.: Energy-optimal gait analysis of quadruped robots. *Artificial Life and Robotics* 6, 3 (Sept. 2002), 120–125.
- [MBBT00] MONZANI J.-S., BAERLOCHER P., BOULIC R., THALMANN D.: Using an intermediate skeleton and inverse kinematics for motion retargeting. *Computer Graphics Forum* 19, 3 (Sept. 2000), 11–19.
- [ML05] MARSLAND S. A., LAPEER R. J.: Physics-based animation of a trotting horse in a virtual environment. In *Proc. InfoVis* (July 2005), pp. 398–403.
- [Muy85] MUYBRIDGE E.: *Horses and other animals in motion*. Dover Publications, 1985.
- [OBRC01] O'NEILL M., BRABAZON T., RYAN C., COLLINS J.: Developing a market timing system using grammatical evolution. In *Proc. GECCO* (2001), pp. 1375–1381.
- [OR03] O'NEILL M., RYAN C.: *Grammatical Evolution*. Kluwer Academic Publishers, 2003.
- [PW99] POPOVIĆ Z., WITKIN A.: Physically based motion transformation. In *Proc. SIGGRAPH* (1999), pp. 11–20.
- [Smi06] SMITH R.: *Open Dynamics Engine v0.5 User Guide*, Feb. 2006. <http://www.ode.org/>.
- [vdB89] VAN DEN BOGERT A. J.: *Computer simulation of locomotion in the horse*. PhD thesis, Univ. Utrecht, 1989.
- [vdPL95] VAN DE PANNE M., LAMOURET A.: Guided optimization for balanced locomotion. In *Computer Animation and Simulation* (Sept. 1995), pp. 165–177.