Discrete Submovments Using Predictive Models

Colin W. Barringer & Andrew G. Barto
Department of Computer Science, University of Massachusetts Amherst
AUTONOMOUS LEARNING LABORATORY
http://www-all.cs.umass.edu/

Abstract

We provide a mechanism for implementing the discrete submovement policy proposed in Fishbach, et al.[9] By using predictive models of movement endpoint, the agent is able to issue discrete, corrective movements to achieve target endpoints. Simulation studies comparing multiple discrete decision strategies show that the policy ferived for primate data by Fishbach, et al. maximizes speed and accuuracy of simulated reaches.

Introduction

As first observed by Woodworth,[32] the tangential velocity trace of a primate's hand in a reaching task often contains multiple peaks. This phenomena is often attributed to the recruitment of multiple, overlapping movement primitives—known as submovements or movement units—by the motor control system.[1,3,5,6,7,9,12,14,21,23,25,26,27,28,32] However, recent continuous control theories suggest that apparent underlying movement units are the result of delays and non-linearity in the controlled system[2,8,11,15,18,30] or the result of a perceptual deadzone.[13,24,31]

Fishbach, et al.[9] proposed a discrete reaching strategy based on endpoint prediction. They showed a correlation between the onset time of the secondary submovement and the normalized amplitude of the primary submovement. This relationship suggests that corrections are made in response to some prediction of primary submovement endpoint, and Fishbach, et al. presented a simple prediction model based on key kinematic variables. Although decisions are made on a continuous basis, changes to the motor command are discrete. That is, the control signal is only altered under certain conditions, though the feedback is monitored continuously. In this way, it is similar to a perceptual deadzone, though not exactly the same.

We extend the work of Fishbach, et al. by providing a mechansitic implementation of this scheme and testing it in a simulated motor system. We evaluated a number of different decision rules based on predictive models of endpoint position, and have found that the rule derived from monkey behavior by Fishbach, et al. provides good speed and accuracy results in comparison to other rules.

Discrete Controller

We propose a model for the initiation of submovements according to endpoint predictions. While Fishbach, et al.[9] described an explanatory model of primate reaching behavior, this model is generative. The model is presented in the accompanying figure.

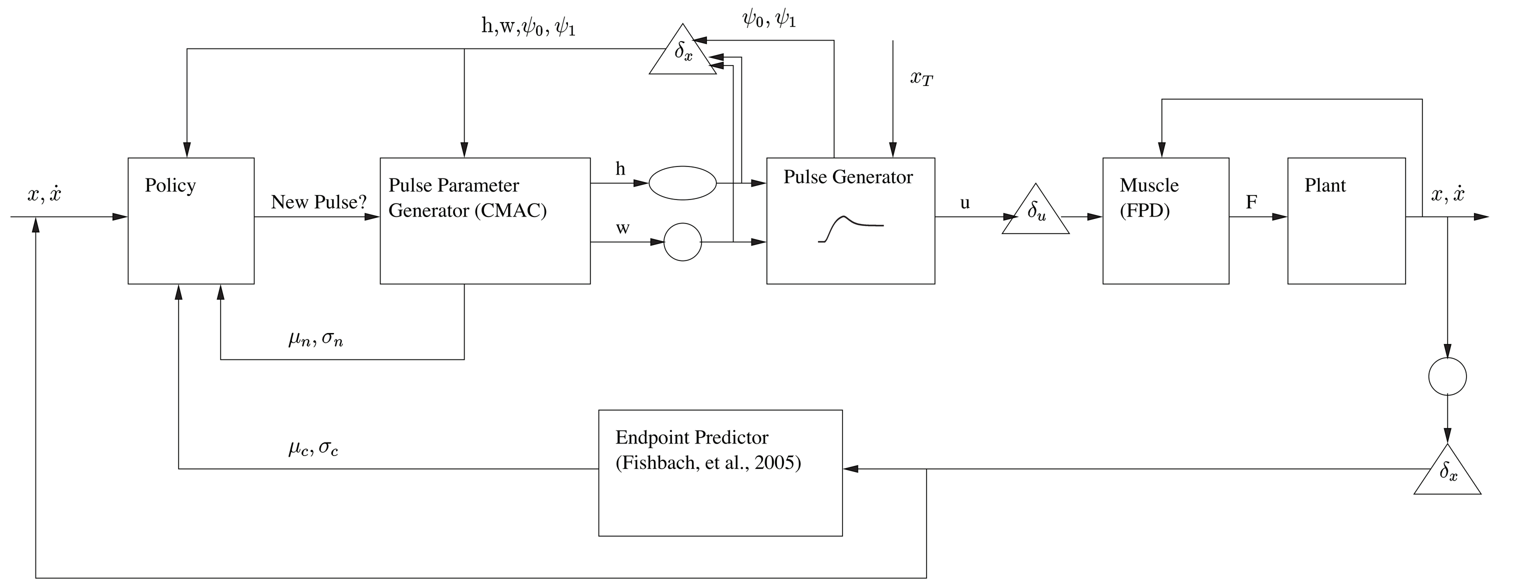

Motor System

The motor plant is a one-dimensional point mass, with position defined as x. It is manipulated by "muscles" that generate a force applied to the point. In initial experiments, we modeled a muscle as a proportional-derivative controller acting upon the point mass by setting an equilibrium point u.

Unfortunately, this simple controller does not exhibit the endpoint variance that Fishbach, et al. observed in their data for experiments with monkeys. The endpoint of movement is uniquely defined as the final value of u. We therefore chose the somewhat more complicated fractional power damping (FPD) model, which was motivated by primate muscles and spinal reflexes,[16,20] but remains relatively simple. Under this model, the force applied to the plant by the muscles is:

$$F = K_p(u - x) + K_v\dot{x}^{\frac{1}{5}}$$

Using the FPD model of force generation results in a region of state space, centered around x=u where the spring forces and the damping forces counteract one another in such a way that the forces applied to the plant become negligible. This region is known as the **stiction region**. The endpoint of movement under this model is defined as the point at which the plant enters the stiction region, and this depends not only on the final value of u, but the entire trajectory.
Additionally, consistent with physiological[4,22] and behavioral[10,14,17,23,28,29,30] evidence, our motor system is subject to both control and feedback noise and delay. We have also filtered the control signal to produce a differentiable acceleration trace. A differentiable acceleration trace is more realistic and necessary for the application of prediction techniques developed by Fishbach, et al.[9]

Pulse-Step Control

Under the FPD framework, it is possible to execute fast, accurate movements by issuing pulse-step commands. A pulse-step command specifies a pulse magnitude (h), a pulse width (w), and a step magnitude (s). The pulse-step is executed by holding u steady at h for a duration of w, and then holding it steady at s for the remainder of the movement. A pulse-step command allows for fast movements by rapidly accelerating the plant at the early stage of a movement, such that it will enter the stiction region at a point close to the target. The accompying figure illustrates a sample pulse-step command executed under the described motor system.

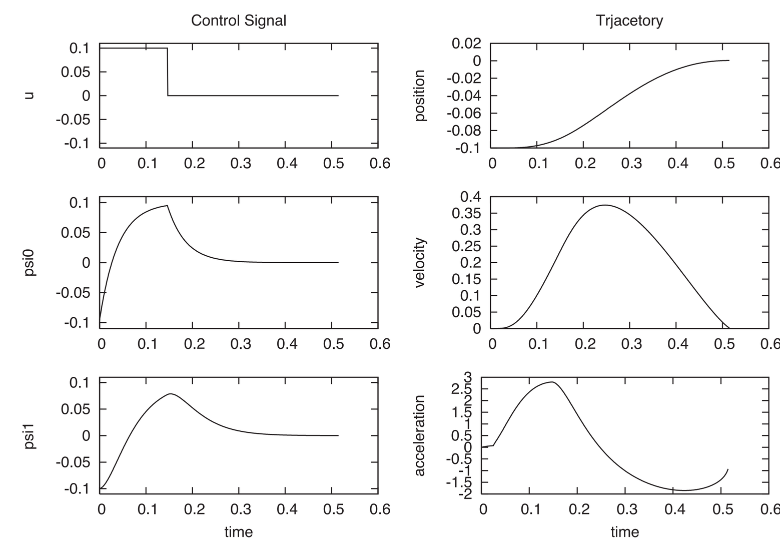

Endpoint Prediction

The control scheme relies on predictions of endpoint position. Predictions are made for the current executing pulse using methods developed by Fishbach, et al.[9] Fishbach. et al. identified four key kinematic variables: the first peak snap (s_1), the first peak jerk (j_1), the first peak acceleration (a_1) and the first peak velocity (v_1). They demonstrated that feedback at these key points accurately predicts the mean and variance of endpoint position. We have found that this method of prediction, developed using data collected from reaching monkeys, is also appropriate for this simulated muscle system: the predictions provide similar levels of accuracy, and the accuracy increases as the movement progresses (see accompanying figure). Predictions for a new pulse executing from the current state are made using a cerebellar model architecture controller (CMAC). Though we use CMACs in our model, it is not our intent to suggest that these predictions are made in the cerebellum.

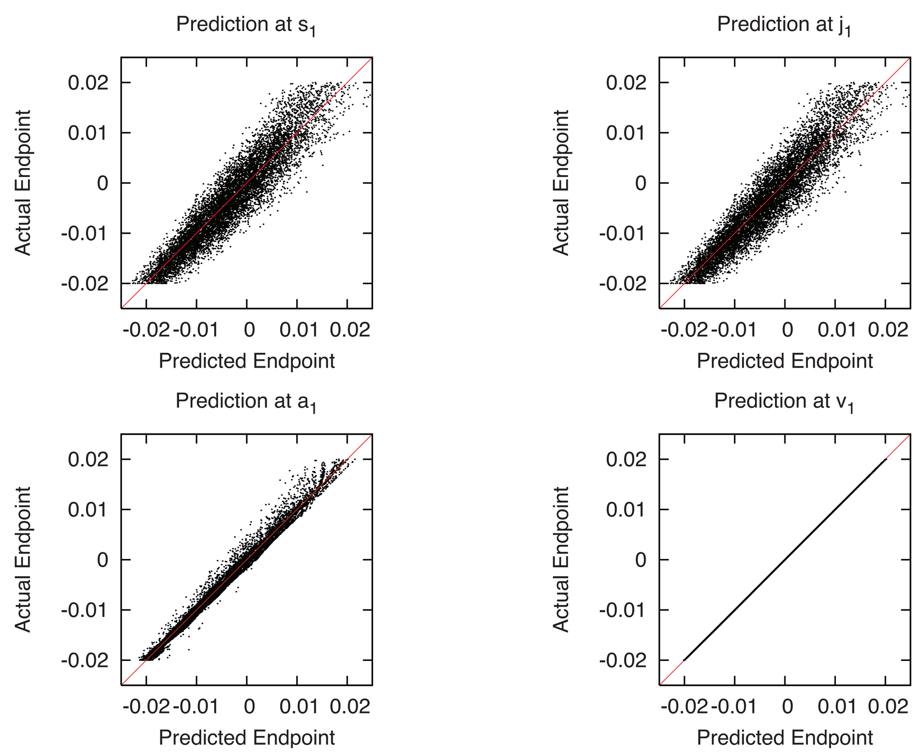

Pulse Initiation Policies

Fishbach, et al. suggest that the primate reacher uses a threshold on the normalized amplitude of the primary submovement divided by the expected variance in this measure as a means of deciding when to initiate a corrective submovement.[9] In addition to this method, we developed a number of different discrete policies that make comparisons between the probability of success with the current pulse and the probability of success with a new pulse.

These policies issue successive pulse-step commands to the described motor system. In each policy we assume that correct parameters for the pulse are stored by the system and not subject to control. The controller need therefore only make decisions regarding *when* to pulse and not *how*. Each controller makes a comparison between the predicted endpoint given the currently executing pulse and the predicted endpoint given a newly issued pulse. The controllers vary only in the manner of this comparison. Because the policies are parameterized, we have used policy gradient methods to optimize each policy with respect to speed and accuracy.

Results

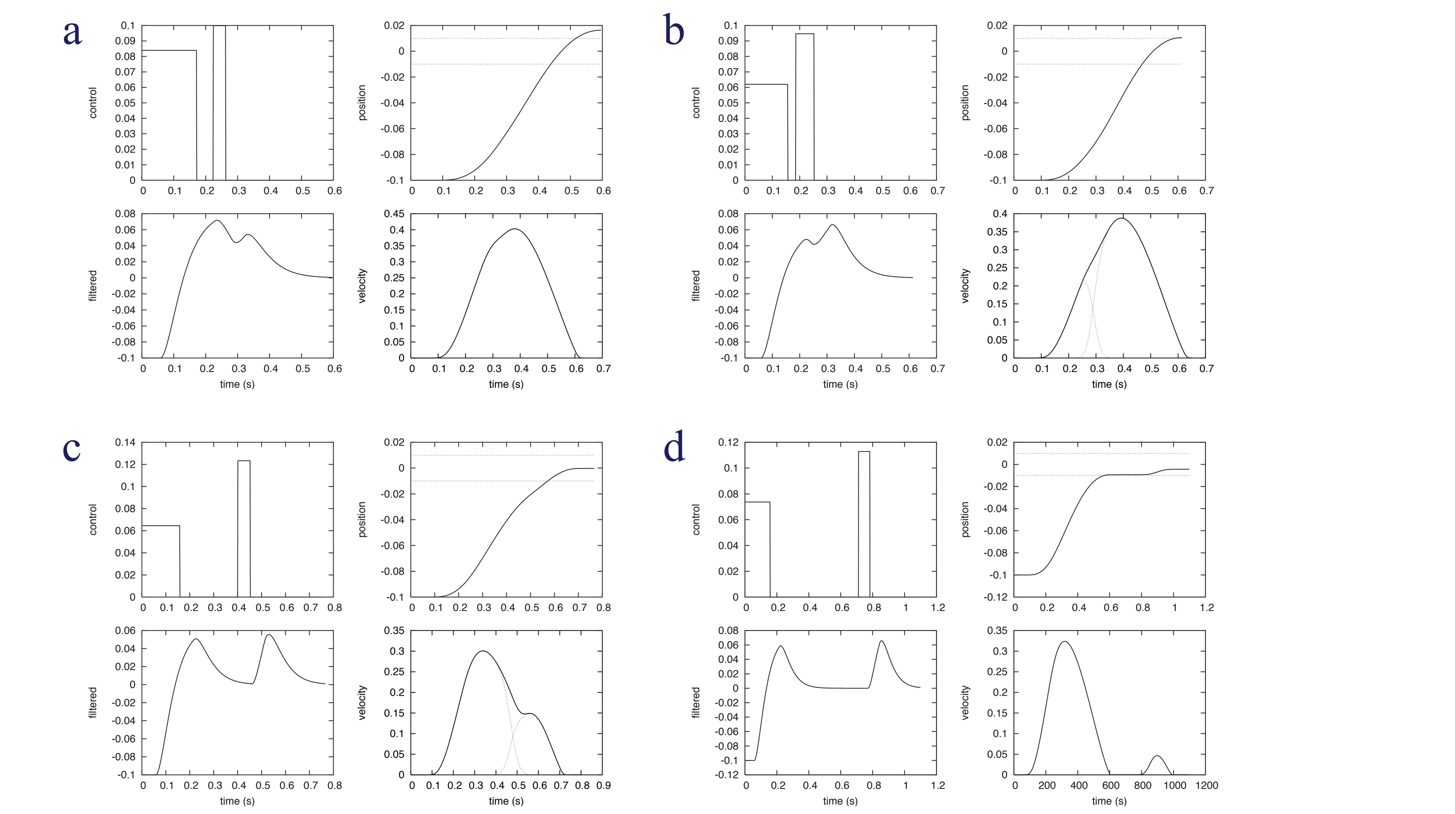

Each discrete control policy, coupled with the FPD muscle system, results in trajectories that are qualitatively similar to primate reaching movements. Illustrated here, we see four different characteristic movements: a) an early correction that fails to properly correct, b) an early correction that successfully corrects , c) a late correction that successfully corrects, and d) a delayed correction.

Because the control signal is subject to multiplicative Gaussian noise, it is often incorrect to issue a correction early in movement; the correction is likely to be just as noisy as the initial pulse. We therefore see that corrections issued later in the movement are more likely to succeed.

Like the experimental data of Fishbach, et al.,[9] the movements generated by this model demonstrate a correlation between primary submovement amplitude and secondary submovement onset time. This was true irrespective of the control policy followed, and likely a result of system dynamics or the nature of the control signal—i.e. discrete pulse-step commands.

Discussion

We presented a simple motor system that produces trajectories qualitatively similar to those seen in primate reaching studies. The controller described here uses predictive models to decide when to issue corrective pulses.

Each of the developed policies that issued corrective movements improved the accuracy of movements, and most of them additionally shortened the duration. Interestingly, the policy that maximized both speed and accuracy was derived from the primate policy hypothesized by Fishbach, et al.[9]

While this does suggest that a control system based on issuing discrete movement primitives is compatible with existing behavioral data, we cannot conclude that a continuous policy could not accomplish this same behavior.

We are currently working on the development of continuous policies under the same motor system. It will be revealing to discover whether or not a continuous policy can accomplish the same result. We expect that the correlation between normalized amplitude and secondary submovement onset time is a feature unique to discrete control policies.

References and Acknowledgements

1. Berthier, N.E. Learning to reach: A mathematical model. Developmental Psychology 32 (1996).
2. Bhushan, N., and Shadmehr, R. Evidence for a forward dynamics model in human adaptive motor control. In *Advances in Neural Processing Systems* (Cambridge, MA, 1999), M.S. Kearns, S.A. Solla, and D.A. Cohn, Eds., vol. 11, MIT Press.
3. Burdet, E. and Milner, T.E. Quantization of human motions and learning of accurate movements. Biological Cybernetics 78, 4 (1998).
4. Clamann, H.P. Statistical analysis of motor unit firing patterns in a human skeletal muscle. *Biophysical Journal 9*, 10 (1969).
5. Craik, K.J.W. Theory of the human operator in control systems. I. The operator as an engineering system. British Journal of Psychology 38 (1947).
6. Crossman, E.R.F.W. and Goodeve, P.J. Feedback control of hand movement and Fitt's law. *Quarterly Journal of Experimental Psychology 35A* (1983).
7. Doeringer, J.A. and Hogan, N. Intermittency in preplanned elbow movements persists in the absence of visual feedback. *Journal of Neurophysiology 80*, 4 (1998).
8. Feldman, A.G. Once more on the equilibrium-point hypothesis (lambda model) for motor control. *Journal of Motor Behavior 18* (1986).
9. Fishback, A., Houk, J.C., Roy, S.A., Bastianen, C. and Miller, L.E. Kinematic properties of on-line error corrections in the monkey. *Experimental Brain Research* (2005).
10. Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology 47*, 6 (1954).
11. Gribble, P.L., Ostry, D.J., Sanguineti, V., and Laboissiére, R. Are complex control signals required for human arm movement? *Jounral of Neurophysiology 79*, 3 (1998).
12. Gross, J., Timmerman, L., Kujala, J., Dirsk, M., Schmitz, F., Salmelin, R. and Schnitzler, A. The neural basis of intermittent motor control in humans. *Proceedings of the National Academy of Sciences of the USA 99*, 4 (2002).
13. Hanneton, S., Berthoz, A., Droulez, J., and Slotline, J.E. Does the brain use sliding variables for the control of movements? *Biological Cybernetics 77* (1997).
14. Harris, C.M. and Wolpert, D.M. Signal-dependent noise determines motor planning. *Nature 394* (1998).
15. Hoff, B., and Arbib, M.A. *Control of arm movement in space: Neurophsiological and computational approaches*. Experimental Brain Research Series. Springer-Verlag, New York, 1991, ch. A model of the effects of spped, accuracy and perturbation on visually guided reaching.
16. Houk, J.C., Fagg, A.H., and Barto, A.G. *Progress in Motor Control:: Structure-Function Relations in Voluntary Movements*, vol. 2. Human Kinetics,2002, ch. Fractional power damping model of joint motion.
17. Jones, K.E., Hamilton, A.F.de C., and Wolpert, D.M. Sources of signaldependent noise during isometric force production. *Journal of Neurophysiology 88*, 3 (2002).
18. Kawato, M. *Attention and Performance*, vol. 14. MIT Press, Cambridge, MA, 1992, ch. Optimization and learning in neural networks for formation and control of coordinated movement.
19. Kohl, N., and Stone, P. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the Institude of Electrical and Electronics Engineers International Conference on Robotics and Automation* (2004), vol. 3.
20. Kositsky, M., and Barto, A.G. Nonlinear damping dynamics and the variability of rapid aimed movements. Tech. Rep. 01-15, University of Massachusetts,Amherst, MA, April 2001.
21. Lee, D., Port, N.L. and Georgopoulos, A.P. Manual interception of moving targets. II. On-line control of overlapping submovements. *Experimental Brain Research 116*, 3 (1997).
22. Matthews, P.B.C. Relationship of firing intervals of human motor units to the trajectory of post-spike after-hyperpolarization and synaptic noise. *Journal of Physiology 492*, 2 (1996).
23. Meyer, D.E., Abrams, RA., Kornblum, S., Wright, C.E. and Smith, J.E.K. Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review 95*, 3 (1988).
24. Miall, R. C., Weir, D.J., and Stein, J.F. Intermittency in human manual tracking tasks. *Journal of Motor Behavior 25*, 1 (1993).
25. Miall, R.C., Weir, D.J. and Stein, J.F. Visuomotor tracking with delayed visual feedback. *Neuroscience 16*, 3 (1985).
26. Neilson, P.D., Neilson, M.D. and O'Dwyer, N.J. Internal models and intermittency: a theoretical account of human tracking behavior. *Biological Cybernetics 58*, 2 (1998).
27. Novak, K.E., Miller, L.E., and Houk, J.C. The use of overlapping submovements in the control of rapid hand movements. *Experimental Brain Research 144* (2002).
28. Schmidt, R.A., Zelaznik, H.N., Hawkins, B., Frank, J.S., and Quinn, Jr., J.T. Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review 86*, 5 (1979).
29. Smith, K.U., and Sussman, H.M. Delayed feedback in steering during learning and transfer of learning. *Journal of Applied Psychology 54*, 4 (1970).
30. Todorov, E. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation 17*, 5 (2005).
31. Wolpert, D.M., Miall, R.C., Winter, J.L., and Stein, J.F. Evidence for an error deadzone in compensatory tracking. *Journal of Motor Behavior 24*, 4 (1992).
32. Woodworth, R.S. The accuracy of voluntary movement. *Psychological Review 3*, 2 (1899).

This research was made possible by NIH grant # NS 044393-04