Realistic human body movement for emotional expressiveness

SIGGRAPH 2009 Course

Friday, 7 August
8:30 AM - 12:15 PM
Auditorium B

Instructors:

Aaron Hertzmann,
University of Toronto
hertzman@dgp.toronto.edu

Carol O'Sullivan,
Trinity College Dublin
Carol.OSullivan@cs.tcd.ie

Ken Perlin,
New York University
perlin@courant.nyu.edu

Course Description:
Humans express their emotions in many ways, in particular through face, eye and body motion. Therefore, the creators of virtual humans strive to convincingly depict emotional movements using a variety of methods. In this course, we focus on the use of realistic human body motion in the service of emotional expressiveness. Applications and research relating to procedural animation of humans with emotion and personality, biomechanical and physical principles of animation, physics-based human motion simulation, and data-driven animation will be reviewed. We will also provide some insights from the field of psychology and discuss issues relating to the perception and evaluation of realistic human body animation.

Prerequisites:
Basic knowledge of human animation.

Syllabus:
Topics covered:

Perlin: Hand-crafted procedural animation:
- Goals: "acting" versus "simulation"
- Advantages/disadvantages of hand-crafted procedural animation
- Path planning
- Walking / Running
- Expressivity of the torso
- Faces and heads
- Arm and hand gestures
- Interaction between characters and with objects in the world
- Connections with A/I, scripting, storytelling
- Challenges for the future

Hertzmann: Biomechanical and physical principles of human animation:
- Kinematic animation
- Biomechanical principles of locomotion
- Efficiency of walking
- Passive-based walking
- Learning objective functions
- Real-time control

O'Sullivan: Expressive Human Motion: Evaluation and Perception:
- Perception of biological motion – an overview
- Emotional body language
- Evaluating the effect of body representation
- Evaluating emotional body motion
- Emotional crowds and multisensory cues

Speaker Bios

Aaron Hertzmann is an Associate Professor of Computer Science at University of Toronto. He received a BA in Computer Science and Art & Art History from Rice University in 1996, and an MS and PhD in Computer Science from New York University in 1998 and 2001, respectively. In the past, he has worked at University of Washington, Microsoft Research, Mitsubishi Electric Research Lab, Interval Research Corporation and NEC Research Institute. His awards include the MIT TR100 (2004), an Ontario Early Researcher Award (2005), a Sloan Foundation Fellowship (2006), a Microsoft New Faculty Fellowship (2006), and a University of Toronto teaching award (2008). His research interests include computer vision, computer graphics, and machine learning.

Carol O'Sullivan is an Associate Professor at Trinity College Dublin. Her research interests include perception, animation, virtual humans and crowds. She has been a member of many IPCs, including the Eurographics and SIGGRAPH papers committee and has published over 100 papers in graphics, especially animation and perception. She is the programme co-chair of the SIGGRAPH Symposium on Applied Perception in Graphics and Visualization 2009, co-Editor in Chief of ACM Transactions on Applied Perception and an editorial board member of IEEE Computer Graphics & Applications. Amongst other conferences, she co-chaired Eurographics'05 in Dublin and SCA’06: Symposium on Computer Animation in Vienna.

Ken Perlin is a professor in the Department of Computer Science at New York University, directs the NYU Games For Learning Institute. He was also founding director of the Media Research Laboratory and director of the NYU Center for Advanced Technology. His research interests include graphics, animation, user interfaces, science education and multimedia. Amongst many honours, he received an Academy Award for Technical Achievement from the Academy of Motion Picture Arts and Sciences for his noise and turbulence procedural texturing techniques, which are widely used in feature films and television, as well as the 2008 ACM/SIGGRAPH Computer Graphics Achievement Award.

Schedule:

8:30 am	Introduction / Overview [All]
8:45 am	Hand-crafted procedural animation [Perlin]
9:45 am	Biomechanical and physical principles of human animation [Hertzmann]
10:15 am	Break
10:30 am	Biomechanical and physical principles - cont.. [Hertzmann]
11:00am	Evaluating expressive human motion, and perceptual issues
12:00 am	Conclusions/Discussion [All]
12:15 am	Close

Representative Bibliography:

Active Learning for Real-time Motion Controllers, Cooper, S. Hertzmann, A. Popović, Z. *ACM Transactions on Graphics* 26(3) (SIGGRAPH 2007)

Learning Physics-Based Motion Style with Nonlinear Inverse Optimization, C. K. Liu, A. Hertzmann, Z. Popović. *ACM Trans. on Graphics* (Proc. SIGGRAPH 2005)

Style-Based Inverse Kinematics, K. Grochow, S. L. Martin, A. Hertzmann, Z. Popović. *ACM Trans. on Graphics* (Proc. SIGGRAPH 2004)

Real Time Responsive Animation with Personality, K Perlin - *IEEE Trans on Visualization and Comp. Graphics*, 1995

Improv: a system for scripting interactive actors in virtual worlds, K.Perlin and A. Goldberg. SIGGRAPH 1996

Building Virtual Actors Who Can Really Act, K. Perlin. *Virtual Storytelling*, LNCS, 2003.

Better acting in computer games: the use of procedural methods. K. Perlin, *Computers and Graphics* 26(1), 2002

Autonomous Digital Actors, K. Perlin and G. Seidman. *Motion in Games*, 2008

Clone Attack! Perception of Crowd Variety. Rachel McDonnell, Micheal Larkin, Simon Dobbyn, Steven Collins and Carol O'Sullivan, *ACM Trans. on Graphics* (SIGGRAPH 2008), 27, (3) (2008)

Evaluating the Effect of Motion and Body Shape on the Perceived Sex of Virtual Characters. R. McDonnell, S. Joerg, J. K. Hodgins, F. Newell, and C. O'Sullivan, *ACM Trans. on Applied Perception*, 5, (4) (2008)

Evaluating the emotional content of human motions on real and virtual characters. R. McDonnell, S. Joerg, J. McHugh, F. Newell and C.O'Sullivan, *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization (APGV'08)*, pp67-74 (2008)

Smooth Movers: Perceptually Guided Human Motion Simulation. R. McDonnell, F. Newell and C. O'Sullivan. *Eurographics/ACM SIGGRAPH Symposium on Computer Animation (SCA'07)*. pp259-269, (2007)

Hand-Crafted Procedural Animation

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

Ken Perlin

Quick Overview of Approaches to Procedural Animation

- Data driven / learning based
- Dynamics / physics based
- Hand-crafted

Simple Example / Demo

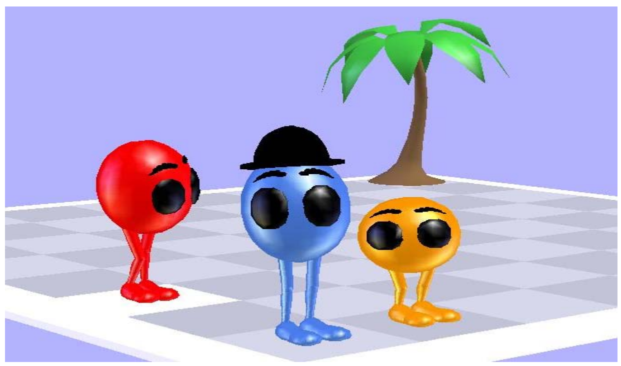

Goals: "Acting" versus "Simulation"

- From a high level user's view
- From an implementation view

Roles of Responsibility

- Different people have different skills

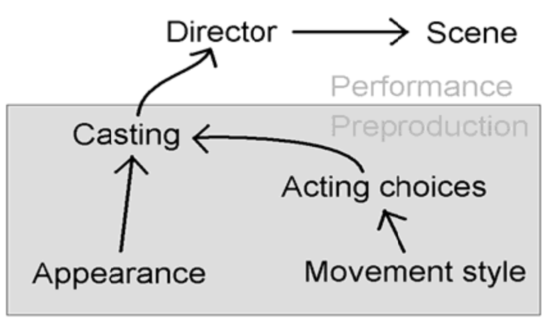

Walking / Running

- Feet
  - User-level parametric control
  - Mechanisms of implementation
  - Physical-world constraints
- Effect on legs, pelvis, upper body, arms

Expressivity of the Torso

- The "line" of the body
- Location of the "acting center"
- Sitting and standing
- Mapping Same expression to different Morphologies

Leaning

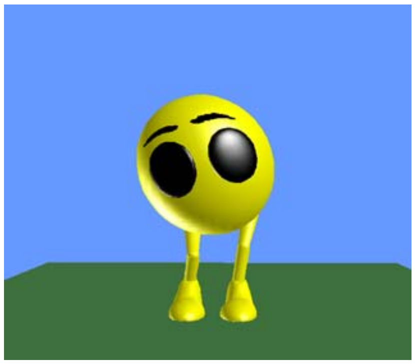
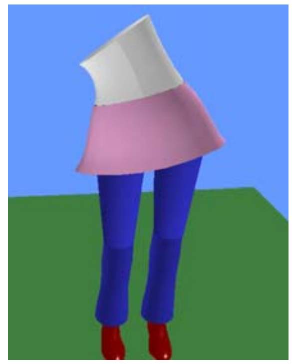

Looking

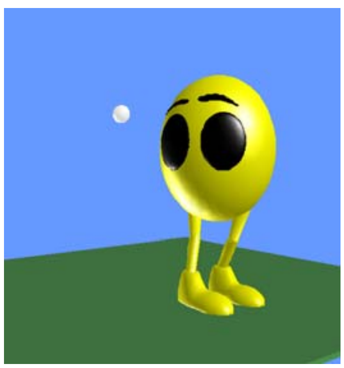


Faces and Heads

- High level expressivity
- Salient parameters - coarse to fine
- Implementation

Arm and Hand Gestures

- Expressivity
- Grasping and holding

Hand Articulation

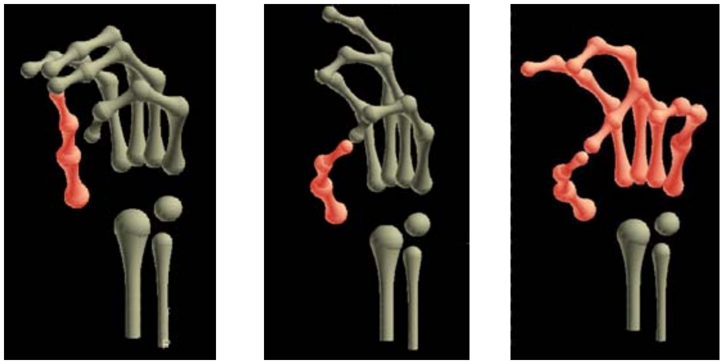

Grasps

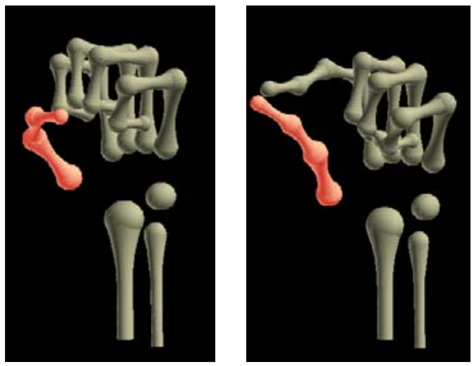

Path Planning

- Alternate approaches
- Algorithmic costs
- Obstacle avoidance
- Dynamic scenes/objects

Foot Steps

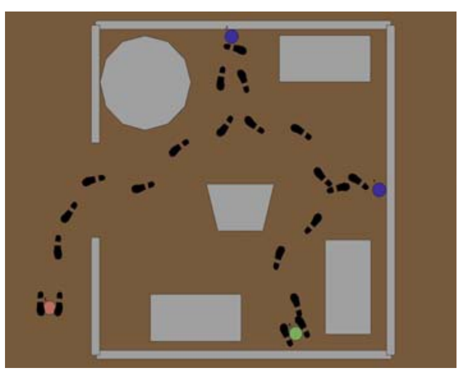

Low Level Personality

- Body Style Parameter settings

High Level Personality

- Acting Choices

Director's View

- Interaction Between characters
- Blocking
- Interaction with objects in the world
- Connections with A/I, scripting, storytelling

Emotive Direction over Time

Mixer for a single shader, i.e. "Happy"

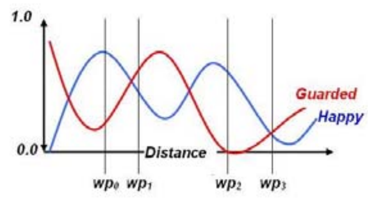

Mixer for multiple shaders

Advantages/Disadvantages of Handcrafted Procedural Animation

- Fidelity versus simplicity
- Scott McCloud's realism scale
- Perils of the uncanny valley

Compromised dynamics

- More flexible and controllable
- CPU
  - Less than physical simulation
  - More than pre-created animation
- Memory
  - Orders of magnitude less than animation

Users/Uses

- Procedural Crowds
- Animation/Film
- Simulation
- Games
- Pre-Vis
- Architectural Walkthroughs
- Live Performance

Challenges for the Future

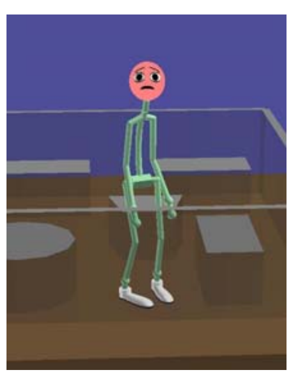

Biomechanical Principles of Motion

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.

Expressive Human Motion:
Evaluation and Perception

Carol O'Sullivan

  

We are all becoming increasingly familiar with, and accepting of computer generated characters. Yet, even in this highly-developed application domain, some suggest that virtual humans have fallen into an "uncanny valley", which was first described as a negative bias for robots when they become extremely close to emulating humans in motivation, characteristics or behaviour [Chaminade et al., 2007; Mori, 1970].

The realism of virtual humans is clearly dependent to a significant degree on their physical appearance, i.e., the quality of the modelling and rendering of their bodies. However, it is in the synthesis of their motions, such as walking, running, smiling, talking, that the biggest research challenges remain. If this is true even for the creation of movies – where significant resources in terms of skilled animators, computational power and motion databases can be allocated to the simulation of realistic humans – then the problems are significantly magnified for human animation in interactive applications running on commodity hardware (such as PCs, game consoles, or even mobile phones). Furthermore, simulating motions that have personality and individual characteristics, without capturing that motion directly from a real human, adds whole new dimensions of complexity. However, it is vital to achieve this goal if we wish to create truly compelling virtual humans with emotional expressiveness.

Chaminade et al. [2007]: The uncanny valley of eeriness, *ACM SIGGRAPH 2007 panels*

Mori [1970]: The valley of eeriness (japanese), *Energy*, 7 (4), 33-35

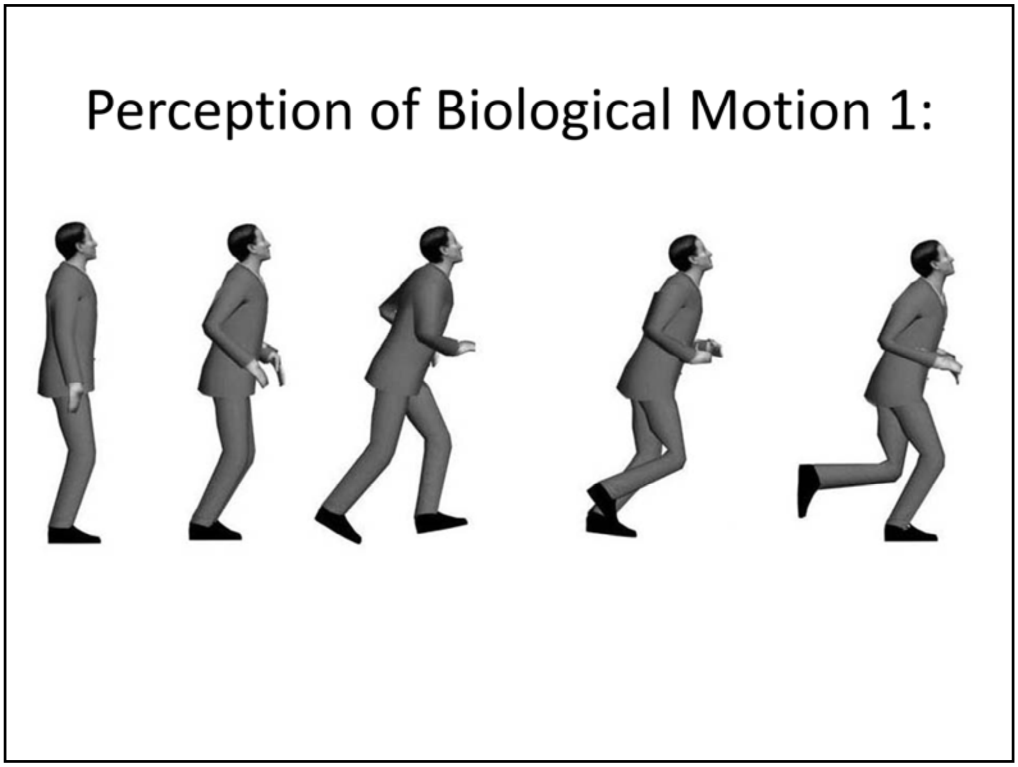

As humans, we are very sensitive to subtle properties of human movements and actions. In fact, the survival of many species is dependent on their ability to recognise complex movements (e.g., of predators, prey and potential mates), as noted by Giese and Poggio. They present a comprehensive review that clearly demonstrates how the brains of humans and other primates have developed to a high level of performance for biological motion recognition. For example, neurons have been discovered that selectively respond to full-body, mouth and hand motion, and facial expressions, while so-called “mirror neurons” have been found that respond both when watching an action *and* performing it.

Giese and Poggio, [2003]: Neural Mechanisms for the Recognition of Biological Movements, *Nature Reviews/Neuroscience,* (4), 179-192

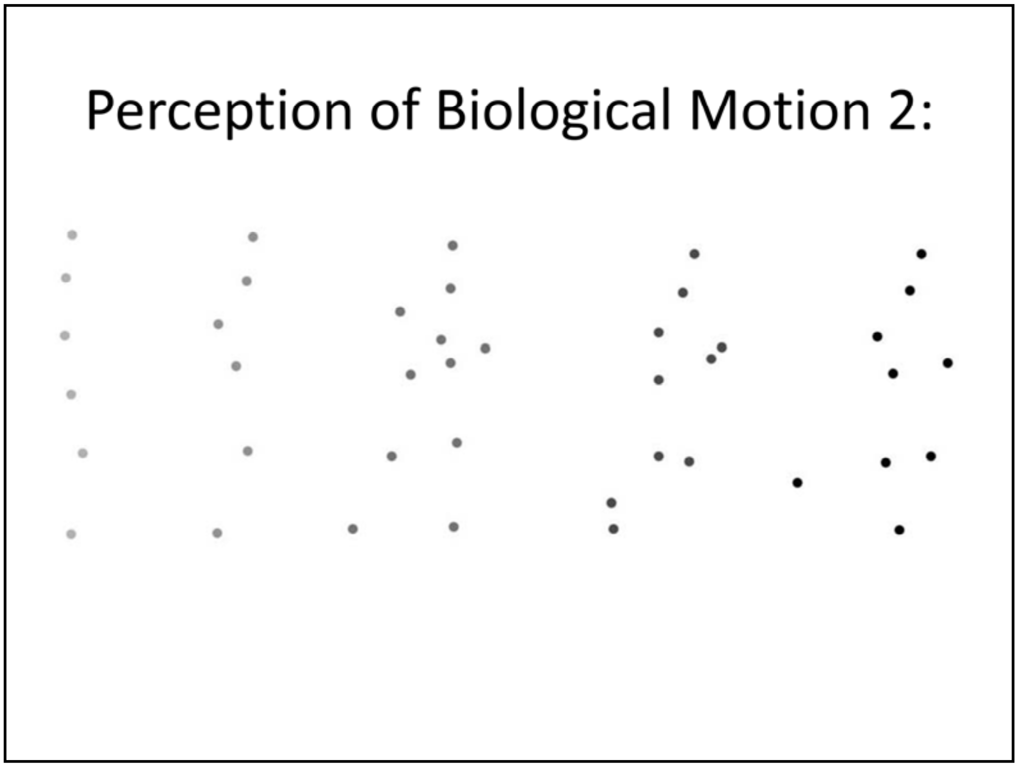

Research in experimental psychology has also demonstrated how finely tuned the human brain is at recognising biological motions and distinguishing between very subtle characteristics of this motion. Johansson [1973,1976] illustrated that natural motion, in the absence of any spatial information, is a sufficient cue to recognise the sex of a walker. His "point-light" displays were designed to separate biological motion information from other sources of information that are normally intermingled with the motion of a human, such as form or outline. He showed that 12 moving light points suffice to create a rich perception of a moving human figure, within a very short space of time (200msec, or five frames of a movie). Many studies since have used point-light displays to show that biological motion perception, an extreme example of sophisticated pattern analysis in the brain, extends even further, e.g., to distinguishing between different movement styles [Pollick et al. 2001] and recognising a *particular* walker [Cutting & Kozlowski 1977] or even one's own walking pattern [Beardsworth & Buckner 1981].

Johansson [1976]: Spatio-temporal differentiation and integration in visual motion perception, *Psychological Research,* 38, 379–393

Johansson [1973]: Visual perception of biological motion and a model for its analysis, *Perception and Psychophysics,* 14, 2, 201–211

Pollick et al. [2001]: Recognizing the style of spatially exaggerated tennis serves", *Perception,* 30, 323-338

Cutting and Kozlowski [1977]: Recognizing friends by their walk: Gait perception without familiarity cues, *Bulletin of the Psychonomic Society*, 9, 5, 353–356

Beardsworth and Buckner [1981]: The ability to recognize oneself from a video recording of ones movements without seeing ones body, *Bulletin of the Psychonomic Society,* 18, 1, 19–22

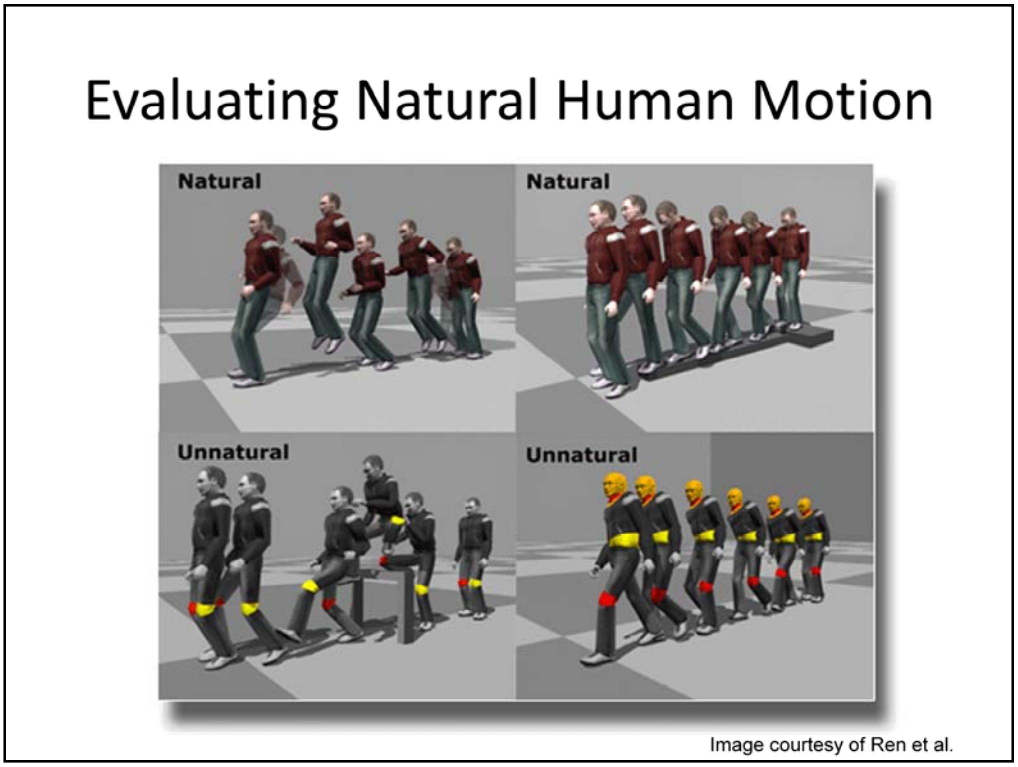

Ren et al. attempted to derive an automatic quantitative metric for the naturalness of human motion by performing a statistical analysis of a large motion capture database. Their metric allowed the particular areas where unnatural motion was present to be pin-pointed.

Ren et al. [2005]: A Data-Driven Approach to Quantifying Natural Human Motion, *ACM Transactions on Graphics (SIGGRAPH 2005),* 24(3), 1090-1097

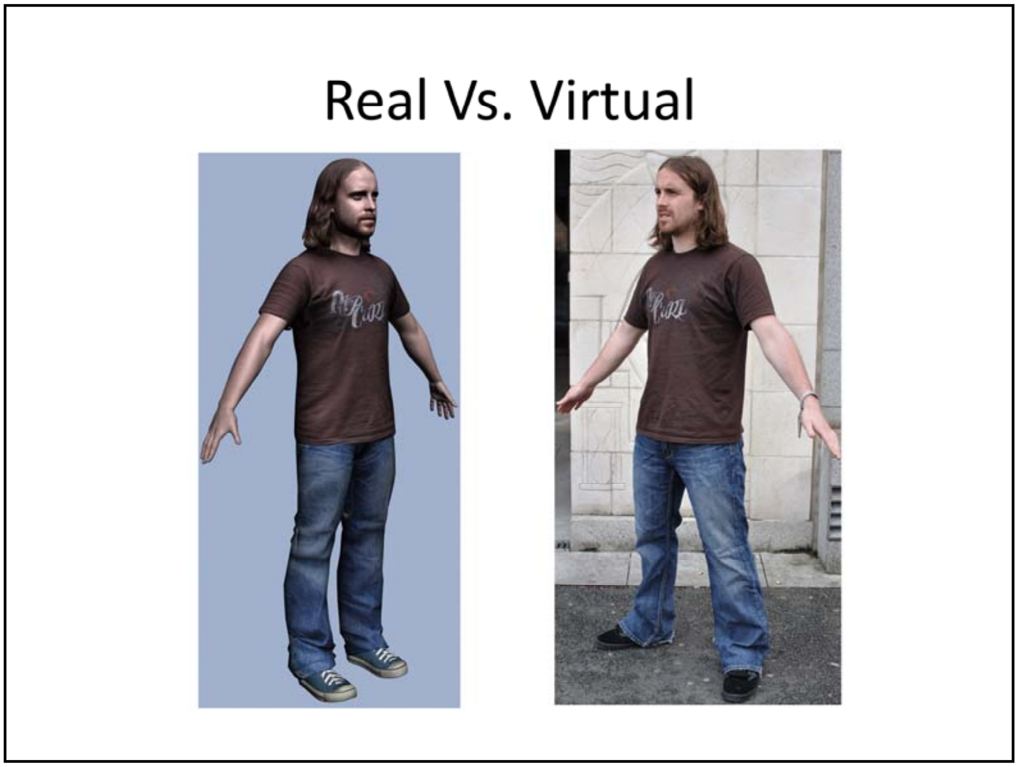

Evidence from the fields of psychology and neuroscience has shown that different neural networks are activated when presented with real and virtual stimuli, e.g., a real hand vs. a virtual reproduction [Perani et al 2001] or cartoons vs. movie clips [Han et al. 2005] . Findings suggest that the human brain functions in a different way when interacting with real people in everyday life than with artificial characters or static social stimuli. When identical biological motion from animated and live-action footage of a movie was displayed to participants, a higher level of activity was found in visual processing areas for real stimuli [Mar et al. 2007]. Other studies have shown that people can respond socially to human and nonhuman entities [Reeves & Naas 1996, Slater & Steed 2002] and that they can engage with virtual humans whether or not they look human [Nowak & Biocca 2003].

Perani et al. [2001]: Different brain correlates for watching real and virtual hand actions. *NeuroImage* 14, 749–758

Han et al. [2005]: Distinct neural substrates for the perception of real and virtual visual worlds. *NeuroImage* 24, 928–935

Mar et al. [2007]: Detecting agency from the biological motion of veridical vs animated agents. *Social Cognitive and Affective Neuroscience Advance Access* 2, 3, 199–205

Nowak and Biocca [2003]: The effect of the agency and anthropomorphism on users' sense of telepresence, copresence, and social presence in virtual environments. *Presence 12, 5, 481–494*

Reeves and Naas [1996]: The media equation: How people treat computers, television, and new media like real people and places. In *Stanford, CA, CSLI Publications*

Slater and Steed [2002]: Meeting people virtually: Experiments in shared virtual environments. In *The social life of avatars: Presence and interaction in shared virtual environments,* 146–171

Emotional Body Language (EBA)

When an emotional body motion is seen, it can unambiguously signal to the viewer not only what emotion the individual is feeling, but also what the reason and resulting actions are (e.g., a terrified body motion of somebody running to escape a predator). De Gelder [2006] presents a thorough survey on the perception of Emotional Body Language and the ways in which emotional whole-body movements are perceived and processed in the human brain. In the EBA literature, the responses to experiments which aim to recognize emotions from body motions are highly consistent [e.g., Wallbott98, Coulson04, Crane07]. It is widely accepted that people can recognize human motion from even a very small set of cues [Johansson73].

Coulson [2004]: Attributing emotion to static body postures: Recognition accuracy, confusions, and viewpoint dependence. *Journal of Nonverbal Behavior,* 28, 2, 117–139

Crane and Gross [2007]: Motion capture and emotion: Affect detection in whole body movement. *Affective Computing and Intelligent Interaction,* 28, 95–101

De Gelder [2006]: Towards the neurobiology of emotional body language. *Nature Reviews Neuroscience,* 7, 242-249

Johansson [1973]: Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics,* 14, 2, 201–211

Wallbott [1998]: Bodily expression of emotion. *European Journal of Social Psychology,* 28, 6, 879–896

Emotional Body Language (EBA)

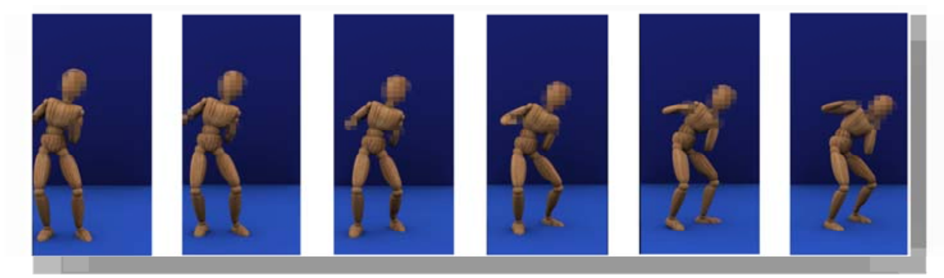

From McDonnell et al. 2009

Atkinson et al. [2004] showed that a small number of cues are enough to recognize every emotion with an above chance rate for point-light animations. They confirmed that emotions can be recognized from body motion even when static form is minimized by use of point-lights. Moreover, they found that the rated intensity of the motion depends more on the motion than on the shape, since point-light motions and real footage were equally rated. Pasch and Poppe [2007] compare the interpretation of static body postures for a realistic virtual character and a mannequin. They show that realism has an effect on the perception of body postures representing the basic emotions, with very different results for each emotion.

Atkinson et al. [2004]: Emotion perception from dynamic and static body expressions in point-light and full-light displays. *Perception,* 33, 6, 717–746

Pasch and Poppe [2007]: Person or puppet? the role of stimulus realism in attributing emotion to static body postures. *Affective Computing and Intelligent Interaction (ACII),* vol. 4738, 83–94

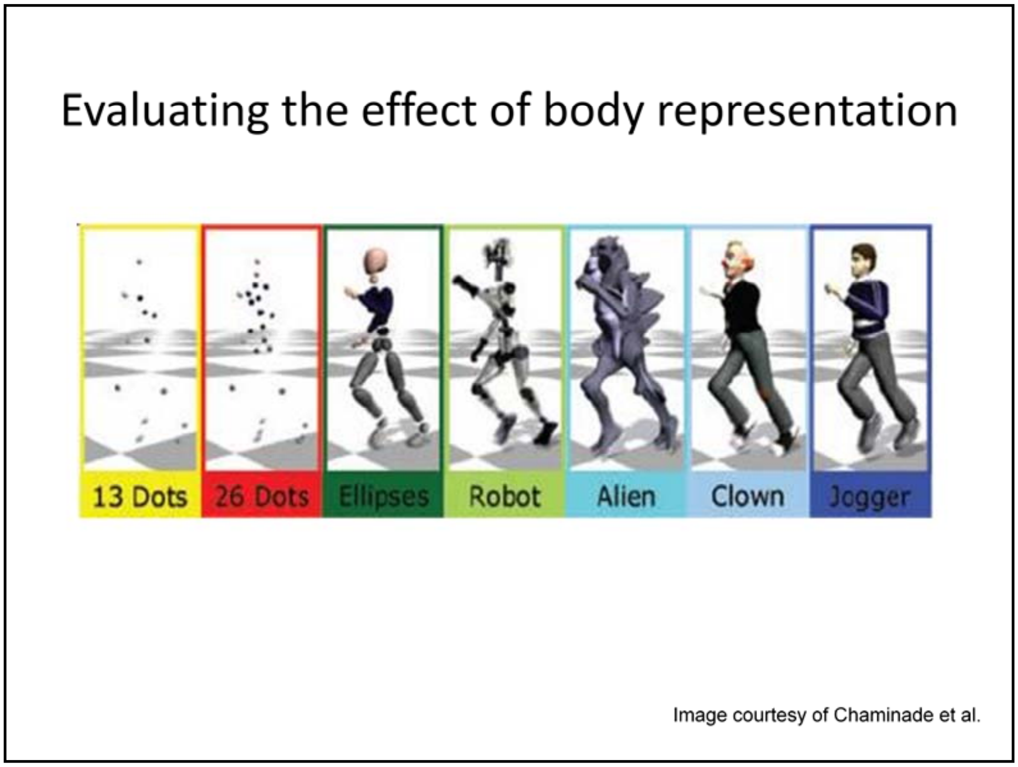

Chaminade et al. [2007] investigated how the appearance of computer animated characters influenced the perception of their actions. They found that the perceived biological nature of a motion decreased with characters' anthropomorphism.

Chaminade et al. [2007] : Anthropomorphism influences perception of computer-animated characters' actions, Social Cognitive and Affective Neuroscience, 2, 206-216

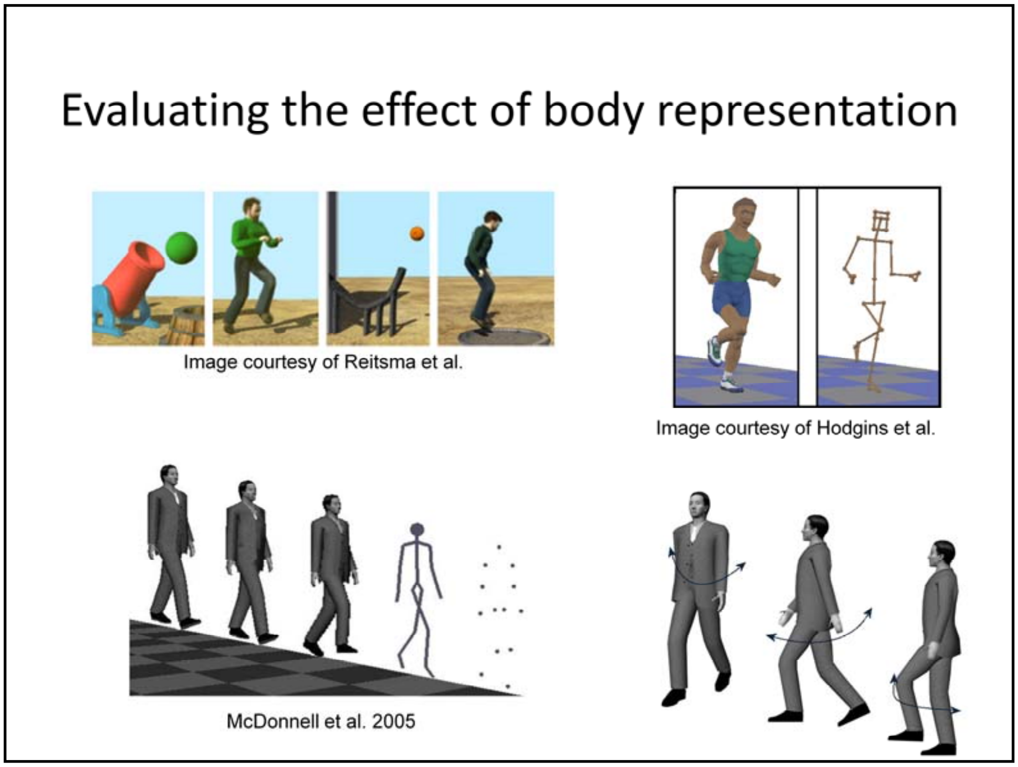

Hodgins et al. [1998] showed that the geometric model type used to represent the human affected people's ability to perceive the difference between two human motions. McDonnell et al. [2005] extended these studies and found that people were most sensitive to differences in human motions for high-resolution geometry and impostor (i.e., image based) representations, less sensitive for low resolution geometry and stick figures, and least sensitive for point-light representations. Finally, Reitsma et al. [2008] demonstrated that participants were more sensitive to errors in motion displayed on human characters than on less detailed human figures or simple geometric objects.

Hodgins et al. [1998]: Perception of Human Motion With Different Geometric Models, IEEE Transactions on Visualization and Computer Graphics, 4(4), 307-316

McDonnell et al. [2005]: LOD Human Representations: A Comparative Study, V-Crowds'05, First International Workshop on Crowd Simulation, 101-115

Reitsma et al. [2008]: Effect of Character Animacy and Preparatory Motion on Perceptual Magnitude of Errors in Ballistic Motion, Computer Graphics Forum (Eurographics), 27(2), 201 - 210

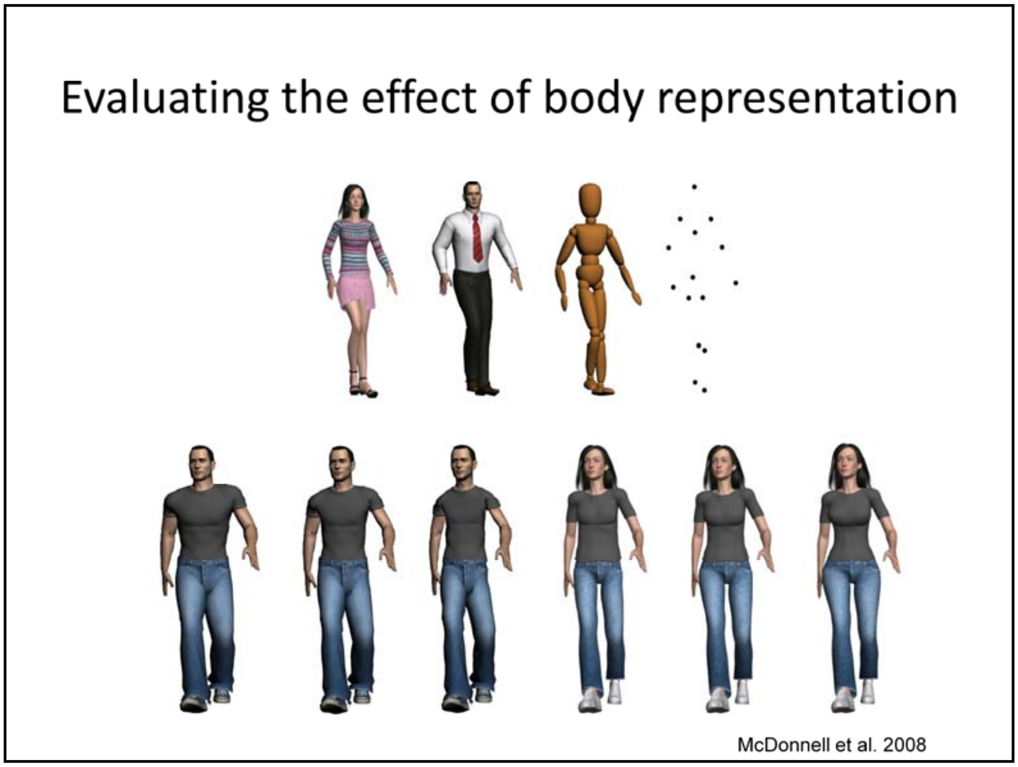

McDonnell et al. [2008] found that both form and motion influence sex perception of virtual characters: for neutral synthetic motions, form determines perceived sex, whereas natural motion affects the perceived sex of both androgynous and realistic forms. A second investigation into the influence of body shape and motion on realistic male and female models showed that adding stereotypical indicators of sex to body shapes influenced sex perception. Exaggerated female body shapes influences sex judgements more than exaggerated male shapes.

McDonnell et al. [2008]: Evaluating the effect of motion and body shape on the perceived sex of virtual characters. *ACM Transactions on Applied Perception*, 5, (4), 21:1 - 21:14

What effect does body type have on perceived emotion?

McDonnell et al. 2009

In order to analyze the emotional content of motions portrayed by different characters, we created real and virtual replicas of an actor exhibiting six basic emotions: sadness, happiness, surprise, fear, anger and disgust. In addition to the video of the real actor, his actions were applied to five virtual body shapes: a low and high resolution virtual counterpart, a cartoon-like character, a wooden mannequin, and a zombie-like character. In a point

light condition, we also tested whether the absence of a body affected the perceived emotion of the movements.

Participants were asked to rate the actions based on a list of 41 more complex emotions. We found that the perception of emotional actions is highly robust and to the most part independent of the character's body so long as form is present. When motion alone is present, emotions were generally perceived as less intense than in the cases where form was present.

McDonnell et al. [2009], Investigating the role of body shape in the perception of emotion, *ACM Transactions on Applied Perception*, 6(3) (In Press)

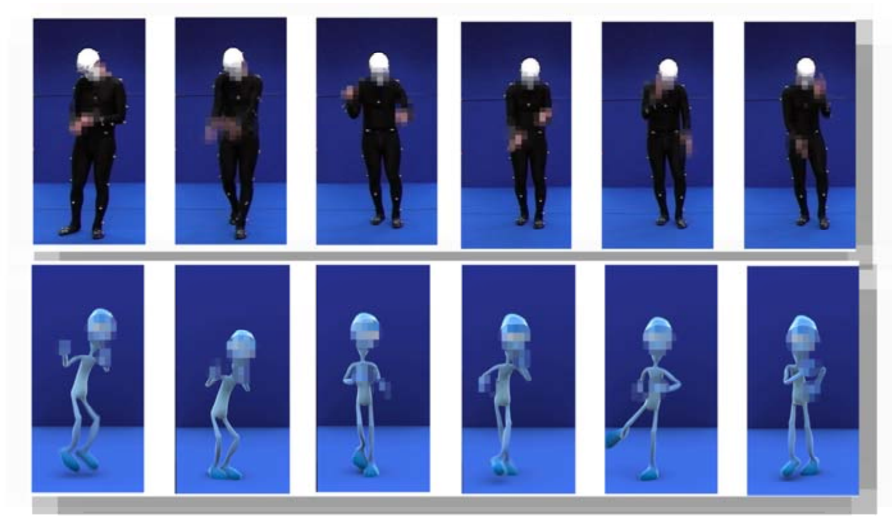

McDonnell et al. 2009

Some examples of the emotional animations used in McDonnell et al. (2009)

Perception of visual emotion in crowds

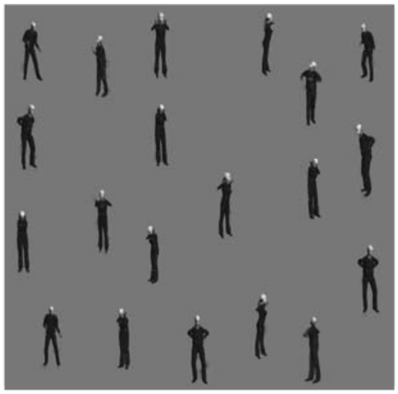

Image courtesy of McHugh & Newell

Participants' task: categorise visually presented emotion

5x3 Within Subjects Design
5 proportions of emotion (0%-100%)
3 sound (majority, minority, absent)

Conclusions:
•Perception of emotion in virtual humans less efficient than in real
•Emotion is recognisable in crowds of virtual humans
•Perception of visual emotion in crowds affected by sound
•Implications for simulating social crowds and groups

McHugh and Newell [2009]: Audition can modulate the visual perception of the emotion of a crowd, *IMRF'09: International Multisensory Research Forum.*

Acknowledgements:

Thanks to the other members of the Metropolis team in Trinity College Dublin, in particular Rachel McDonnell, Micheal Larkin, Peter Loksek, Joanna McHugh and Fiona Newell. Metropolis is funded by Science Foundation Ireland.