

Notes on Balancing an Inverted Pendulum

Yoke Peng Leong

January 8, 2014

1 Introduction

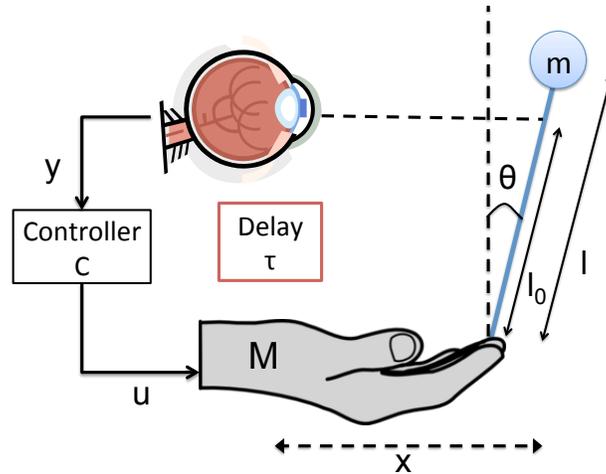


Figure 1: A schematic of balancing an inverted pendulum on one's palm where m is the pendulum mass, M is the arm mass, l is the location of center of mass (COM), l_0 is the measurement height from the hand y is position measurement using the eye, u is the control force, θ is the pendulum tilt angle from the vertical, and x is the horizontal displacement of the arm.

Balancing an inverted pendulum on a person's hand (Figure 1) illustrates fundamental limits that constrain what robust efficiency is achievable and can be reasonably thoroughly investigated using office materials and undergraduate mathematics, plus a little ingenuity. There is a large and growing literature on related balancing tasks in sensorimotor control neuroscience [1–6], and we hope our approach will help clarify some of the confusion surrounding the experiments. Surprisingly, the most essential issues can be illustrated experimentally with nothing more than an extendable pointer whose length can be changed, and then balanced on the palm or finger tip using visual feedback.

The first experimental observation is that it is trivial to hold the pointer down without active control, even with closed eyes, but difficult to balance the pointer in the up position. In the up case, the most important parameters are the lengths l from the bottom tip to the center of mass

(CoM) and l_0 to where the person looks at the pointer, and for the top tip $l_0 > l$. The next experimental result is that when $l_0 > l$ balancing becomes harder and then impossible as l gets sufficiently short. For most healthy adults the failure point comes for $.25m < l < .5m$. An even more intriguing result is that keeping $l = 1m$ but shifting the point l_0 down also causes failure, but more dramatically at much higher l_0 , even $.8m < l_0$. To get this effect, it helps to use the other hand to shield the eyes from using peripheral vision to see the pointer above l_0 . Subjects vary greatly in their use of peripheral vision, so simply preventing it is the easiest way to get consistent results. These apparently cryptic experimental results can be explained in terms of standard sensorimotor control concepts, basic and known neuroanatomy, and elementary undergraduate math, largely independent of additional and highly complex neural details that remain unresolved.

There are several points to be made with this example. The most essential features are that physical efficiency and total robustness have hard tradeoffs, and that both efficiency and delays in the cyber (communication and control) components necessarily aggravate achievable robustness. One universal and familiar feature of real systems in engineering and biology is that improving efficiency in terms of minimizing both the consumption and waste of material, energy, and other resources inevitably reduces the intrinsic stability of the systems. Modern airplanes, rockets, comms networks, supply chains, etc and most importantly, future smartgrids, automobiles, highways, etc. will crash without automatic control systems. Complex life forms will also crash without complex systems to maintain homeostasis and control movement. But precisely what the costs in increased crash fragility due to improved efficiencies are has been inadequately addressed until now. This issue is captured in a highly stylized way in the difference between the stable down and unstable up conditions, and in the increased difficulty of preventing a “crash” (the pointer falling over) as the pointer is made shorter (analogous to using less resources). A further universal issue that this simple example illustrates is the tradeoff between cost of cyber implementation as reflected in delay versus robustness.

2 Inverted Pendulum Model

The standard control model of a 1D inverted pendulum on a moving cart is given by

$$\begin{aligned}
 (M + m)\ddot{x} + ml(\ddot{\theta} \cos \theta - \dot{\theta}^2 \sin \theta) &= u + r \\
 m(\ddot{x} \cos \theta + l\ddot{\theta} - g \sin \theta) &= 0 \\
 z = x + l_0 \sin \theta \quad y = z + n &
 \end{aligned} \tag{1}$$

where y is position measurement using the eye of z , the position of interest, u is the control force, r is actuation noise, and n is sensor noise, θ is the pendulum tilt angle from the vertical, and x is the horizontal displacement of the arm. M is the mass of the cart, m is the mass of the pendulum, g is gravitational acceleration, and l and l_0 are the COM and measurement heights. The absolute value of z should be small because it corresponds to how far the pendulum has drifted away from the desired equilibrium point. As a model of real 3D balancing it is grossly oversimplified, but will prove to be surprisingly useful. We will systematically oversimplify the modeling but in ways that also would simplify the task itself, and argue that hard limits we derive would overestimate the

real robustness (or underestimate fragility). In fact, we can get quite close to real experiments with very simple, even analytically tractable models.

To study local stability, the set of equations (1) are linearized around the up and down equilibria, and its linearized and Laplace-transformed forms are shown in Table 1.

Linearized	Laplace transformed
$(M + m)\ddot{x} \pm ml\ddot{\theta} = u + r$ $m(\pm\ddot{x} + l\ddot{\theta} \mp g\theta) = 0$ $z = x \pm l_0\theta$ $y = z + n$	$\begin{bmatrix} \hat{x} \\ \hat{\theta} \end{bmatrix} = \frac{1}{D(s)} \begin{bmatrix} ls^2 \mp g \\ \mp s^2 \end{bmatrix} (\hat{u} + \hat{r})$ $\hat{z} = \frac{(l-l_0)s^2 \mp g}{D(s)} (\hat{u} + \hat{r})$ $\hat{y} = \hat{z} + \hat{n}$

Table 1: Linearized and Laplace-transformed (1) where $D(s) = s^2(Mls^2 \mp (M + m)g)$. The top (bottom) sign in \pm or \mp corresponds to linearization around up (down) equilibrium.

3 Stability Analysis

In this section, we will discuss how the poles, and zeros relate to the behavior of this system. The poles and zeros of the open loop plant are given by Table 2.

First, consider the poles for both upright and downward positions. For the downward position, as l increases, the magnitude of the poles also decreases. However, because they are imaginary, the magnitude of the poles is the resonance frequency of the system, and the resonance frequency increases as l decreases, as is easily verified experimentally. For the upright position, as l decreases, the magnitude of the poles increases, but here the poles are real so the positive one is more unstable, and thus harder to control. This is aggravated further by time delays in the nervous system to sense, decide, and act. It is further aggravated by an unstable zero, which accounts for the additional difficulty of control when $l_0 < l$.

When $l_0 = l$, zero does not exist for both positions. But, when $l_0 \neq l$, the upright position has an opposite effect to the downward position for the same measurement distance away from the

Open loop plant: $P(s) = \frac{(l-l_0)s^2 \mp g}{D(s)}$		
Positions	Poles	Zeros
Upright	$0, \pm\sqrt{\frac{(M+m)g}{Ml}}$	$\pm i\sqrt{\frac{g}{l_0-l}}$ if $l_0 > l$ none if $l_0 = l$ $\pm\sqrt{\frac{g}{l-l_0}}$ if $l_0 < l$
Downward	$0, \pm i\sqrt{\frac{(M+m)g}{Ml}}$	$\pm\sqrt{\frac{g}{l_0-l}}$ if $l_0 > l$ none if $l_0 = l$ $\pm i\sqrt{\frac{g}{l-l_0}}$ if $l_0 < l$

Table 2: Poles and zeros of an inverted pendulum on a moving cart model.

hand. A real unstable zero represents the pendulum swing frequency whereby the pendulum will look stationary although it is moving. Notice that for the upright position, as l_0 moves from the center of mass of the inverted pendulum towards the bottom of it, the magnitude of zero moves from infinity to a small real number. When balancing an inverted pendulum, the pendulum swing has a relatively low frequency. This shift from large to small zero implies that the pendulum will more likely to look stationary as the measurement height decreases because the zero will get closer to the nominal pendulum swing frequency. The next section will further quantify this phenomena using the Bode's integral formula.

Because the interesting case is the upright position, from here onwards, the analysis will focus only on the upright position.

4 Robust Performance Analysis

The dynamical model of an inverted pendulum on a moving cart can be derived using three "laws" well-known since Newton involving mechanics, gravity, and optics. However, these laws (and their modern generalizations) say nothing about the limits of robustness and efficiency or how they interact with each other and other system parameters. We need an additional "law", known to control theorists as Bode's integral formula [7], to study the relationship among robustness, efficiency, delays, and sensing.

The closed-loop output is derived from equations in Table 1

$$\hat{z} = \frac{1}{D(s) + ((l - l_0)s^2 - g)C(s)} \begin{bmatrix} (l - l_0)s^2 - g & -C(s)((l - l_0)s^2 - g) \end{bmatrix} \begin{bmatrix} \hat{r} \\ \hat{n} \end{bmatrix}. \quad (2)$$

Define the sensitivity function, $S(s)$, as

$$\begin{aligned} S(s) &= \frac{1}{1 + P(s)C(s)} \\ &= \frac{D(s)}{D(s) + ((l - l_0)s^2 - g)C(s)} \end{aligned} \quad (3)$$

and the complimentary sensitivity function, $T(s)$, as

$$\begin{aligned} T(s) &= \frac{P(s)C(s)}{1 + P(s)C(s)} \\ &= \frac{((l - l_0)s^2 - g)C(s)e^{-\tau s}}{D(s) + ((l - l_0)s^2 - g)C(s)}. \end{aligned} \quad (4)$$

Then, (2) can be rewrite as

$$\hat{z} = P(s)S(s)\hat{r} - T(s)\hat{n}. \quad (5)$$

For any given noises r and n , z is desired to be small. More concretely, define the performance measure as the infinity norm of the closed-loop transfer functions from noise sources to z . A robust

performance is to have the infinity norm of these transfer functions be small (i.e. to have the magnitude of the transfer function at all frequency be small). This performance measure quantifies the effects of noises on the output.

A useful way of analyzing the transfer function is by using the Bode's integral formula. A general form of this formula is given by

$$\ln |G_{mp}(s_0)| = \frac{1}{\pi} \int_{-\infty}^{\infty} \ln |G(j\omega)| \frac{\sigma_0}{\sigma_0^2 + (\omega - \omega_0)^2} d\omega \quad (6)$$

where $s_0 = \sigma_0 + \omega_0 j$ with $\sigma_0 > 0$, $G(s)$ is any transfer function, and $G_{mp}(s)$ is a minimum phase transfer function of $G(s)$.

First, consider the transfer function from $n \rightarrow z$ that quantifies the effects of measurement noise on the performance measure. For this transfer function, Bode's integral formula is bounded as follows

$$\frac{1}{\pi} \int_{-\infty}^{\infty} \ln |T(j\omega)| \frac{p}{p^2 + \omega^2} d\omega \geq \tau p + \begin{cases} 0 & l_0 \geq l \\ \ln \left| \frac{z+p}{z-p} \right| & l_0 < l \end{cases} \quad (7)$$

where $T(s)$ is the transfer function from measurement noise n to output z , $T_{mp}(s)$ is the minimum phase part of $T(s)$, $p = \sqrt{\frac{(M+m)g}{Mt}}$ is the RHP pole, $z = \sqrt{\frac{g}{l-l_0}}$ is the RHP zero, and τ is the total delay in the control system from measurement to actuator. This integral places a hard lower bound on the net achievable fragility of a control system in the presence of measurement noise. A small integral implies a less fragile (i.e. more robust) system is possible, and sensor noise does not necessarily prevent stabilizing the pendulum.

Proof. Factor the complimentary sensitivity function as

$$T(s) = T_{mp}(s)T_{ap}(s)$$

where mp is minimum-phase and ap is all-pass.

When $l_0 < l$, there is only one RHP zero,

$$T_{ap}(s) = \frac{z-s}{z+s} e^{-\tau s}.$$

Then,

$$T_{mp}(p) = T_{ap}(p)^{-1} = \frac{z+p}{z-p} e^{\tau p}.$$

Substitute into (6),

$$\begin{aligned} \frac{1}{\pi} \int_{-\infty}^{\infty} \ln |T(j\omega)| \frac{p}{p^2 + \omega^2} d\omega &= \ln |T_{mp}(p)| \\ &= \ln \left(\left| \frac{z+p}{z-p} \right| e^{\tau p} \right) \\ &= \tau p + \ln \left| \frac{z+p}{z-p} \right|. \end{aligned}$$

When $l_0 \geq l$, there is no RHP zero,

$$\begin{aligned} \frac{1}{\pi} \int_{-\infty}^{\infty} \ln |T(j\omega)| \frac{p}{p^2 + \omega^2} d\omega &= \ln |T_{mp}(p)| \\ &= \ln (e^{\tau p}) \\ &= \tau p. \end{aligned}$$

□

For the other transfer function from $r \rightarrow z$ that quantifies the effect of actuation noise on the performance measure, only $S(s)$ is considered because $P(s)$ can be thought as a weight on $S(s)$. Thus, the Bode's integral is bounded by

$$\begin{aligned} \frac{1}{\pi} \int_{-\infty}^{\infty} \ln |S(j\omega)| \frac{z}{z^2 + \omega^2} d\omega &\geq \begin{cases} 0 & l_0 > l \\ \ln \left| \frac{z+p}{z-p} \right| & l_0 < l \end{cases} \\ \frac{1}{\pi} \int_{-\infty}^{\infty} \ln |S(j\omega)| d\omega &\geq 2p \quad l_0 = l \end{aligned} \quad (8)$$

where $S(s)$ quantifies the effect of actuation noise r to output z , $S_{mp}(s)$ is the minimum phase part of $S(s)$, $p = \sqrt{\frac{(M+m)g}{Ml}}$ is the RHP pole, and $z = \sqrt{\frac{g}{l-l_0}}$ is the RHP zero. When $l_0 = l$, the Bode's integral is bounded only by the RHP pole which characterizes the amount of instability of the system. This bound is consistent with the intuition that as a system becomes more unstable, the same amount of actuation noise results in more detrimental effect to the robustness of the system.

Proof. When $l_0 \neq l$, the proof follows the same steps as the proof for bounding the Bode's integral of $T(s)$.

When $l_0 = l$, there is no RHP zero (i.e. the RHP zero is at infinity). Then,

$$\frac{1}{\pi} \int_{-\infty}^{\infty} \ln |S(j\omega)| \frac{z}{z^2 + \omega^2} d\omega = \ln \left| \frac{z+p}{z-p} \right|.$$

Multiply both side by z ,

$$\frac{1}{\pi} \int_{-\infty}^{\infty} \ln |S(j\omega)| \frac{z^2}{z^2 + \omega^2} d\omega = z \ln \left| \frac{z+p}{z-p} \right|.$$

Let $z = \frac{1}{\epsilon}$. Then,

$$\begin{aligned} \frac{z+p}{z-p} &= \frac{1+\epsilon p}{1-\epsilon p} \\ &= (1+\epsilon p)(1+\epsilon p + (\epsilon p)^2 + \dots) \\ &= 1 + 2\epsilon p + \dots \end{aligned}$$

$$\begin{aligned}\ln \left| \frac{z+p}{z-p} \right| &= \ln |1 + 2\epsilon p + \dots| \\ &= 2\epsilon p - \frac{(2\epsilon p)^2}{2} + \dots\end{aligned}$$

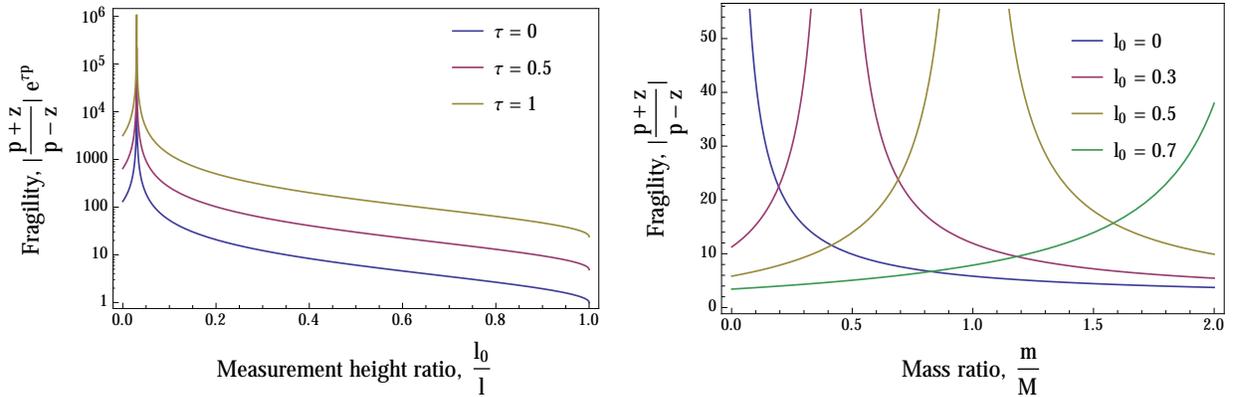
Let $z \rightarrow \infty$ to obtain

$$\begin{aligned}\frac{1}{\pi} \int_{-\infty}^{\infty} \ln |S(j\omega)| d\omega &= z(2\epsilon p - \frac{(2\epsilon p)^2}{2} + \dots) \\ &\geq 2p.\end{aligned}$$

□

Figure 2 are visualizations of the integral bounds. In Figure 2(a), the system becomes more fragile and harder to control when the measurement height decreases or delays increases. The plot in Figure 2(b) shows that when looking at the very bottom of the pendulum while trying to balance the pendulum, the fragility is huge if the pendulum mass is small (e.g. when balancing a stick on one's hand). As the mass increases, measuring closer to the bottom increases the fragility.

More informatively, Figure 3(a) has a plot of the Bode's integral bound as a function of measurement or pointer height under three different assumptions. The middle green curve is fragility versus pendulum length when $l_0 = l$ is varied and delay is 0.3s, a typical value from the literature [8–10]. The red curve is when pendulum length is fixed at 1m, but measurement l_0 is varied, with delay still 0.3s. The blue curve is when $l_0 = l$ and delay is 0.2s, which is likely to be unrealistically fast, but illustrates the extreme impact of small changes in delay. Note that Bode's integral formula gives a lower bound on the net fragility of a system (in this case to sensor noise) that holds for *all* controllers (and is also achievable, so tight). Thus the dashed line indicates a level of fragility that



(a) Fragility v.s. measurement height and delays

(b) Fragility v.s. mass and measurement height

Figure 2: (a) Assume $M = 3.25kg$ and $m = 0.1kg$. As the measurement height decreases, the system becomes more fragile and harder to control. When delays increases, the fragility increases. (b) Assume $l = 1m$. As the measurement height increases, the fragility curve shifts rightwards.

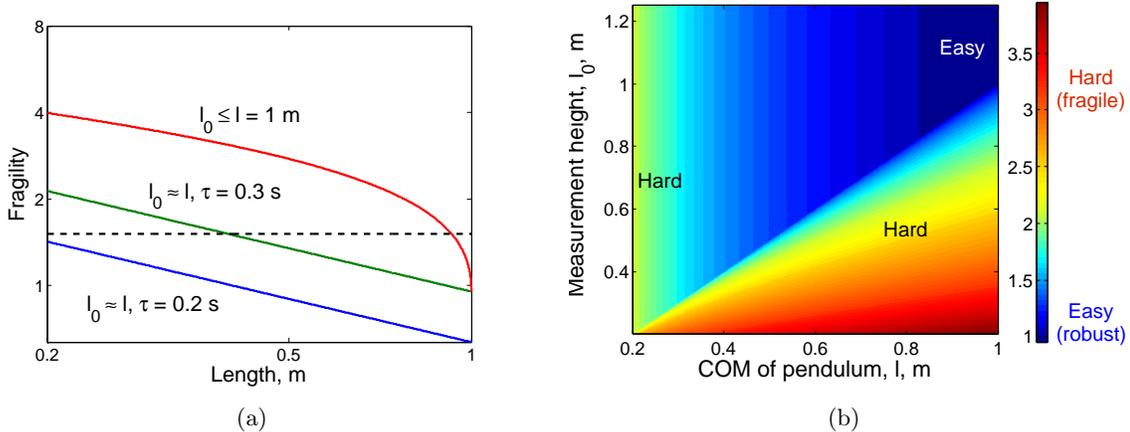


Figure 3: Plots of fragility in $T(s)$ when COM of pendulum, measurement length, and delay are varied. Assume $M = 3.25kg$, and $m = 0.1kg$. (a) The green curve is fragility versus pendulum length when $l_0 = l$ is varied and delay is 0.3s, the red curve is when pendulum length is 1m, l_0 is varied, and delay is 0.3s, and the blue curve is when $l_0 = l$ and delay is 0.2s. (b) The color corresponds to fragility.

a typical sensorimotor system (in this case Professor John Doyle) can handle without crashing, and is at least consistent with experiments, though only the red and green cases can be done experimentally.

Thus we can analytically and easily explore the impact of parameters l , l_0 , and τ on the intrinsic fragility of the resulting system without having to make further decisions about controller design. The integral constraints apply to all controllers, so define a conserved quantity (the integrand) that in turn describes an unavoidable lower bound on fragility. Thus we can predict roughly when a scenario is likely to be too fragile and will crash, without having detailed knowledge of either the sensor noise in the visual system, nor the complex neural implementation of the sensorimotor controller. All that is needed is an estimate of the delay from vision to actuation, which is well known. Because the dependencies on parameters are so strong, we actually can get reasonably quantitative predictions, and certainly qualitative insight which is perhaps more important.

Figure 3(b) show fragility as a function of both l and l_0 . It nicely captures the qualitative features of the problem, and even matches quantitatively what can be observed experimentally. This exact experiment has not been studied in the literature in the detail we would like, but it is easy for the reader to try to stabilize an extendable pointer in various (l, l_0) points on the figure, and they will find that leaving the dark blue region makes stabilization increasingly difficult and ultimately impossible.

5 Implications

Essentially, this case study highlights the following important issues in system design:

There exist a fundamental tradeoff between robustness and efficiency. By balancing a long pendulum and a short pendulum, one can easily conclude that a short pendulum is harder to balance than a long pendulum. This observation is explained by the green curve in Figure 3(a). Note that when the delay is zero, the Bode’s integral bound is constant and zero. But, when a delay is considered, the bound explicitly increases with the RHP pole p , which in turn increases as pendulum length, l , decreases. The further role of delay will be discussed in detail below. However, both experimental observations and theoretical analysis show that given sensorimotor delays, it is impossible to be both robust (balance despite noise) and efficient (small pointer). Robots or other organisms would face similar tradeoffs, with the absolute levels determined by controller delay and sensor noise.

Sensing location can have a significant impact on robustness. If the eyes’ focus point (while blocking peripheral visions) is moved from the tip of the pendulum to a lower point along the pendulum, the inverted pendulum becomes harder to balance (explained by the red curve in Figure 3(a)). Theoretically, as the measurement height decreases, the zero moves from infinity towards the pole thus increasing the bound, and thus the fragility. Interestingly, balancing a pendulum while focusing at a lower point without using peripheral vision much harder than balancing a shorter pendulum while focusing at the tip (compare green and red curves in Figure 3(a)). This observation implies that eyes’ measurement height degrades robustness even more than pendulum length (i.e. RHP zero is more detrimental than RHP pole). Thus poor sensing can greatly degrade achievable robustness beyond what is limited by other tradeoffs. Simple modifications while balancing a pointer raise interesting questions that we can answer, such as the effect of closing one eye, standing on one leg, moving and walking, a darker room, etc..

Delays have an enormous impact on robustness and efficiency tradeoff. In Figure 3(a), increasing delay, τ , moves the curves towards the northeast corner of the plot, and therefore, delay degrades the fundamental limit of this system. This result fits the observation that a trained person can balance the inverted pendulum better than an untrained person because generally, a trained person has a shorter delay in the sensorimotor system. This is an issue that pervades biology and neuroscience as well as engineering. What is particularly important about sensorimotor control is that it is *distributed*, an issue that few theories address at all. That is, the sensors, processing, communications, decision making, and actuation are distributed throughout the body in modules that communicate with each other with internal delays. So far in this case study we have only consider a single lumped source of delay, primarily due to vision, as discussed below. This idealization does not happen to be particularly limiting here, but will be in general.

Delays in this system are primarily due to vision. In fact, we can study the origin and impact of delay and how nature has intelligently coped with it through another simple demo of the vestibulo-ocular reflex (VOR) [11]. In this demo, place your hand in front of your face a comfortable distance so that you can easily see the details of your palm sharply. Then shake either your head (as if vigorously signaling “no”) or shake your hand. As you increase the frequency of shaking, you will notice that the palm blurs much more with hand motion than with head motion. Most people we have asked, including neuroscientists, predict the opposite, so are surprised by this demo. This

simple experiment shows an important feature of the human nervous system – tracking a moving object is much harder than stabilizing your gaze in a moving head. To understand this phenomena, we need to explain the different mechanisms involve for the two seemingly similar but actually quite distinct tasks, and then more deeply why this is necessary and not an evolutionary accident.

When tracking a moving object, the retina sends its signals through the ganglion cell to the cortex, and then the cortex sends signals to the muscles around the eyes as a response. This whole process takes about 100-200ms depending on light levels [8,9] from sensing to actuating the eye muscles. On the other hand, when stabilizing gaze on a moving head, the vestibulo-ocular reflex (VOR) is involved. VOR is a reflex mechanism of eye movements to stabilize gaze when the head is moving. This reflex involves a short pathway from vestibular system to brainstem, and then to eye muscles, and it does not include the cortex (i.e. less processing time), and has large heavily myelinated axons. Therefore, the latency is in the order of 9ms [12]. Because VOR doesn't use vision at all, it works in the dark, though it would be hard to easily verify this without special equipment. Both mechanisms use the same fast (< 9 ms reaction time) eye muscles so most of the above 100-200ms is in vision, computation, and communication, not actuation. Based on these simple observations, we can see the huge effects of delays on the performance of a system (i.e. ability to see the hand). Also, we can see why it is plausible that the delay from vision to hand for the balancing problem is around 300ms and is mostly due to delay thru the visual and sensorimotor systems. The muscle motion itself has abundant bandwidth for the task as illustrated by the ability of hand motion to be so fast as to completely blur the hand to vision. There are other tasks (e.g. sprinting or throwing, and much of cell biology) where actuator saturation will be more important.

Vision and VOR illustrate several other issues about layered architecture. One is the necessity of tuning the components and their networks to the system requirements. Here, the VOR system must be fast enough to allow vision to track the world despite the kind of head motion a top predator like humans would have. In contrast, vision itself is too slow to do this tracking, but fast enough for much but not all external motion. But vision is vastly more flexible illustrated by its use in reading this sentence. VOR is fast but inflexible, as it does nothing else but sense and react to head orientation and motion, whereas vision is slow and flexible. While balancing a pointer, myriad reflexes and low layer processes allow for control of body position and movement, speech production and understanding, without interfering with the main task. This universal aspect of layered architectures is arguably its most important benefit, and one that holds in our cells and technologies as well.

If you are reading this on a computer, there are myriad lower level “reflex” activities in the computer and the networks it is connected to that you need not be aware of. Indeed it is essential that you not be aware, as these would overwhelm your ability to do any task. This speed versus flexibility or reflex versus reflect tradeoff is at the heart of adaptation, is the next important tradeoff after efficiency and robustness, and is at the heart of layered architectures. The lower layer VOR system is a reflex, and operates so effectively and rapidly that we are unaware of it until something makes us dizzy and we realize it is there. We hope our research could further explain the tradeoffs in axon size and density versus delay and bandwidth, and the impact on sensorimotor performance. While challenging, the theory we have illustrated here and will further explore is essential to make sense of these tradeoffs.

References

- [1] N. P. Reeves, P. Pathak, J. M. Popovich, and V. Vijayanagar, “Limits in motor control bandwidth during stick balancing,” *Journal of Neurophysiology*, vol. 109, no. 10, pp. 2523–2527, 2013. [Online]. Available: <http://jn.physiology.org/content/109/10/2523.abstract>
- [2] J. Milton, J. L. Cabrera, T. Ohira, S. Tajima, Y. Tonosaki, C. W. Eurich, and S. A. Campbell, “The time-delayed inverted pendulum: Implications for human balance control,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 19, no. 2, pp. 026 110–026 110–12, Jun. 2009. [Online]. Available: http://chaos.aip.org/resource/1/chaoh/v19/i2/p026110_s1
- [3] J. G. Milton, “The delayed and noisy nervous system: implications for neural control,” *Journal of Neural Engineering*, vol. 8, no. 6, p. 065005, Dec. 2011. [Online]. Available: <http://iopscience.iop.org/1741-2552/8/6/065005>
- [4] J. L. Cabrera and J. Milton, “Human stick balancing: Tuning lévy flights to improve balance control,” 2004.
- [5] I. D. Loram and M. Lakie, “Human balancing of an inverted pendulum: position control by small, ballistic-like, throw and catch movements,” *The Journal of Physiology*, vol. 540, no. 3, pp. 1111–1124, 2002.
- [6] T. Cluff, M. A. Riley, and R. Balasubramaniam, “Dynamical structure of hand trajectories during pole balancing,” *Neuroscience Letters*, vol. 464, no. 2, pp. 88 – 92, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0304394009011197>
- [7] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory*. Macmillan Publishing Co., 1990.
- [8] P. Lennie, “The physiological basis of variations in visual latency,” *Vision Research*, vol. 21, no. 6, pp. 815–824, 1981. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0042698981901802>
- [9] R. Nijhawan, “Visual prediction: Psychophysics and neurophysiology of compensation for time delays,” *Behavioral and Brain Sciences*, vol. 31, no. 02, pp. 179–198, 2008.
- [10] P. Cavanagh and P. Komi, “Electromechanical delay in human skeletal muscle under concentric and eccentric contractions,” *European Journal of Applied Physiology and Occupational Physiology*, vol. 42, no. 3, pp. 159–163, 1979. [Online]. Available: <http://dx.doi.org/10.1007/BF00431022>
- [11] J. R. Lackner and P. DiZio, “Vestibular, proprioceptive, and haptic contributions to spatial orientation,” *Annu. Rev. Psychol.*, vol. 56, pp. 115–147, 2005.
- [12] H. Collewijn and J. B. J. Smeets, “Early components of the human vestibulo-ocular response to head rotation: Latency and gain,” *Journal of Neurophysiology*, vol. 84, no. 1, pp. 376–389, 2000. [Online]. Available: <http://jn.physiology.org/content/84/1/376.abstract>