

Learning Humanoid Arm Gestures

Bryan Adams

Humanoid Robotics Group
MIT Artificial Intelligence Laboratory
200 Technology Sq. Cambridge, MA 02139
bpadams@ai.mit.edu

Abstract

While biological motion control systems are generally simple and robust, their robotic analogs tend to be just the opposite. While function has driven many of the control architectures to date, we feel that a biologically-inspired system for monitoring the energy consumption of virtual muscles can lead to the development of more humanoid motion and gesture.

Animal and Robotic Motion Control

While animals of all shapes and sizes are able to successfully move their bodies to perform complicated tasks such as running or flying, robotic research in modeling animal movements has progressed slowly. In part, this may stem from the fact that most motion control structures have taken a very functionalist approach. In essence, this approach analyzes the physical system of a given robot and creates a control structure that directly manipulates these physical properties to create the desired movement. While this approach makes a good deal of sense, the difference between robotic motion (involving motors and electricity) and animal motion (often involving muscles and a chemical energy source) has caused the resulting systems to deviate enormously from their biological inspiration.

We have approached this problem looking to the biological system in humans for inspiration and direction. This approach, though, makes the problem twofold: first, how does the animal motion control system work and second, how can we model that system using computational architectures, software, and a robot? Because we are still struggling to understand the biological system, this would appear to needlessly complicate the task. However, we feel that expanding the robotic motion control model will result in more effective and human-like gesture creation.

Functional Focus

Because robotic motion generally focuses on manipulating the environment, the most common type of robotic arm control places an emphasis on the desired state of the end effector (typically a claw or hand). Some simple

trigonometry produces a simple transform that produces joint angles from a desired position and orientation of the hand. By changing the hand position slightly and re-performing the transform, a series of contiguous angles creates a trajectory. This method is not completely without biological inspiration. Cruse (1987) examined gestures in several different planes of motion and postulated that human motor planning is done in part by minimizing the “cost” associated with certain “uncomfortable” joint angles during a trajectory. While this work does not directly say that planning is done on a joint-by-joint basis, it does imply that joint angles are meaningful parameters to monitor and command during gesture formation.

Many robots have successfully employed this idea of using the orientation of the end effector as the basis for gesture formation. For example, the humanoid robot DB in the Kawato Dynamic Brain Project has used this type of control to perform several humanoid tasks ranging from oculomotor responses to more full-bodied tasks such as drumming and dancing. Schaal and Tevatia (1999) explicitly say:

“... [M]ovement planning and learning in task space ... require appropriate coordinate transformation from task to actuator space before motor commands can be computed.”

Their recent work with different computational forms of the Jacobian (Schaal and Vijayakumar 2000) has created a system that is capable of learning fairly general tasks using this transform-based method. Their paper cites the robot’s ability to learn rhythmic and discrete gestures by monitoring position, velocity, and acceleration profiles for each of the joint angles. Through both faster computers and mathematical simplification, the thrust of this approach was to make this method computationally tractable.

While this work is strong from a task-execution standpoint, it is highly unlikely that this method holds a strong analog to the biological motion control structure. The first and biggest difference is that movement comes not from simple actuated joints but from the muscles that attach to the limbs themselves. This distinction is critical because the mapping from muscles to joints is not simple: some muscles span across several joints, thereby causing motion in several joints from a single muscle (or muscle

pair). And while muscles do have sensory organs that provide some notion of position and velocity, biological research has suggested that this feedback is not used directly, but instead contributes to a more complex control system. McMahon (1984) demonstrated how some of the pre-cortex control structures might work, and Hogan's (1990) research on the mechanics of arm movement suggested biologically sensible "spring-like" model for limb movement.

At the Humanoid Robotics Group at MIT, our humanoid Cog employs a simple, spring-based control architecture as the basis for movement. While sensors in the arm joints sense torque and position, these values are used as feedback to a simple linear spring law. Under this system, joint angles are not specified directly, but instead are the result of the parameters of the software spring (equilibrium position, stiffness, and damping) and properties of the environment and limb (gravity, inertia, end load). Using one simulated spring at each joint, Williamson was able to implement some simple ballistic gestures using a postural primitives model (1996) as well as rhythmic gestures using a simple neural oscillator (1999).

While the system was able to successfully learn to reach for a visual target with some accuracy (Marjanovic, Scassellati, and Williamson, 1996), this simple spring law system has some limitations. Although the system creates biologically inspired movement, the sensory information from the arm is far less complex than the feedback provided by human muscles. Additionally, the robot's "muscles" have no memory of the past; the motion of the arm at any given time is determined entirely by its state at that instant. These limitations are particularly debilitating when the robot is attempting to learn novel, humanoid gestures.

Biological Modeling Focus

In response to these shortcomings, we have initiated work expanding the underlying control architecture in hopes of creating a more humanoid learning environment for the robot. The first step has been to broaden our basic muscle model to include a model for energy consumption.

The biochemistry of how muscles turn chemical energy into movement is understood and generally considered a closed question. However, despite thorough knowledge in this area, power consumption issues in robotic arm control are generally ignored. This is because, obviously, motor energy consumption is not a critical engineering issue. Either the robot is tethered and therefore afforded an unlimited energy supply or, if the energy supply is limited, consumption is dominated by factors other than action selection, such as mechanical efficiency or engineering design choice. However, a lack of functional impact in the robotic world does not mean that feedback about energy consumption is not important to the process of movement organization and action selection.

In Adams (2000), the argument is made that a model of energy consumption will help robotic learning in two direct ways. First, a model of energy consumption allows for greater equivalency between the robot and those interacting

with it. This is particularly important for imitation-based learning. If a human attempts to perform a task, but is unsuccessful due to energy constraints (such as attempting to hold a heavy object in an awkward position), the robot must have a concept of why the task was unsuccessful for the human. Without some basis for understanding the "cost" associated with exertion, the robot is unable to differentiate an intentional act from a conspecific's failure due to fatigue.

Second, limitations imposed by this model will help the robot develop along human lines. This is important in instances where typical robotic ability is more functional than that of a human. By either providing or failing to curb superhuman abilities, we run the risk of failing our research goals by allowing the robot to learn human tasks in decidedly non-human ways. For example a camera that can sense variations in temperature would be helpful in locating people in the visual field, yet one is not employed because such a device could very well alter the robot's social development in a fundamental way. In the same way, creating a model of energy metabolism in order to prevent robot's virtual muscles from exerting themselves in a superhuman way provides humanoid boundaries on learning new gestures.

Robotic Implementation

We have implemented such a system, called *meso*, on Cog. Using our message-passing architecture (called *sok*), this system simulates the behavior of the major organs involved in energy production. Energy is "expended" by monitoring the torque values sent to the motors; as more torque is commanded, the metabolic system draws more "energy" from the various organs.

The system, in its current implementation, provides a small set of variables representing chemical levels at different points in the human energy metabolism. Local energy stores in each of the muscles (or *glycogen*), general energy supply in the bloodstream (or *glucose*), and different longer-term energy stores such as fatty tissue and liver glycogen are all maintained. These variables communicate with each other through simulated chemical messengers such as insulin, glucagons, and epinephrine. When the robot exerts a force in one of its joints, this causes the local energy store for that muscle to be depleted in proportion with the strength of the exertion. As this fuel source is depleted, a variety of chemical messengers are triggered, causing different energy stores to provide energy to "fuel" the motion.

Because healthy humans never "run out" of energy, this system doesn't typically interfere with the robot's motion. The system does have two major impacts on the behavior of the robot. First, as the robot moves, the rise and fall of the different chemical levels provides the robot with meaningful feedback about the nature of the gesture. If the robot is required to suddenly exert a high level of force, the metabolic system will react differently from a slower but lengthier motion. By using these cues, the robot can, for example, differentiate one type of gesture over another as

being more “energy efficient”. The second impact the system has is to prevent exertions that would be superhuman in nature. If an exertion causes the short-term energy stores to be depleted beyond their limits, the system intercedes in the motor command system and reduces the force output. This introduces the humanoid limits on exertion that encourage proper learning.

Any implementation of the metabolic system must deal with the issue of complexity. While the metabolic system is well understood at the chemical-reaction level, the interplay between each of the reactions, if modeled explicitly, would create a system with unacceptable complexity. In fact, for this application, the requirements are even more stringent: the model must operate on a robot in real-time; hence the complexity of the model must allow the system to create the proper feedback on the proper timescale.

However, creating a model with a reasonable level of complexity can be achieved given our fairly modest goals. The current system only recreates two aspects of human metabolism. First, it provides the robot with humanoid behavioral limits by placing appropriate restrictions on the nearly unlimited power that the robot is capable of exerting. Second, the model creates the accompanying metabolic experience that goes along with testing these limits. Without entering into the debate of whether a robot actually has “experiences”, it is enough to say that this system provides a set of metabolic variables that correspond to the robot’s actions and provide information regarding the level of exertion. *Meso* accomplishes these goals by treating the metabolism as a simple control mechanism. While the values of the variables in the model are not meant to represent actual values found in humans (these vary too greatly for specific values to be useful), they do change in proportion with the human reactions. The levels of the variables in the various processes are then applied to the electro-mechanical system to achieve the second goal.

The time scale for these processes is also significant. In the current environment, interactions with Cog are, from a metabolic standpoint, short (i.e., less than one hour). While a small set of individuals do interact with the robot over a period of weeks, months, and years, the robot has not yet been designed to accumulate any sort of state over a period longer than an hour. As such, this implementation of *meso* focuses on modeling metabolic effects that happen over roughly an hour. Because developing these longer-term effects could prove beneficial in the future, some consideration is paid to allowing for future development in this direction. Many of the long-term metabolic effects can be represented in this model by dynamically (but slowly) changing the coefficients of the reactions established in *meso*. Other long-term metabolic changes could include the results of a trauma: long periods without nutrients or an unbalanced diet. But because the overall fitness level rarely (if ever) changes in perceptible ways over short-term interactions, neglecting this part of the model will not change the nature of the short-term behaviors.

Also, from a biological standpoint, any model of the metabolic system must recognize that each person’s metabolism is entirely different. While the basic chemical reactions are the same in all people, the higher level relationships (for example, the amount that the heart rate goes up for a given amount of work) vary greatly not only from person to person, but vary for a given person over the course of his life. Factors such as genetic makeup, environmental quality, and general fitness level change the relationships greatly, in some cases by an order of magnitude. As such, there is no single “right” behavior, but instead a range of values that the system can emulate. In *meso*, the right set of parameters and associations are established, but the specific relations can be manipulated. With that basic framework, future work can model the influence of these other more distant (and often longer-term) factors.

Finally, the *meso* system creates a “virtual” metabolic state for the robot, but stops short of providing easy emotional or behavioral cues. Like the metabolic state in humans, the sensing of the chemical state of the body is vague and poorly understood. These senses do not result in concrete thoughts, but instead are thought to create a feeling that may or may not be acted upon by a higher level of control. While *meso* provides parameters that correspond to nebulous feelings such as “tiredness”, the correct use of this variable to create humanoid behavior is an open question.

Acknowledgements

This work was supported by ONR and DARPA under MURI N00014-95-1-06000 and by DARPA under contract DABT 63-99-1-0012.

References

- Adams, B. 2000. *Meso: A virtual musculature for humanoid robots*. M.Eng Thesis, Department of Electrical Engineering and Computer Science, MIT.
- Cruse, H., Brüwer, M. 1987. The human arm as a redundant manipulator: the control of path and joint angles. *Biological Cybernetics*, 57:137-144.
- Hogan, N. 1990. Mechanical Impedance of Single- and Multi-Articular Systems. Winters, J. M. and Woo, S. L.-Y., eds. *Multiple Muscle Systems: Biomechanics and Movement Organization*, Springer-Verlag.
- Marjanovic, M., Scassellati, B. and Williamson, M. 1996. Self-Taught Visually-Guided Pointing for a Humanoid Robot. In Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB-96), Cape Cod, MA.

McMahon, T. 1984. *Muscles, Reflexes, and Locomotion*. Princeton University Press, Princeton, NJ.

Schaal, S., Vijayakumar, S. 2000. Fast and efficient incremental learning for high-dimensional movement systems. In Proceedings of the International Conference on Robotics and Automation (ICRA2000). San Francisco, CA.

Schaal, S., Tevatia, G. 1999. Inverse kinematics for humanoid robots. In Proceedings of IEEE International Conference on Computational Intelligence in Robotics and Automation (CIRA '99).

Williamson, M. 1996. Postural Primitives: Interactive Behavior for a Humanoid Robot Arm. In Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB-96), Cape Cod, MA.

Williamson, M. 1999. Robot Arm Control Exploiting Natural Dynamics. Ph.D. diss., Department of Electrical Engineering and Computer Science, M.I.T.