Biomechanical Principles of Motion

SIGGRAPH 2009 Course Notes: Realistic Human Body Movement for Emotional Expressiveness

Aaron Hertzmann

May 20, 2009

Despite some early enthusiasm, early attempts at physics-based character animation foundered due to the difficulty of creating realistic and expressive models, as well as significant computational burdens. In comparison, motion capture and keyframe animation can give good results, provided enough time and effort are spent with them. However, neither method gives a fully realistic and flexible model of motion that can be used to generate highly-realistic motions in new circumstances. For example, motion capture provides little ability to create new motions that are very different from the data, and modifications of mocap suffer noticeable violations of physics, such as footskate and implausibly large forces. There is now a small but growing resurgence of interest in physics-based character animation. Whereas mocap and animation systems require significant labor to create motion, physics-based animation offers the promise of greater generality and flexibility. Many open research problems remain before physics-based animation becomes practical.

The principles of human motion connect many fields of scientific research, including biomechanics, optimal control, machine learning, robotics, motor neuroscience, psychology, and others, as well as theatre, animation, and dance. Each of these fields can give a different perspective on motion, each of which is useful for understanding how we move. This lecture aims to survey the most relevant principles from these areas, with an emphasis on human locomotion (especially walking).

Even from a physical point of view, there are many ways we can look at motion. We can inspect all the individual forces involved in a motion, or we can look at the forces on the system as a whole and the center-of-mass, or we can look at higher-level properties such as energy and work.

Research questions: why do we move the way we do, and what role does physics play? How do we model human rewards and objectives in movement? How can we deal with the difficult optimization problems that arise? How do we use physical models in artistic contexts?

1 Basic physical models and simulation

- The body is typically modeled as an articulated rigid-body system. Most of the details here are standard in physical simulation, and well-known at SIGGRAPH. However, different choices in the body model affect the body's natural modes of movement.
- For example, even topological choices make a difference: you cannot have toe-off — a crucial component of walking — without toes. Yet most physics-based models to date do not have toes. As a further thought experiment, imagine (or observe) walking in ski boots (where the toe is rigid and the ankle is nearly rigid), as compared to walking barefoot, in formal shoes, or high heels.
- When building a physical model, one has a number of choices in terms of degree of realism. Simpler methods are usually adequate for very simple problems (e.g., ragdoll simulation), but usually fail at capturing nuances of active locomotion, especially when optimizing motion and the motion can "exploit" modeling approximations and errors.
  - An important choice is constraint handling, at contacts and joint limits.

  - One can use exact constraint handling, or penalty models. While exact methods are much more complex, penalty methods can lead to serious issues, especially if motion is being optimized. For example, if a penalty-method is used for ground contact, the character can walk efficiently by bouncing on the ground as if walking on a trampoline. It is usually easiest to begin with simple models, but be prepared to move to more complex models when problems arise.

- Parameterization either with generalized coordinates $\mathbf{q}$ (i.e., joint angles and root position/orientation), or positions of each link. Latter has simple equations of motion, but requires rigid-link constraints, which is usually more trouble than it's worth.

- Other important constraints: joint limits, torque limits, interpenetration constraints (these are usually ignored).

- Joint angle parameterization: it is often observed that bad parameterization leads to bad motion, and this is true for simulation, and especially control. On every physics-based animation project that I have worked on or witnessed, the student begins with Euler angles because they are simplest, and eventually concludes that they have too many problems, and eventually ends up using exponential maps.

- Segment (link) masses: these can be taken from standard tables. These tables usually show little variance in relative masses (e.g., proportion of arm mass to total body mass), and so can be used to scale according to body mass. However, this data may be somewhat limited in range; e.g., datasets based on cadavers of military personnel. Segmental parameters may be heuristically used to determine inertia tensors.

- Given the physical model, simulation may be performed using the robotics equations-of-motion (EOMs), usually written:

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) + \mathbf{G}(\mathbf{q}) = \boldsymbol{\tau} \tag{1}$$

  where $\mathbf{M}$ is the mass matrix, $\mathbf{C}$ is Coriolis/centrifugal forces, and $\mathbf{G}$ are gravitational forces, and $\boldsymbol{\tau}$ are the joint torques. These equations can be thought of as a generalized form of Newton's second law $=$ $\mathbf{ma}$, where the individual terms depend on the current configuration $\mathbf{q}$ and generalized velocities $\dot{\mathbf{q}}$.

- A controller is a mapping from state $\mathbf{q}, \dot{\mathbf{q}}$ to torques $\boldsymbol{\tau}$:

$$\boldsymbol{\tau} = f(\mathbf{q}, \dot{\mathbf{q}}) \tag{2}$$

  The controller may have a persistent internal state as well, e.g., a state machine. The big question in simulation of characters is: what are the joint torques $\boldsymbol{\tau}$? This depends on the character's control mechanism, and will be discussed in more detail below.

- Given the body model, a controller, and an initial state $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$ the basic steps of simulation are:

  1. Compute control torques from the controller
  2. Solve EOMs to estimate accelerations $\ddot{\mathbf{q}}$
  3. Simulate forward one time-step, while resolving ground contacts, as necessary. Explicit or implicit integration may be used.

- The EOMs and their solution can be derived either from Lagrange's Principle of Least of Action, or using Featherstone's Algorithm. Featherstone's algorithm is the most efficient algorithm, though more involved, and there is no single tutorial paper/book that explains every aspect of it clearly (see Mirtich's thesis, the review by Featherstone and Orin, and Featherstone's book).

- In practice, detailed solution of the EOMs are too complex to do by hand, especially since you will often need to experiment with different parameterizations, ways of handling contact and so on. Hence, it is usually necessary to employ a symbolic math data structure (often home-grown) that can compute the EOMs, necessary Jacobians, and derivatives automatically.

2 Sources of data

- When modeling motion it is crucial to look at real movement. Generate hypotheses by watching motion. Evaluate models by comparing them to new data.
- Sources of data: video, motion capture, force plate data, electromyography (measures muscle activation)
- There is no existing way to measure every force, mass, and velocity in the body. Some biomechanicists have turned to physical simulation as one way to test hypotheses.
- Personal laboratory: you can learn a lot simply by "observing" your own motion as you move around; sit in a coffeeshop and watch people walk. What trajectories do your feet take when they walk? When are they applying forces? What role does arm-swing play: if you hold your arms rigid by your sides and try to walk or jog, what happens? etc.

3 Optimality principles of movement

- Evolution has an amazing ability to craft biological systems to get ahead and multiply. This can be mathematically modeled as a process of optimization. The most natural objective function to describe biological organisms is reproductive fitness: e.g., number of grandchildren. But we can often use more short-term goals that organisms might have. Optimality principles have been used to explain many aspects of biological systems, such as bone densities, foraging behaviors, to movement.
- What is the objective function? Basic concerns in movement include cost-of-transport, stability, balance, and achieving goals (e.g., get from one place to another).
- Optimality, in principle, allows significant generalization that other models cannot: if you know my objective function, then you should be able to predict how I move in a wide variety of new situations, just by adding appropriate new constraints.
- The cost of transport is defined as the energy used divided by distance travelled. Note that measuring external work is a lower-bound (e.g., if you end up where you started, your total work is zero, but you have consumed energy nonetheless). Dimensionless cost of transport (cost of transport divided by weight) can be used to compare efficiency different robots/characters.
- In experiments with humans, energy consumption is normally measured in terms of the amount of oxygen consumed.
- Two main approaches:
  - Trajectory optimization ("Spacetime constraints"): obtain the single-best trajectory. Requires offline optimization/planning, and cannot respond to real-time commands and perturbations (e.g., new forces, user control)
  - Control synthesis/Optimal Control: define a mapping from state to torques. Can theoretically run in real-time, but needs to be able to "plan ahead." More on this below.

- Limitations of optimality theory:
  - Optimality can be controversial in biology; not all evolution is adaptation.
  - We are not always "optimal," e.g., compare a novice athlete or dancer to an expert. If one receives poor training, one may converge to a poor local minimum, e.g., you can get very good at snowplow skiing without ever making the leap to parallel. Dance instructors try to get you out of bad habits as soon as they can. For such reasons, optimality approaches can be controversial.
  - Simple objective terms say nothing about "style" (e.g., social signalling behaviors). We can sometimes embed "style" in physics: e.g., dragging your feet by modeling shoes as literally heavy; jaunty, bouncing style as the product of overdamped spring-mass systems.
  - Local optima and search problems are formidable: search space is very high-dimensional, dynamics are nonlinear (esp. due to ground contact), local optima are very prevalent. For low-energy motions (such as walking), the local minima problems get even worse.

4 Phases of walking

- It is useful to observe in detail the phases of walking. (You can find every aspect of it observed in minute detail in the biomechanics literature; much biomechanics research is very descriptive).
- Walking: Swing phase, heel strike, full-foot contact, roll, heel-off, toe-off. Try this at home.
- The foot: extremely complex; the more you can model, the better. Ankle provides approx 53% of energy consumption during walking.

5 Simplified Mechanical Models

- Study simplified mechanical systems for insight into movement
- Most powered robots (e.g., Asimo) appear very stiff and energy efficient; they use very inefficient control strategies.
- Inverted pendulum: a simple one-parameter system matches human center-of-mass (COM) motion in walking
- Bouncing ball: matches human COM during running
- Passive-dynamic walking: McGeer showed a simple passive robot that walks downhill, powered only by gravity. Very similar to human movement. Much work on analysis of motion and stability has been done in the context of these simplified 2D models; e.g., see work by Art Kuo.
- Passive-based walking: Collins et al. describe level-ground walkers based on passive principles. These robots are extremely energy efficient as compared to commercial robots, and only a little less efficient than humans.
- Templates and anchors: an approach to simplifying locomotor systems to their essence, which allows comparing related biological models (e.g., different animals that move in similar ways)
- Relevance for animation: da Silva et al. and Brubaker et al. demonstrate animation/tracking models that relate a low-degree-of-freedom mechanical model to a high-DOF kinematic model (without physics). Provides a natural "dimension reduction" or "mechanical model," but may be difficult to enforce consistency between the models, especially when you deviate from the model.
- Kry et al. apply modal analysis to obtain a low-dimensional model of locomotion (though ignoring contacts).

6 Muscles, Bones, Ligaments

- The detailed workings of muscles, bones, ligaments, and tendons is itself quite involved, and the subject of many papers in both biomechanics and animation. One often abstracts away these details for full-body motion, with a corresponding loss of fine detail. Many highly-detailed models of individual substructures (e.g., neck, knee, hands, face, etc.) have been devised, illustrating the importance of these body parts
- Muscles supply active forces. Muscles attach to bones via tendons, and bones connect to each other via ligaments.
- Muscles, tendons, and ligaments all have passive elastic (spring and damper properties). Compression and expansion of passive elements conserves about 30% of energy during running.
- Muscles can only contract (pull, not push). Muscles come in agonist/antagonist pairs: one muscle pulls one way, and the other pulls the other way. Both muscles can contract to create stiffness/tension. Stiffness depends on task, e.g., compare walking on cement to sand/mud or ice, and has a noticeable affect on style and expressiveness of motion.
- Bones connect and move relative to each other in fairly complex ways. For example, one might model the knee as a simple hinge joint for full-body motion, but the actual range of movement and degrees of freedom are quite a bit more complex.
- Muscles, tendons, and ligaments may attach to bones in various ways which affect their mechanical efficiency for various tasks, e.g., a muscle may have better leverage to move the bone in one direction than another.
- Muscle activation models are based on somewhat simple experiments, e.g., extract a muscle from a frog and run a current through it in the laboratory to see how much it contracts.
- The Hill muscle model is often used in biomechanics as an approximate muscle model of suitable accuracy. It includes spring and damper elements in parallel and in series with an active force element.
- Tendons, muscles, and ligaments can tear when overly stressed. It is theorized that avoiding injury is a significant factor in some animals' preferred movements.
- Fingers move in concert. Try to flex your index finger all the way without moving any other finger.

7 Building controllers

- Mapping from state to actions: $\boldsymbol{\tau} = f(\mathbf{q}, \dot{\mathbf{q}})$, possibly with persistent state.
- Joint-space control: PD-servo at each DOF:

$$\boldsymbol{\tau} = k_s(q - \bar{q}) + k_d\dot{q} \tag{3}$$

  where k_s is the positional gain (analogous to a spring constant), and k_d is damping coefficient.
- SIMBICON: state machine model, with PD control at each joint, plus balance controller. Hand-tuned; produces robust, stable motions, but lacking several key features of human motion (e.g., long strides).
- It is very difficult to tune joint-space controllers for complex tasks with coordinated motions, where all joints are highly interdependent. How does the activation of my knee affect the style of my gait?

- As a general non-linear mapping, an artificial neural network or basis function representation may be used.
- Commercial robotic walking (ASIMO, zero-moment point control). Very conservative; much less efficient than human motion.
- Online optimization control: at each instance, optimize a local objective function based on task goals (such as COM trajectory, matching mocap, end-effector constraints)
- Composable controllers

8 Optimal control

- Given a high-level reward/cost function: determine controller that minimizes an objective function.
- Movement is subject to random variations of many types, e.g., nervous system noise, random variations in ground contact, unexpected external perturbations, sensory ambiguity, etc. (Without uncertainty, all we need is trajectory optimization).
- Objective function: achieve goals (e.g., don't fall down), minimize effort, or expected value of reward. Analogy to rational choice economics.
- Classical optimal control gives theoretically-optimal control, but makes unrealistically restrictive assumptions (e.g., linear dynamics).
- Machine learning approaches to control try to learn from experience, but usually lack theoretical guarantees
- Reinforcement learning is machine learning in which agents attempt to update controllers from experience, including methods such as Q-learning and value iteration. While guaranteed to work for any system, in practice they have a reputation for being impractical for systems with non-trivial dimensionality. (The term "curse of dimensionality" was coined by Bellman, an economist, to describe the difficulty of modeling value functions in high-dimensional action or state spaces).
- Monte Carlo policy evaluation: given a stochastic system and a reward function, average many simulations to approximate expected reward.
- LQR (linear-quadratic regulator): optimal for linear systems with quadratic reward functions. DDP (differential dynamic programming) uses LQR based on linearization around an optimal trajectory.