12-8-1987

Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling

Martin Buehler

Daniel E. Koditschek
University of Pennsylvania, kod@seas.upenn.edu

Recommended Citation

Martin Buehler and Daniel E. Koditschek, "Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling", . December 1987.

Buehler, M. & Koditschek, D. E. (1987). Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling. *26th IEEE Conference on Decision and Control.*

NOTE: At the time of publication, author Daniel Koditschek was affiliated with Yale University. Currently, he is a faculty member at the School of Engineering and Applied Science at the University of Pennsylvania.

Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling

Abstract
We explore a very simple representative of a class of of robotic tasks which require "dynamic dexterity", among them the task of "juggling". In this initial paper we propose a formal definition of a "vertical one juggle", report a few preliminary analytical results, and offer illustrative simulations. This analysis is being currently applied to the design of and experimentation with "juggling algorithms" for a one-degree of freedom "robot" operative in the gravitational field.

Disciplines
Systems and Communications

Comments
Buehler, M. & Koditschek, D. E. (1987). Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling. *26th IEEE Conference on Decision and Control.*

NOTE: At the time of publication, author Daniel Koditschek was affiliated with Yale University. Currently, he is a faculty member at the School of Engineering and Applied Science at the University of Pennsylvania.

Robotics in an Intermittent Dynamical Environment: A Prelude to Juggling

A Talk Delivered at the
26^{th} IEEE Conference on Decision and Control
Los Angeles, California
December 11, 1987

Martin Bühler and Daniel E. Koditschek [1]

Center for Systems Science
Yale University, Department of Electrical Engineering

December 8, 1987

[1]This work is supported in part by the National Science Foundation under grant DMC-8552851, as well as PMI Corporation and Inmos Corporation.

Abstract

We explore a very simple representative of a class of robotic tasks which require "dynamical dexterity", among them the task of juggling. In this initial paper we propose a formal definition of a "vertical one juggle", report a few preliminary analytical results, and offer illustrative simulations. This analysis is being currently applied to the design of and experimentation with "juggling algorithms" for a one-degree of freedom "robot" operating in the gravitational field.

1 Introduction

In this paper we consider a simple representative from a range of robotic tasks associated with dexterous capabilities that might be grouped under the general rubric of "juggling". We understand this term to include those tasks requiring throwing and catching, or (as in this paper) beating and batting, or any other interaction with an object (or multiple objects) which would otherwise fall freely in the earth's gravitational field. Such tasks share the property of presenting non-trivial dynamical environments whose characteristics change intermittently subject to excitation from the robot. It seems fair to say that the only systematic work in this realm to date has been the pioneering research of Raibert [3].

We find it useful to distinguish between the robot and the environment. The former consists of the familiar set of actuated rigid links whose kinematic constraints give rise to a reasonably well understood dynamical relationship between actuator commands and physical forces and movements. The latter consists of all motions and forces of those objects relevant to some abstract task, and, in particular, the *manner in which they may be dynamically coupled to the actuated robot* . Thus, we may speak of an "environmental control problem" in isolation from the particular details of a given robot. Or we may consider the more particular (and generally more difficult) "robotic problem" requiring the translation of abstract goals into robot actuator commands with a guarantee that the resulting closed loop interaction of robot and environment dynamics will achieve those goals. Of course, these notions are quite familiar in the world of robotics, and, in particular have been carefully examined in the work of Hogan [2].

The present discussion concentrates exclusively on the environmental control problem. Specifically, we are interested in developing analytical insight into the environment of the "robotic apparatus" depicted in Figure 4. This is a one degree of freedom revolute arm actuated by a direct drive electric motor, whose axis of rotation is perpendicular to a "frictionless" plane which may be turned into the gravitational field. This arm will engage a "frictionless" puck which slides along the plane as acted upon by gravity and the forces of impact.

In Section 2, we develop a model for the interaction of this primitive juggling robot with the puck. Since the dynamical interaction between the robot and environment is intermittent, the model is discrete. Our simple model views the position and velocity of the puck as a controlled dynamical system whose control inputs correspond to the position and velocity at impact of the arm. We treat the robot simply as a feedback agent which makes observations regarding the state of the environment [1] and takes appropriate action according to our exact specifications. Thus, we consider the environmental control problem in isolation. Since we are about to begin a series of experiments with the actual physical apparatus, the shortcomings of these conceptual distinctions and simplifying assumptions will become readily apparent in the near future.

In Section 3, we define a simple class of vertical juggling tasks as a subset in state space, and the notion of a "vertical one-juggle" — a feedback law with respect to that task subset. We find necessary and sufficient conditions under which a specified member of that set could, indeed, represent the limiting behavior of the controlled dynamical environment — i.e. when it

[1]The first physical experiments will involve pucks with an active signal source permitting complete state information to be used in the control law. Subsequent trials, of course, will involve less than full state feedback.

may constitute a fixed point of some closed loop map.

In Section 4 we turn our attention to the stabilizability of fixed points in the task subset. Specifically, we show that the system is locally controllable at any point in this set. We finally propose a specific strategy for the vertical one-juggle (which takes the form of a linear feedback algorithm) and demonstrate that the algorithms are indeed, correct — i.e. they locally stabilize the system around a specified task. Illustrative simulations are presented at the end of the Section.

2 Modeling the Environment

Here we develop a simple model for the "task" confronting a one degree of freedom juggling robot. The term "one-juggle" expresses the fact that we consider, at this stage, batting about only one object. We presume that the generalization to n objects will require a thorough analysis of this simpler task. For, simple as the task may be, the model developed below is a fourth order nonlinear discrete dynamical control system. Thus, its analysis is not trivial, and it will likely exhibit phenomena that turn out to be fundamental to more general instances of dynamic interaction with an intermittent environment. [2] The reader who is unfamiliar with the analysis of discrete dynamic systems might wish to consult the recent text by Guckenheimer and Holmes [1], to see just how complex apparently simple models may be.

In the following discussion, refer to the sketch of the system given in Figure 4. Let p_j denote the ball position at impact; $\dot{p}_j$, the ball velocity just before; and $\dot{p}'_j$, the velocity just after impact, all as seen from frame $\mathcal{F}_0$ in the robot's base. As stated above, we ignore the robot's dynamics: actuator positions and velocities are assumed to be specified arbitrarily and are not at all influenced by the interaction with the ball. This derivation is based on the following additional simplifying assumptions.

First, we assume here that all interactions between ball and robot during impact can be modelled adequately as an instantaneous event. Namely, the velocity component of the ball perpendicular to the bar, after impact j is given as in [4],

$$ {}^1\dot{y}'_j = -\alpha\, {}^1\dot{y}_j + (1+\alpha) u_{2,j} \tag{1} $$

where ${}^1\dot{y}'_j$ is the perpendicular velocity component of the ball after impact in the $\mathcal{F}_1$ coordinate frame, $\alpha \in (0,1)$ is the "coefficient of restitution", and $u_{2,j} = \|p_j\|\dot{\theta}$ is the linear velocity of the robot at the impact point p_j. Next, it is assumed that the parallel velocity components to the robot bar Finally, spin effects and friction during flight are neglected as well.

Under these conditions, the velocity of the ball after impact in the $\mathcal{F}_1$ coordinate frame with origin at p_j is

[2] Even with a single object, there are, evidently, many more interesting tasks than the vertical one-juggle. For instance, such tasks as an asymmetric side-to-side juggle must be modeled as orbits of period $n > 1$, and are a present subject of active investigation.

$$ {}^1\dot{p}'_j = \begin{bmatrix} 1 & 0 \\ 0 & -\alpha \end{bmatrix} {}^1\dot{p}_j + \begin{bmatrix} 0 \\ 1+\alpha \end{bmatrix} u_{2,j} = A\,{}^1\dot{p}_j + b' u_{2,j}. \tag{2} $$

To transform this back into $\mathcal{F}_0$ we use

$$ R_j = \begin{bmatrix} \cos\theta_j & -\sin\theta_j \\ \sin\theta_j & \cos\theta_j \end{bmatrix}. $$

While $\theta_j = \arctan \frac{y_j}{x_j}$ is a function of the ball position at the moment of impact, we assume that the robot velocity is independent of the ball velocity, thus

$$ {}^1\dot{p}_j = R_j^T \dot{p}_j, $$

and, (2) may be written as

$$ \dot{p}'_j = M_j \dot{p}_j + b_j u_{2,j}, \tag{3} $$

where

$$ M_j \triangleq R_j A R_j^T; \qquad b_j \triangleq R_j b'. $$

Denote the space of control inputs as $\mathcal{U} \triangleq \mathbb{R}^2$, the state space as $\mathcal{Z} = \mathbb{R}^4$ and the complete state of the ball as

$$ z = \begin{bmatrix} p \\ \dot{p} \end{bmatrix} = \begin{bmatrix} x \\ y \\ \dot{x} \\ \dot{y} \end{bmatrix}. $$

This defines a discrete dynamical system,

$$ \begin{aligned} z_{j+1} &= f(z_j, u_j) = \begin{bmatrix} p_j + \dot{p}'_j u_{1,j} - a u_{1,j}^2 \\ \dot{p}'_j - 2a u_{1,j} \end{bmatrix}, \\ &= \begin{bmatrix} p_j + (M_j \dot{p}_j + b_j u_{2,j}) u_{1,j} - a u_{1,j}^2 \\ M_j \dot{p}_j + b_j u_{2,j} - 2a u_{1,j} \end{bmatrix} \end{aligned} \tag{4} $$

with $a = \begin{bmatrix} 0 \\ -\frac{1}{2} g \end{bmatrix}$.

The quantity u_1 denotes the time elapsed between impact j and (j+1). Evidently, this interval uniquely specifies θ_{j+1}, the angle of the robot's joint at the next impact: we may regard $u_{1,j}$ as a robot control input to the environment. Recall, that $u_{2,j}$ is the linear velocity of the robot at the instant of impact at point p_j — a second control input. Thus, the robot acts upon the environment at intermittent times and positions which are a function of their mutual dynamical interaction.

3 The Vertical One-Juggle

Probably the simplest systematic behavior of this environment imaginable (after the rest position), is a periodic vertical motion of the puck in its plane. We are led to the following definition. Let the task subspace of the vertical one-juggle be the plane

$$\mathcal{T} \triangleq span \left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \right\} = \{ z \in \mathcal{Z} : \dot{x} = 0, \ y = 0 \}.$$

We will say that a feedback law, $g : \mathcal{Z} \to \mathcal{U}$, constitutes a *vertical one-juggle* with respect to the task, $z^* \in \mathcal{T}$, if z^* is a fixed point of the closed loop system,

$$z^* = f_g(z^*); \qquad f_g(z) \triangleq f(z, g(z)),$$

and is a stable attractor of the resulting discrete dynamics.

Proposition 1 *Given the discrete dynamical control system, (4), and a point, $z^* \in \mathcal{Z}$, there exists a feedback law, $g : \mathcal{Z} \to \mathcal{U}$ such that z^* is a fixed point of the closed loop map,*

$$z^* = f_g(z^*),$$

if and only if

$$\begin{array}{ll} (i) & z^* \in \mathcal{T}; \\ (ii) & g(z^*) = u^* = \begin{bmatrix} -2/g \\ -\frac{1-\alpha}{1+\alpha} \end{bmatrix} \dot{y}^*. \end{array}$$

Proof: Consider the fixed point condition

$$z = f(z, u) = \begin{bmatrix} p + \dot{p}' u_1 - a u_1^2 \\ \dot{p}' - 2 a u_1 \end{bmatrix}$$

Elimination of $\dot{p}'$ directly results in

$$\begin{array}{ll} (1) & \dot{x}^* = 0 \\ (2) & u_1^*(z^*) = -2\dot{y}^*/g. \end{array}$$

With these conditions and

$$\dot{p}' = M\dot{p} + b u_2 = a u_1$$

we obtain two more conditions

$$\begin{array}{ll} (3) & y^* = 0 \\ (4) & u_2^*(z^*) = -\frac{1-\alpha}{1+\alpha} \dot{y}^* \end{array}$$

where conditions (1),(3)[(2),(4)] are equivalent to (i)[(ii)].

□

This result shows, on the one hand, that only a point in $\mathcal{T}$ may be fixed by feedback, and, on the other hand, that an appropriate constant, u^*, may be found to fix *any* point of $\mathcal{T}$. We remark, in passing, that the same constant offset, u^*, fixes *every* point on a line in the task plane, hence by itself, could not possibly stabilize any particular point on that line (here, stabilize is being used to denote the property of attractivity as well as stability).

4 Local Stabilizability of the Task Plane

We finally observe that the system is locally controllable at any point in the task set, and provide a few examples of feedback control algorithms which locally stabilize the system around a particular task.

Proposition 2 *If*

$$z^* \in \mathcal{T} - 0$$

then the system 4 is locally controllable. That is, if

$$A = D_z f[z^*, u^*]; \quad B = D_u f[z^*, u^*]$$

then (A,B) is a completely controllable pair.

Proof: Taking partial derivatives of (4) gives

$$A = \begin{bmatrix} 1 & 2\omega\frac{\dot{y}^{*2}}{x^*} & \omega\dot{y}^* & 0 \\ 0 & 1 & 0 & -\alpha\omega\dot{y}^* \\ 0 & 2\frac{\dot{y}^*}{x^*} & 1 & 0 \\ 0 & 0 & 0 & -\alpha \end{bmatrix}; \quad B = \begin{bmatrix} 0 & 0 \\ \dot{y}^* & (1+\alpha)\omega\dot{y}^* \\ 0 & 0 \\ -g & 1+\alpha \end{bmatrix}$$

where ω denotes the constant

$$\omega = -\frac{2}{g}.$$

It suffices to show that four of the eight columns of the matrix (B, AB, A^2B, A^3B) are linearly independent. The four columns (B, A^2B) we consider are

$$(B, A^2B) = \begin{bmatrix} 0 & 0 & \frac{4\dot{y}^{*3}}{gx^*}(2\alpha - 3) & \frac{8\dot{y}^{*3}}{g^2x^*}(1+\alpha)(3-\alpha) \\ \dot{y}^* & -\frac{2}{g}\dot{y}^*(1+\alpha) & \dot{y}^*(2\alpha^2 - 2\alpha + 1) & -\frac{2}{g}\dot{y}^*(1-\alpha+\alpha^2) \\ 0 & 0 & 4\frac{\dot{y}^2}{x^*}(1-\alpha) & -\frac{4\dot{y}^{*2}}{gx^*}(1+\alpha)(2-\alpha) \\ -g & 1+\alpha & -\alpha^2 g & \alpha^2(1+\alpha) \end{bmatrix}.$$

The determinant of this matrix,

$$\frac{16}{g^2}\frac{\dot{y}^{*6}}{x^{*2}}\alpha(1+\alpha)^2$$

is nonzero for any $z^* \in \mathcal{T} - 0$.

□

According to Proposition 1 , any feedback law which fixes a point, z^* must include the constant term, u^* On the other hand, the present result assures us of the existence of some affine feedback law which makes the system locally asymptotically stable. The most general affine feedback structure for some $z^* \in \mathcal{T}$ is

$$u(z) = g_{z^*}(z) + g_{fb}(z) = g_{z^*}(z) + K(z - z^*)$$

with

$$g_{z^*}(z) = \begin{bmatrix} -2/g \\ -\frac{1-\alpha}{1+\alpha} \end{bmatrix} \dot{y}$$

and

$$K = \begin{bmatrix} k_1^T & k_2^T \\ k_3^T & k_4^T \end{bmatrix}$$

For ease of analysis, we will not consider this general algorithm, but present the simpler version,

$$K = \begin{bmatrix} k_{11} & k_{12} & k_{21} & 0 \\ 0 & 0 & 0 & k_{42} \end{bmatrix}; \qquad z^* = \begin{bmatrix} x^* \\ 0 \\ 0 \\ \dot{y}^* \end{bmatrix}.$$

It is worth noting, again in passing, that, unlike the constant term, u^*, this feedback fixes z^* uniquely as long as neither k_{11} nor k_{42} are set to zero. Thus, it makes sense to examine the local stability properties of the resulting closed loop map,

$$f_g(z) = f(z, g_{z^*}(z) + g_{fb}(z)),$$

in the neighborhood of z^*.

In order to do so, we compute the Jacobian,

$$D_z f_g(z^*) = \begin{bmatrix} 1 & 2\omega\frac{\dot{y}^{2*}}{x^*} & \omega\dot{y}^* & 0 \\ k_{11}\dot{y}^* & 1 + k_{12}\dot{y}^* & k_{21}\dot{y}^* & k_{42}(1+\alpha)\omega\dot{y}^* \\ 0 & 2\frac{\dot{y}^*}{x^*} & 1 & 0 \\ \frac{2k_{11}}{\omega} & \frac{2k_{12}}{\omega} & \frac{2k_{21}}{\omega} & 1 + k_{42}(1+\alpha) \end{bmatrix}. \qquad (5)$$

Note that choosing $k_{11} = k_{42} = 0$ (eliminating the terms in the feedback law which specify a unique fixed point in the task plane) results in a spectrum for the array (5) which has two unit eigenvalues $\lambda_{1,2}$ with eigenvectors in the task plane. Thus, the closed loop system is, at best, marginally stable — we already know it cannot be asymptotically stable since the fixed points are not isolated when $k_{11} = k_{42} = 0$. Examination of the remaining eigenvalues in this case shows that they may be readily assigned by appropriate choice of feedback parameters. Of course, for this non-hyperbolic case, the behavior of the linearized system contains no information about even the local stability properties of the true nonlinear closed loop system. However, in this case, simulations suggest that the fixed points are locally stable when the remaining eigenvalues are placed inside the unit circle. Of course, it is easy to find initial conditions outside the stable region for any of these parameter settings. These considerations begin to suggest that horizontal position and vertical velocity (or, at least estimates thereof) are necessary for stabilizing a point in the task plane $\mathcal{T}$.

Clearly, to achieve a vertical one-juggle — an attracting stable fixed point at $z^* \in \mathcal{T} - 0$ — we will require at least the generality of the reduced affine feedback law which led to (5). Unfortunately, at the time of this preliminary analysis, we have not yet determined whether the structure of the linearized system admits any simple choice of parameter values in order to place the poles of the resulting linearized closed loop system with any precision. It is, however, not hard to find stabilizing values.

To demonstrate this last point, and as a matter of idependent interest, we have simulated the full nonlinear model with a variety of theoretically stabilizing parameter values. We may select a set of feedback parameters K_I for which the (numerically evaluated) spectrum is

$$\sigma_I = (0.11, 0.716 \pm 0.168i, 0.9).$$

The response to the initial state z_0 and the desired state z^*:

$$z_0 = \begin{bmatrix} 4 \\ 0 \\ -0.5 \\ -6 \end{bmatrix}; \quad z^* = \begin{bmatrix} 5 \\ 0 \\ 0 \\ -5 \end{bmatrix}; \quad K_I = \begin{bmatrix} -0.01 & 0.33 & -0.05 & 0 \\ 0 & 0 & 0 & 0.05 \end{bmatrix},$$

is shown in Figures 2-5 These simulations confirm our expectation of stabilizing (4) around an arbitrary point in the task space, $\mathcal{T}$-0. However, as was to be expected as well, we only have local stability: the simulation responses given sufficiently large initial errors, $z - z^*$, evince unstable behavior.

Selecting feedback

$$K_{II} = \begin{bmatrix} -0.01 & 0.15 & -0.05 & 0 \\ 0 & 0 & 0 & 0.05 \end{bmatrix}$$

such that the spectrum is

$$\sigma_{II} = (0.81 \pm 0.716, 0.825, 0.9)$$

which has two poles just outside the unit circle, with the same z_0, z^* as before, results in the unstable responses shown in Figures 6-9.

Figures

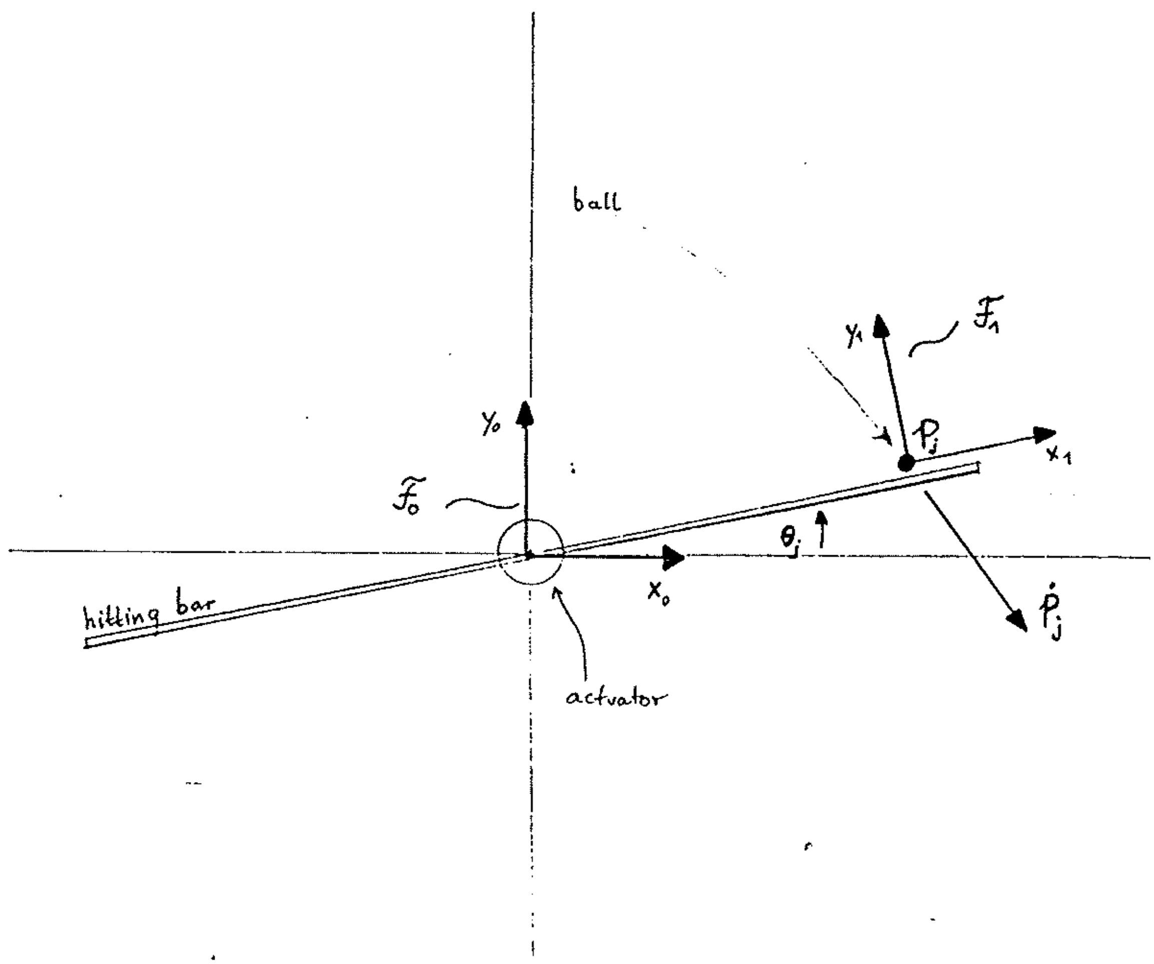

Figure 1: The Juggler

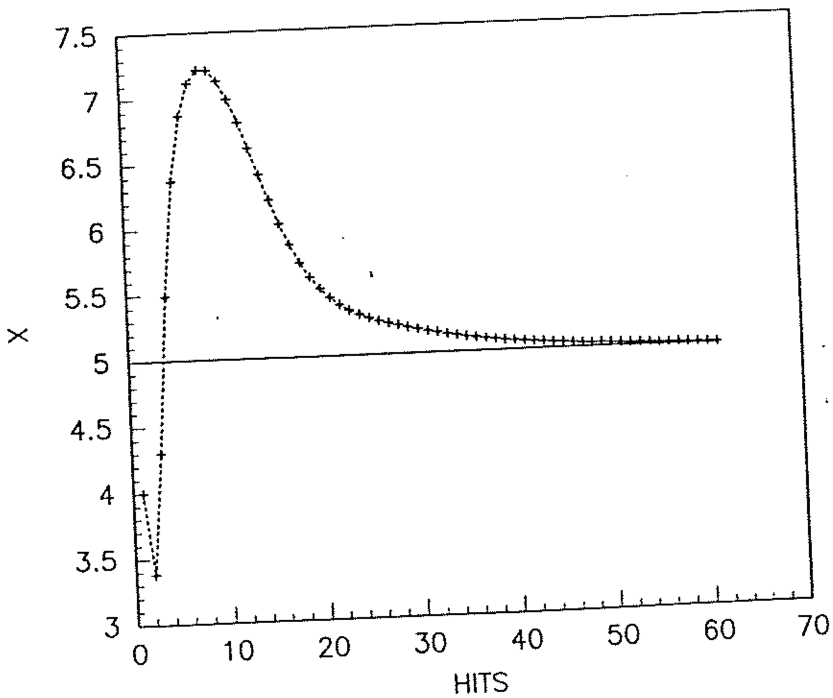

Figure 2: Stable Vertical Juggle: x coordinate

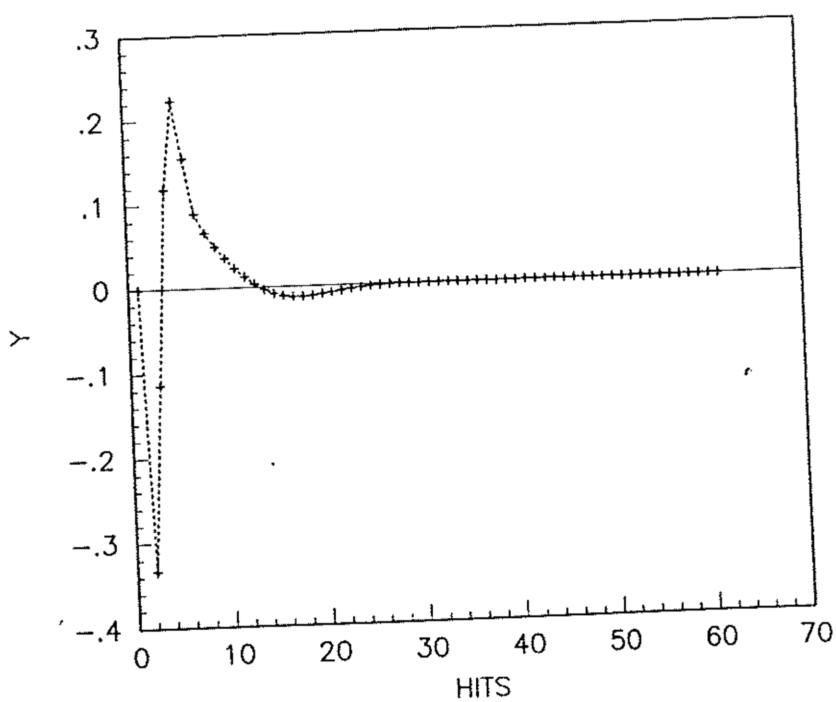

Figure 3: Stable Vertical Juggle: y coordinate

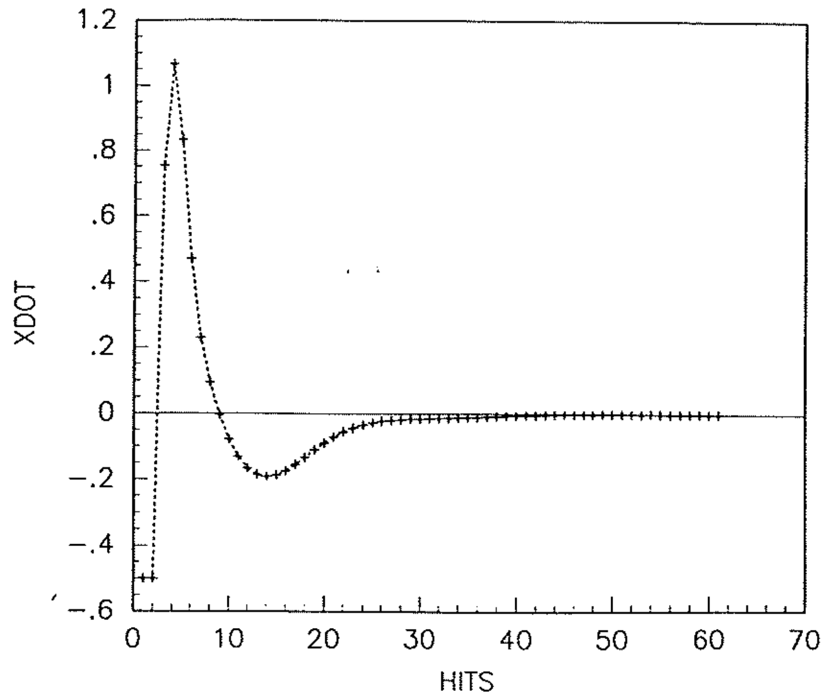

Figure 4: Stable Vertical Juggle: $\dot{x}$ coordinate

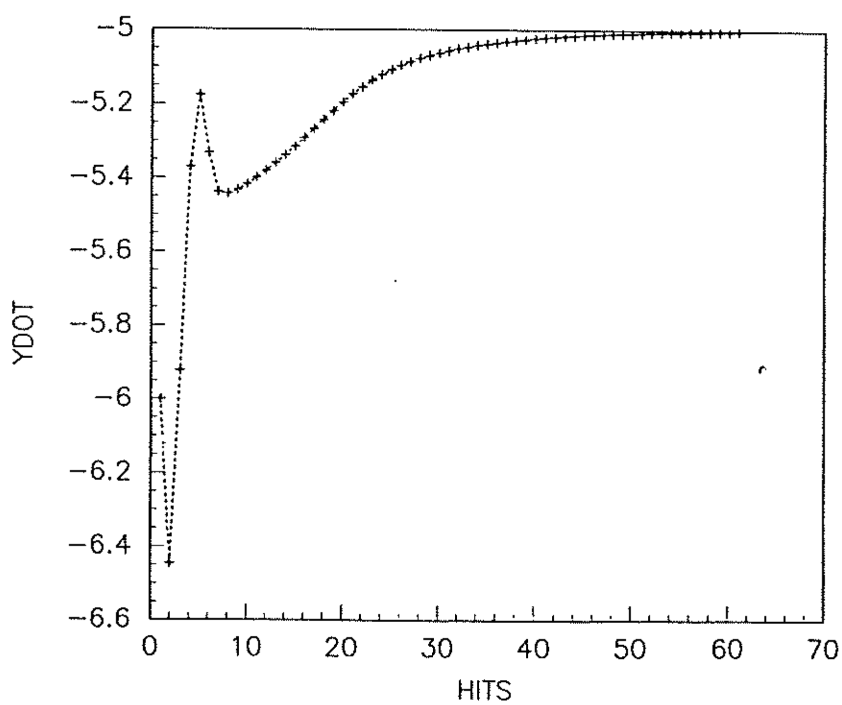

Figure 5: Stable Vertical Juggle: $\dot{y}$ coordinate

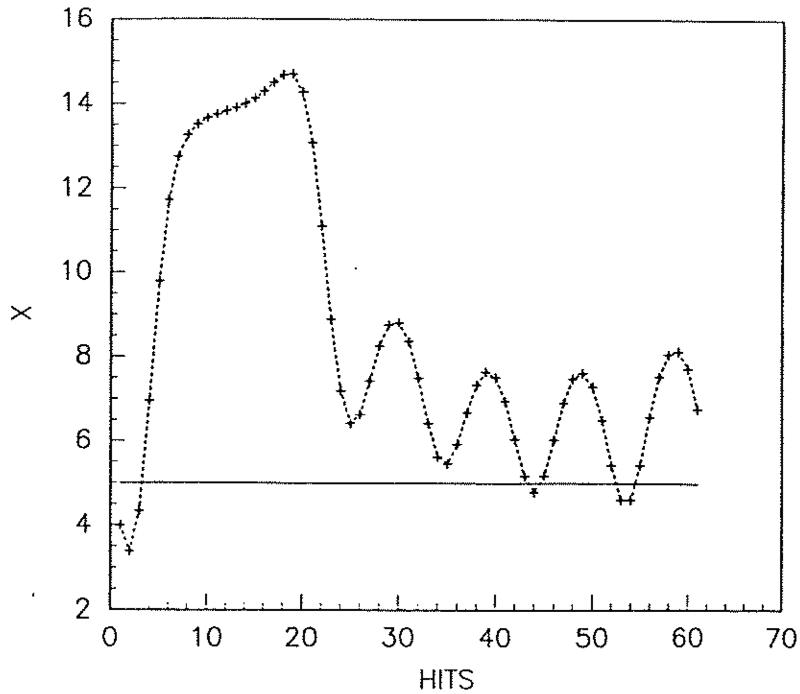

Figure 6: Unstable Vertical Juggle: x coordinate

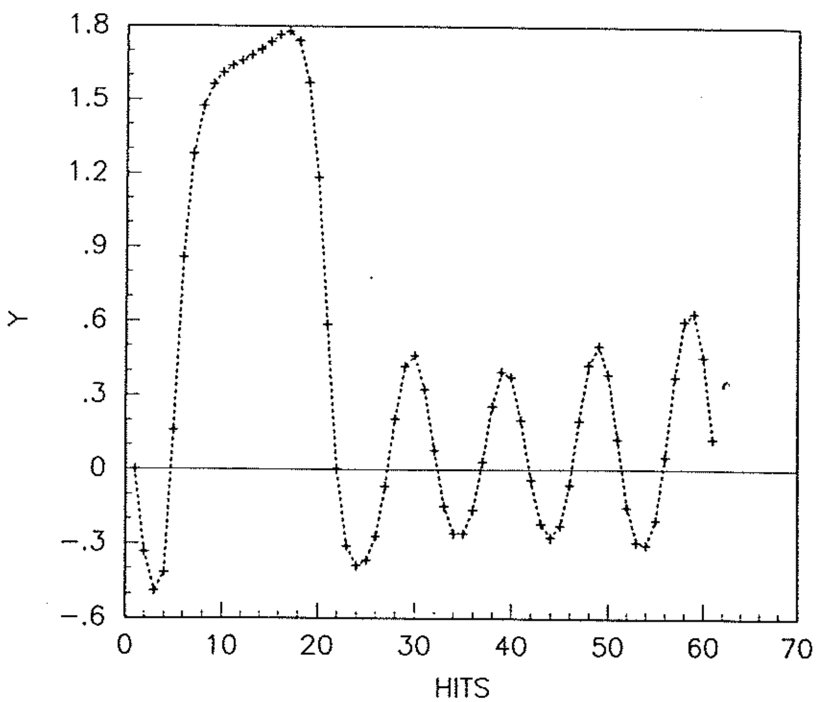

Figure 7: Unstable Vertical Juggle: y coordinate

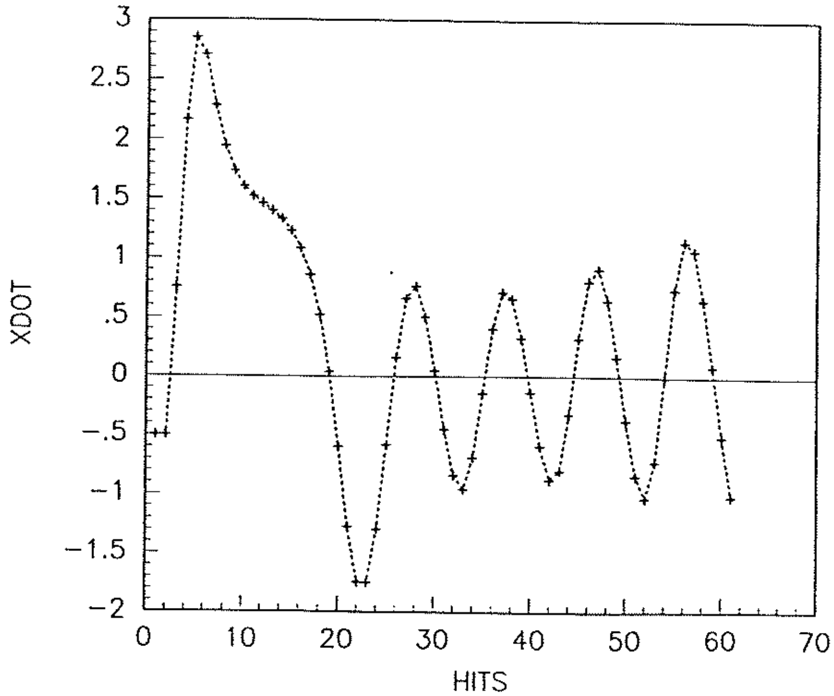

Figure 8: Unstable Vertical Juggle: $\dot{x}$ coordinate

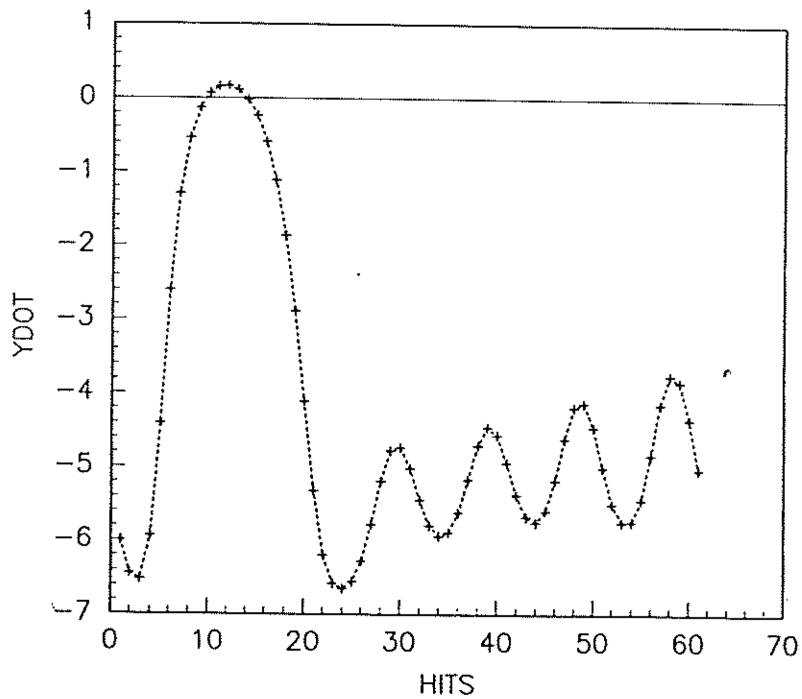

Figure 9: Unstable Vertical Juggle: $\dot{y}$ coordinate

References

[1] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* Springer-Verlag, New York, 1983.

[2] Neville Hogan. Impedance control: an approach to manipulation. *ASME Journal of Dynamics Systems,Measurement, and Control*, 107:1–7, Mar 1985.

[3] Marc H. Raibert. *Legged Robots That Balance.* MIT Press, Cambridge, MA, 1986.

[4] J. L. Synge and B. A. Griffith. *Principles of Mechanics.* McGraw Hill, London, 1959.