

e - c o m p a n i o n

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—“Maintaining Diagnostic Knowledge-Based Systems: A Control-Theoretic Approach” by Alain Bensoussan, Radha Mookerjee, Vijay Mookerjee, Wei T. Yue, *Management Science*, DOI 10.1287/mnsc.1080.0908.

E-companion to: Maintaining Diagnostic Knowledge-based Systems: A Control Theoretic Approach

Alain Bensoussan, Radha Mookerjee, Vijay Mookerjee, Wei T. Yue

School of Management, The University of Texas at Dallas, Richardson, Texas 75083, USA.

(alain.bensoussan, radham, vijaym, wei.yue)@utdallas.edu

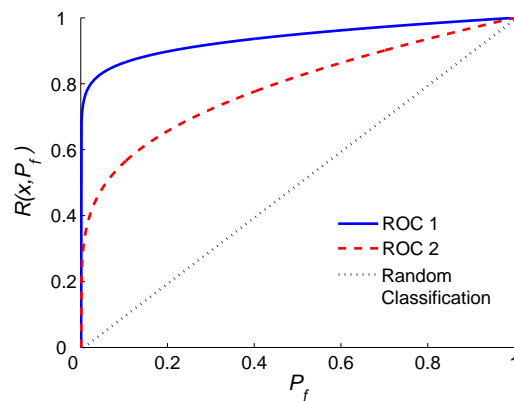
Appendix A: Performance Measurement and Maintenance

A.1. Performance Measurement Using ROC Curves

There is a rich body of work on the subject of measuring the discrimination ability of a diagnostic system. Most of early work has been on measuring the discrimination ability of a medical diagnostic system. Swets (1988) introduced a quantitative measure of discrimination ability as the area under the ROC curve. Since the endpoints of the ROC curve are usually assumed to be (0,0) and (1,1), the maximum possible area under the curve is 1. Also, because the ROC curve must represent a system whose discrimination ability is at least as good as a random classifier, the minimum possible area under the curve is 0.5. The area under the curve (AUC) is preferred as an discrimination ability measure instead of simple point-measures (such as the rates of the two errors) because the area incorporates both error rates and the inherent trade-off between the two kinds of error (Obuchowski 1997). However, one problem with the AUC measure is that two curves with different characteristics may possess the same area (Venkatraman and Begg 1996). Since the utilities of two curves with the same area may be distinguishable from a decision-theoretic perspective (depending on the cost of the two errors), it has been proposed to check the equality of the entire curve in order to compare two diagnostic systems (Campbell 1994). A ROC curve that is more skewed more to the upper left corner corresponds to a system of higher discrimination ability because a higher true positive rate can be obtained for the same false positive rate. For instance, in Figure A-1, ROC 1 represents a system of higher discrimination ability.

A.2. Maintenance Methods

Like most systems, diagnostic systems also deteriorate in performance after they have been constructed. The discrimination ability of the system may fall from natural causes such as changes in the concepts over time. The diagnostic system should be able to track such changes and adapt to them relatively quickly. This problem, known in the literature as concept drift (or the related notion of population drift), can lead initially to a high rate of deterioration (Kelly et al. 1999, Widmer and Kubat 1996) that appears to saturate later (Wang et al. 2003). In addition to concept drift that occurs from natural causes, in some diagnostic

Figure A-1 Three ROC Curves with Different Accuracy.

applications such as spam detection, intrusion detection, and surveillance, the input data used for diagnosis may be actively modified by an adversary causing system discrimination ability to fall. The challenge for adversarial learning is to be robust to a noisy diagnostic system or a system that is frequently updated (Lowd and Meek 2005).

Because diagnostic systems can be costly to modify once they are built, it is necessary to build the system in a manner that facilitates modification. Jacob and Froscher (1990) provide a software engineering methodology to construct a rule-based diagnostic system. The basic approach here is to make the system more modular by dividing the system knowledge into similar groups and carefully specifying the flow of information between these groups. Other similar approaches include using “Well Structured Rules” Higa (1990) or using better knowledge representation schemes (Lee and O’Keefe 1996) so that system maintenance is facilitated. Outside of the software engineering approaches to maintaining a diagnostic system, there have been many studies that deal with the verification and validation of newly updated system knowledge (Liu and Pierce 1997, O’Leary 1996). A systems engineering perspective to maintenance is provided by (Kokar et al. 1999), where a feedback control architecture is presented.

To reduce the cost of human effort involved in maintenance, machine learning techniques have been applied to partly automate the maintenance of diagnostic systems (Black and Hickey 1999). Widmer and Kubat (1996) present a method for learning in changing environments (i.e., with concept drift) as well as with noisy input data. The goal of a learner is to: (1) quickly adapt to concept drift, (2) distinguish between noise and a true change in the concept, and (3) adapt to recurring and seasonal changes in concepts. Another strategy to deal with concept drift is to “forget” (or delete) old concepts and instead, generate new system knowledge from recent data (Maloof and Michalski 2004). Although the maintenance of a diagnostic system is

supported by the use of machine learning techniques, researchers acknowledge that the complete automation of diagnostic system maintenance may never be possible (Dalvi et al. 2004). As a result, human effort is exerted to repeatedly bring a system back to its desired performance level (Lowd and Meek 2005, Dalvi et al. 2004). The resulting costs can be high, often significantly exceeding the initial cost to develop the system (Gill 1995, Harmon and Sawyer 1990). However, there has been little research to report on the optimal extent of maintenance effort to keep the performance of a diagnostic system at its desired level keeping in mind the cost of human intervention.

References

- Black, M., R. J. Hickey. 1999. Maintaining the performance of a learned classifier under concept drift. *Intelligent Data Analysis* **3**(6) 453–474.
- Campbell, G. 1994. Advances in statistical methodology for the evaluation of diagnostic and laboratory tests. *Statistics in Medicine* **13** 499–508.
- Dalvi, N., P. Domingos, S. Sanghai Mausam, D. Verma. 2004. Adversarial classification. *In KDD'04: Proceedings of the 2004 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 99–108.
- Gill, T. G. 1995. Early expert systems: Where are they now? *MIS Quarterly* **19**(1) 51–81.
- Harmon, P., B. Sawyer. 1990. *Creating Expert Systems for Business and Industry*. John Wiley & Sons, New York, NY.
- Higa, Kunihiro. 1990. Methodical design and maintenance of well structured rule base. Elias M. Awad, ed., *Proceedings of the 1990 ACM SIGBDP Conference on Trends and Directions in Expert Systems*. 377–395.
- Jacob, R. J. K., J. N. Froscher. 1990. A software engineering methodology for rule-based systems. *IEEE Transactions on Knowledge and Data Engineering* **2**(2) 173–189.
- Kelly, M. G., D. J. Hand, N. M. Adams. 1999. The impact of changing populations on classifier performance. *In KDD'99: Proceedings of the 1999 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 367–371.
- Kokar, M. M., K. Baclawski, Y. A. Eracar. 1999. Control theory-based foundations of self-controlling software. *IEEE Intelligence Systems* **14**(3) 37–45.

- Lee, S., R. M. O’Keefe. 1996. The effect of knowledge representation schemes on maintainability of knowledge-based systems. *IEEE Transactions on Knowledge and Data Engineering* **8**(1) 173–178.
- Liu, C., D. Pierce. 1997. Verification and validation of artificial intelligent systems with applications to power systems. *IEE Colloquium on Artificial Intelligence Techniques in Power Systems*. 2/1 – 2/4.
- Lowd, D., C. Meek. 2005. Adversarial learning. In *KDD’05: Proceedings of the 2005 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 641–647.
- Maloof, M. A., R. S. Michalski. 2004. Incremental learning with partial instance memory. *Artificial Intelligence* **154** 95–126.
- Obuchowski, N. A. 1997. Nonparametric analysis of clustered roc curve data. *Biometrics* **53**(2) 567–578.
- O’Leary, D. E. 1996. Verification of uncertain knowledge-based systems: An empirical verification approach. *Management Sci.* **42**(12) 1663–1675.
- Swets, J. A. 1988. Measuring the discrimination ability of diagnostic systems. *Science* **240**(4857) 1285–1293.
- Venkatraman, E. S., C. B. Begg. 1996. A distribution-free procedure for comparing receiver operating characteristic curves from a paired experiment. *Biometrika* **83**(4) 835–848.
- Wang, H., W. Fan, P. S. Yu, J. Han. 2003. Mining concept-drifting data streams using ensemble classifiers. In *KDD’03: Proceedings of the 2003 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 24–27.
- Widmer, G., M. Kubat. 1996. Learning in the presence of concept drift and hidden contexts. *Machine Learning* **23**(1) 69–101.

Appendix B: Non-Linear, Time Dependent

In general, the optimal control should take the following form

$$u(t) = \begin{cases} 1 & z < 0, \\ u = \left[\left(1 - \frac{z}{\mu k} \right) \right]^{\frac{1}{\mu-1}} & 0 \leq z \leq \mu k, \\ 0 & z > \mu k. \end{cases}$$

A consolidated form of the above optimal control is

$$u = \left[\left(1 - \frac{z^+}{\mu k} \right)^+ \right]^{\frac{1}{\mu-1}}$$

where $z = \mu k + \lambda \alpha$. Upon analysis, however, it can be shown that for this problem, the control $u = 0$ can be ruled out. The structure of the optimal policy is characterized below

$$u^*(t) = \begin{cases} 1 & 0 < t < \tau, \\ \left(1 - \frac{z^+}{\mu k} \right)^{\frac{1}{\mu-1}} & \tau < t < T \end{cases}$$

where $z = k\mu + \lambda \alpha$, and τ is the solution to $z(\tau) = 0$.

Proof: The Hamiltonian is given by $H = f(x) + ku^\mu + \lambda(u\alpha(x, t) - \beta(x, t))$, and $H_u = \mu \cdot k \cdot u^{\mu-1} + \lambda \alpha$.

Therefore

$$\mu k u^{\mu-1} + \lambda \alpha < 0, \quad u = 1 \tag{A-1}$$

$$\mu k u^{\mu-1} + \lambda \alpha = 0, \quad 0 < u < 1 \tag{A-2}$$

$$\mu k u^{\mu-1} + \lambda \alpha > 0, \quad u = 0 \tag{A-3}$$

Let $z = \mu k + \lambda \alpha$. Then it can be seen that

$$u(t) = \begin{cases} 1 & z < 0, \\ u = \left[\left(1 - \frac{z}{\mu k} \right) \right]^{\frac{1}{\mu-1}} & 0 \leq z \leq \mu k, \\ 0 & z > \mu k. \end{cases} \tag{A-4}$$

The corner solutions ($u = 1, u = 0$) result from substituting the appropriate value of z in (A-1) and (A-3). When $0 < z < \mu k$ substituting $u = 0$ in (A-2) results in an inconsistency. Similarly, substituting $u = 1$ in (A-2) also results in an inconsistency, leading us to the interior solution in (A-4). The value of u for any value of z can then be written as

$$u = \left[\left(1 - \frac{z^+}{\mu k} \right)^+ \right]^{\frac{1}{\mu-1}} \tag{A-5}$$

The system under the optimal control is described by the following conditions

$$\dot{x} = -\beta + \alpha \cdot \left[\left(1 - \frac{z^+}{\mu k} \right)^+ \right]^{\frac{1}{\mu-1}},$$

$$\begin{aligned}
x(0) &= x_0, \\
-\dot{z} + \left(r + \beta_x - \beta \frac{\alpha}{\alpha_x} + \frac{\alpha_t}{\alpha} \right) z &= h(x), \\
z(T) &= k\mu
\end{aligned}$$

where $h(x) = \alpha f'(x) + \mu k \left(r + \beta_x - \beta \frac{\alpha}{\alpha_x} + \frac{\alpha_t}{\alpha} \right)$

We first write the state equation with the u as defined in (A-5). Next, because $z = \mu k + \lambda \alpha$, differentiating we get, $\dot{z} = \dot{\lambda} \alpha + \lambda (\alpha_x \dot{x} + \alpha_t)$. In this equation, we substitute for $\dot{\lambda} = \lambda r - H_x$, where H_x is the partial derivative of the Hamiltonian with respect to the state variable. We also substitute for $\dot{x} = \alpha u - \beta$ using the state equation. With some algebra we obtain

$$\dot{z} = -\alpha f_x + (z - k\mu) \left[r + \beta_x - \beta \frac{\alpha_x}{\alpha} + \frac{\alpha_t}{\alpha} \right]$$

If we set $\gamma = r + \beta_x - \beta \frac{\alpha_x}{\alpha} + \frac{\alpha_t}{\alpha}$ and $w = z - k\mu$, we get

$$-\dot{w} + w\gamma = h - \mu k \gamma = \alpha f_x \leq 0$$

Since $w(T) = 0$, it implies that $w \leq 0$ on $(0, T)$. Thus $z \leq \mu k$ and hence the control $u = 0$ can be ruled out. ■

Appendix C: Non-Linear, Time Independent

C.1. Steady State Analysis

Here we look for steady state solutions. We denote the steady state values of the control, state, and the adjoint variable as \hat{u} , \hat{x} and $\hat{\lambda}$ respectively. We present this case for an infinite planning horizon.

Using $z = k\mu + \lambda\alpha$, we get

$$\dot{z} = \dot{\lambda}\alpha + \lambda\alpha_x\dot{x}$$

The adjoint equation provides us with,

$$-\dot{\lambda} + \lambda r = H_x$$

Using $\dot{z} = 0$ we get,

$$-\alpha f'(x) + \lambda(\alpha r + \alpha\beta x - \beta\alpha_x) = 0$$

Substituting $\lambda = \frac{z-k\mu}{\alpha}$ and $z = k\mu \left[1 - \left(\frac{\beta}{\alpha}\right)^{\mu-1}\right]$, we get an equation that can be solved to yield the steady state value of the system discrimination ability:

$$\alpha^{\mu+1} f'(x) + k\mu\beta^{\mu-1}(\alpha r + \alpha\beta x - \beta\alpha_x) = 0 \quad (\text{A-6})$$

The discrimination ability that solves the above equation (\hat{x}) is constant, i.e., $\dot{x} = 0$. Using the state equation, we therefore find that the control exerted in steady state is also a constant and is given by the expression below.

$$\hat{u} = \frac{\beta(\hat{x})}{\alpha(\hat{x})} \quad (\text{A-7})$$

The value of the adjoint variable in steady state ($\hat{\lambda}$) can be obtained by substituting the value of \hat{x} in the equation $z = 0$.

$$\hat{\lambda} = -\frac{k\mu}{\alpha(\hat{x})} \quad (\text{A-8})$$

C.2. Phase Diagram Analysis

The phase diagram differential equation, obtained by dividing \dot{z} by \dot{x} , is given by

$$\frac{dz}{dx} = \frac{\left(r + \beta_x - \beta \frac{\alpha}{\alpha_x}\right) z - h(x)}{-\beta + \alpha \cdot \left[\left(1 - \frac{z^+}{\mu k}\right)^+\right]^{\frac{1}{\mu-1}}} \quad (\text{A-9})$$

with the initial condition $z(\hat{x}) = \hat{z}$. Because $\frac{dz}{dx}$ at $x = \hat{x}$ is indeterminate (a zero by zero form), we find the value of $z'(\hat{x})$ by asymptotic expansion.

First we let

$$N(\hat{x}) = \left(r + \beta_x - \beta \frac{\alpha}{\alpha_x}\right) \hat{z} - h(\hat{x})$$

$$D(\hat{x}) = -\beta + \alpha \cdot \left[\left(1 - \frac{\hat{z}^+}{\mu k}\right)^+\right]^{\frac{1}{\mu-1}}$$

Then, $\frac{dz}{dx}$ can be written as

$$\frac{N(\hat{x} + \epsilon)}{D(\hat{x} + \epsilon)} \sim \frac{N(\hat{x}) + \epsilon N'(\hat{x})}{D(\hat{x}) + \epsilon D'(\hat{x})} = \frac{N'(\hat{x})}{D'(\hat{x})}$$

because $N(\hat{x}) = D(\hat{x}) = 0$.

Given that

$$N'(\hat{x}) = \hat{z} \left(\beta_{xx} - \beta_x \frac{\alpha}{\alpha_x} - \beta \frac{\alpha_x^2 - \alpha \alpha_{xx}}{\alpha_x^2} \right) + z'(\hat{x}) \left(r + \beta_x - \beta \frac{\alpha}{\alpha_x} \right) - h'(\hat{x})$$

$$D'(\hat{x}) = -\beta_x - \frac{\alpha z'(\hat{x})}{\mu k (\mu - 1)} \left(1 - \frac{\hat{z}}{\mu k} \right)^{\frac{1}{\mu-1} - 1} + \alpha_x \left(1 - \frac{\hat{z}}{\mu k} \right)^{\frac{1}{\mu-1}}$$

$z'(\hat{x})$ can therefore be obtained from the quadratic equation that is given by

$$z'(\hat{x}) = \frac{N'(\hat{x})}{D'(\hat{x})}$$

From $z'(\hat{x})$, we can then obtain $z(\hat{x} + \epsilon) = \hat{z} + \epsilon z'(\hat{x})$. We can also compute $x = \hat{x} + \epsilon$, which allows us to obtain the value of $z'(x + \epsilon)$ using the general form of $z'(x)$ (as shown in (A-9)). We then iteratively compute the value of $z(x_0)$ and obtain the value of the optimal control (u) between $x = x_0$ and $x = \hat{x}$.

■

C.3. Optimal Policy

The state of the system under the optimal policy can be characterized as below

(a) $x_0 < \hat{x}$ implies

$$\dot{x}(t) > 0 \quad \forall x(t) < \hat{x}$$

$$\dot{u}(t) < 0 \quad \forall x(t) < \hat{x}$$

(b) $x_0 > \hat{x}$ implies

$$\dot{x}(t) < 0 \quad \forall x(t) < \hat{x}$$

$$\dot{u}(t) > 0 \quad \forall x(t) < \hat{x}$$

Proof: The proof consists of the following steps.

1. $\frac{dz}{dx}$ has a constant sign.

For the differential equation (A-9) to have a smooth solution we need that the denominator have only one zero, necessarily at \hat{x} . That is,

$$\begin{aligned} \left(\left(1 - \frac{z^+}{\mu k} \right)^+ \right)^{\frac{1}{\mu-1}} &= \frac{\beta(\hat{x})}{\alpha(\hat{x})} \\ \Rightarrow z(x) = \mu k + \frac{\alpha(x)f'(x)}{\left(r + \beta_x - \beta \frac{\alpha}{\alpha_x} \right)} &\Rightarrow x = \hat{x}. \end{aligned}$$

Conversely, if the numerator is zero, \hat{x} is uniquely defined, but $1 - \frac{z(x)}{\mu k} = \left(\frac{\beta(x)}{\alpha(x)} \right)^{\mu-1}$ satisfies this condition.

Hence the numerator and denominator can vanish only at point \hat{x} . Therefore, $\frac{dz}{dx}$ has a constant sign.

2. From 1. we have $z'(\hat{x}) > 0 \Rightarrow z'(x) > 0 \forall x$.
3. The denominator of (A-9), \dot{x} is decreasing in x (from 2.).
4. $x_0 < \hat{x} \Rightarrow \dot{x} > 0$. (from 3. and $\dot{x} = 0$ when $x = \hat{x}$).
5. From 2. and (A-4), $x_0 < \hat{x} \Rightarrow u$ is decreasing in x .
6. Using 4. and 5., $x_0 < \hat{x}$ implies u is decreasing in time.

Similarly, when $x_0 > \hat{x}$ is true, it can be shown that $\dot{x} < 0$ and $\dot{u} > 0$.

■

Appendix D: Linear, Time Independent

D.1. Steady State Conditions

Define $z = k + \lambda\alpha$. Differentiating z with respect to t yields,

$$\dot{z} = \dot{\lambda}\alpha - \lambda\alpha_x\dot{x}.$$

Substituting in $\dot{x} = \alpha u - \beta$ and $\dot{\lambda} = \lambda r - H_x$, where H_x is the partial derivative of the Hamiltonian with respect to x , given by $H_x = f'(x) + \lambda(\alpha_x u - \beta_x)$. Setting $H_u = z = 0$, and $\dot{z} = 0$, we get obtain an equation that can be solved for the steady state level of discrimination ability \hat{x}

$$\alpha f'(x) = -k(r + \beta_x - \alpha_x \frac{\beta}{\alpha}).$$

Solving for the steady state value \hat{x} , we can use the fact that $\dot{x} = 0$ to obtain $\hat{u} = \frac{\alpha(\hat{x})}{\beta(\hat{x})}$ and $z = 0$ to obtain $\hat{\lambda} = \frac{-k}{\alpha(\hat{x})}$.

D.2. Necessary Conditions

Define $z = k + \lambda\alpha$. Differentiating z with respect to t yields,

$$\dot{z} = \dot{\lambda}\alpha - \lambda\alpha_x\dot{x}.$$

Substituting in $\dot{x} = \alpha u - \beta$ and $\dot{\lambda} = \lambda r - H_x$, where H_x is the partial derivative of the Hamiltonian with respect to x , given by $H_x = f'(x) + \lambda(\alpha_x u - \beta_x)$. Also, given that $\lambda = \frac{z-k}{\alpha}$, the above equation becomes

$$\dot{z} = -\alpha f'(x) + (z - k)(r + \beta_x - \alpha_x \frac{\beta}{\alpha}).$$

Let $h(x) = (\alpha)^2 f'(x) + k(\alpha r + \alpha\beta_x - \alpha_x\beta)$. Differentiating $h(x)$, we get

$$h'(x) = 2\alpha\alpha_x f'(x) + (\alpha)^2 f''(x) + k(\alpha_x r + \alpha\beta_{xx} - \alpha_{xx}\beta)$$

Requiring that $h'(x) > 0$, we write

$$\dot{z} = -\frac{h(x)}{\alpha} + \frac{\alpha r + \alpha\beta_x - \alpha_x\beta}{\alpha} \cdot z \quad (\text{A-10})$$

i) Suppose $x_0 < \hat{x}$, we look for the optimal strategy as follow

$$u(t) = \begin{cases} 1 & 0 < t < \tau_0, \\ \hat{u} & \tau_0 < t < \theta, \\ 0 & \theta < t < T. \end{cases}$$

In this case, $\dot{x} = \alpha - \beta$, where $x(0) = x_0$ and $x(\tau_0) = 0$. We need to establish that $z(t) < 0$ for $t < \tau_0$, but $z(\tau_0) = 0$. Hence, we have to solve (A-10) between t and T . We get

$$z(t) = \int_t^{\tau_0} \frac{h(x(s))}{\alpha(s)} e^{-\int_s^t \frac{\omega(\sigma)}{\alpha(\sigma)} d\sigma} ds. \quad (\text{A-11})$$

where $\omega(x) = \alpha r + \alpha \beta_x - \alpha_x \beta > 0$. Since $x(s) < \hat{x}$ for $s < \tau_0$, it is enough to assume that $h(x) < 0$ for $x < \hat{x}$.

When $u = 0$, we know $\dot{x} = -\beta(x)$, for $\theta < t < T$, $z(\theta) = 0$, and $z(T) = km$. We need to check that $z(t) > 0$ for $\theta < t < T$. We write

$$z(t) = - \int_{\theta}^t \frac{h(x(s))}{\alpha(s)} e^{-\int_s^t \frac{\omega(\sigma)}{\alpha(\sigma)} d\sigma} ds.. \quad (\text{A-12})$$

For $s > \theta$, $x(s) < \hat{x}$, hence, $h(x(s)) < h(\hat{x})$. Therefore $z(t) > 0$.

ii) In the case which $x_0 > \hat{x}$, given that $h'(x) > 0$, we know that $h(x) > 0$. Therefore, from (A-11) we learn that $z(t) > 0$.

■

D.3. Sufficiency Conditions

We limit ourselves to the infinite horizon version of the problem. To prove that the control is optimal in the absence of convexity assumptions, we rely on DP arguments. It is well known that if the value function is smooth, the corresponding feedback leads to an optimal solution. The smoothness of the solution of the Bellman equation is an interesting issue as such.

The Bellman equation is the following:

$$-rV + \min_{0 \leq u \leq 1} [V_x(\alpha u - \beta) + f(x) + ku] = 0,$$

which can be further developed into

$$-rV - V_x \beta + f(x) + \min_{0 \leq u \leq 1} [(V_x \alpha + k)u] = 0$$

Differentiating with respect to x , we can write

$$\lambda' \cdot (\alpha - \beta) + f'(x) - (r + \beta_x - \alpha_x)\lambda = 0$$

where $\lambda(x) = V_x$. When $k + V_x \alpha < 0$, it implies $x < \hat{x}$. When $k + V_x \alpha > 0$, minimization yields zero, hence

$$-\lambda' \cdot \beta + f'(x) - (\beta_x + r)\lambda = 0$$

it implies $x > \hat{x}$.

In \hat{x} , $k + \hat{\lambda} \alpha(\hat{x}) = 0$,

$$\frac{f'(\hat{x}) - (\beta_x(\hat{x}) + r)\hat{\lambda}}{\beta(\hat{x})} = \frac{f'(\hat{x}) - (r + \beta_x - \alpha)\hat{\lambda}}{\alpha(\hat{x}) - \beta(\hat{x})}$$

Let $z = \lambda\alpha + k$, hence $z' = \lambda'\alpha + \lambda\alpha_x$. Let $z' = -h(x) + za(x)$, where $h(x) = (\alpha)^2 f'(x) + k(\alpha r + \alpha\beta_x - \alpha_x\beta)$, and $a(x) = \frac{\beta_x + r\alpha/\alpha_x}{(\alpha - \beta)\alpha}$. Then,

$$\begin{aligned} \frac{d}{dx}(ze^{-\int_0^x a(\xi)d\xi}) &= (z' - az)e^{-\int_0^x a(\xi)d\xi} \\ &= -h(x)e^{-\int_0^x a(\xi)d\xi} \end{aligned}$$

We then write

$$-ze^{-\int_0^x a(\xi)d\xi} = -\int_x^{\hat{x}} h(\eta)e^{-\int_0^\eta a(\xi)d\xi}$$

and

$$z = \int_x^{\hat{x}} h(\eta)e^{-\int_x^\eta a(\xi)d\xi}$$

To get $z < 0$, $h(x) < 0$ for $x < \hat{x}$. We can use similar logic to get $z > 0$, which means $h(x) > 0$ for $x > \hat{x}$. We are able to make the proofs by assuming $h(x) < 0$, when $x < \hat{x}$, and $h(x) > 0$ when $x > \hat{x}$, the costate of DP is smooth, V is C^2 . The optimal feedback is simply

$$u(x) = \begin{cases} 1 & x < \hat{x} \\ 0 & x > \hat{x} \end{cases} \quad (\text{A-13})$$

With such a feedback and smoothness of V , it is easy to check that

$$V(x) = J_x(u(\cdot))$$

where $J_x(u(\cdot))$ is the cost function corresponding to the feedback rule $u(x)$ and starting from an initial point x . This is sufficient to prove that the control is optimal. ■

Appendix E: Discrete Model

E.1. Optimal Policy Structure

Consider the state equation for the continuous model $\dot{x} = c_0u(1-x) - c_1x$. In the case of the discrete model, when U_n is applied at θ_n , there is an increase in discrimination ability during an impulse. Let us assume that the application of effort U_n leads to an improvement in system discrimination ability from $x(\theta_n^-)$ to

$$x(\theta_n) = 1 - e^{-c_0U_n}(1 - x(\theta_n^-))$$

The deterioration in discrimination ability then occurs in $\theta_n \leq t \leq \theta_{n+1}$ at the rate of

$$\dot{x} = -c_1x(t).$$

Let $I = \{(\theta_n, U_n), \dots\}$ be the impulse control policy. Then, the cost function for the discrete model can be written as

$$J_x(I) = \int_0^\infty e^{-rt} f(x(t)) dt + \sum_{n=1}^\infty (kU_n + K)e^{-r\theta_n}$$

where a fixed cost of K is incurred in each maintenance. Let $V(x) = \inf J_x(I)$ satisfies the following DP inequalities

$$V(x) \leq K + \inf_{u \geq 0} \{kU + V(x)(1 - e^{-c_0U}(1-x))\} \quad (\text{A-14})$$

$$f(x) - rV(x) - c_1xV'(x) \geq 0 \quad (\text{A-15})$$

where (A-14) applies when maintenance occurs and (A-15) applies when no maintenance occurs. These two equations are governed by the complementary and slackness conditions given by

$$(f(x) - rV(x) - c_1xV'(x))(V(x) - K - \inf_{U \geq 0} \{kU + V(x)(1 - e^{-c_0U}(1-x))\}) = 0. \quad (\text{A-16})$$

We make the following transformation by first writing $y = 1 - e^{-c_0U}(1-x)$, then $y \geq x$ and we get

$$U = \frac{1}{c_0} \{\log(1-x) - \log(1-y)\}$$

(A-14) can be written as

$$V(x) \leq K + \inf_{y \geq x} \left\{ \frac{k}{c_0} \log(1-x) - \frac{k}{c_0} \log(1-y) + V(y) \right\},$$

It becomes

$$V(x) - \frac{k}{c_0} \log(1-x) \leq K + \inf_{y \geq x} \left\{ V(y) - \frac{k}{c_0} \log(1-y) \right\}.$$

Setting $z(x) = V(x) - \frac{k}{c_0} \log(1-x)$, (A-14), (A-15) and (A-16) become

$$\begin{aligned} z(x) &\leq K + \inf_{y \geq x} \{z(y)\}, \\ g(x) - rz(x) - c_1xz'(x) &\geq 0, \\ (g(x) - rz(x) - c_1xz'(x))(z(x) - k - \inf_{y \geq x} z(y)) &= 0. \end{aligned}$$

where $g(x) = f(x) - r \frac{k}{c_0} \log(1-x) + \frac{kc_1x}{c_0(1-x)}$.

A (s, S) policy is a solution for the transformed equations in the following manner

$$\begin{aligned} z(x) &= z(s), & x \leq s, \\ g(x) - rz(x) - c_1xz'(x) &= 0, & x \geq s, \end{aligned}$$

Note that the slackness condition $(g(x) - rz(x) - c_1xz'(x))(z(x) - z(s)) = 0$; is automatically satisfied. For a given s and a given value of $z(s)$, we can solve the differential equation in $(s, 1)$. That is

$$z_s(x) = x^{-\frac{x}{c_1}} \left[z_s(s) s^{\frac{x}{c_1}} + \frac{1}{c_1} \int_s^x g(\xi) \xi^{\frac{x}{c_1}-1} d\xi \right]$$

Let $S(s)$ be the smallest minimum of $z_s(x)$ that is larger than or equal to s

$$z_s(S) = \inf_{y \geq s} z_s(y)$$

We determine $S(s)$ by $z'_s(S) = 0$, i.e., $g(S) = rz_s(S)$. We still need to define $z_s(s)$ and s , which can be obtained by the two equations

$$\begin{aligned} z_s(s) &= K + z_s(S) \\ z'_s(s) &= 0 \end{aligned}$$

It is convenient to introduce $H_s(x) = z'(x)$, then we have

$$\begin{aligned} H_s(x) &= 0, & x \leq s, \\ g'(x) - (r + c_1)H_s(x) - c_1xH'_s(x) &= 0, & x \geq s, \end{aligned}$$

and we have the equations

$$\begin{aligned} H_s(S) &= 0 \\ -K &= \int_s^S H_s(x) dx \end{aligned}$$

Clearly $H_s(x) = \frac{x^{-\frac{r+c_1}{c_1}}}{c_1} \int_s^x g'(\xi) \xi^{\frac{r}{c_1}} d\xi$, we get

$$H_s(x) = \frac{r}{c_1 x} \left\{ g(x) - r x^{-\frac{r}{c_1}} \left[\frac{g(s)}{r} s^{\frac{r}{c_1}} + \frac{1}{c_1} \int_s^x g(\xi) \xi^{\frac{r}{c_1}-1} d\xi \right] \right\}$$

$$z_s(x) = \left(\frac{g(s)}{r} s^{\frac{r}{c_1}} + \frac{1}{c_1} \int_s^x g(\xi) \xi^{\frac{r}{c_1}-1} d\xi \right) x^{-\frac{r}{c_1}}$$

which leads to the following two equations that solve for the optimal (s, S) policy

$$g(S)S^{\frac{r}{c_1}} = g(s)s^{\frac{r}{c_1}} + \frac{r}{c_1} \int_s^S g(\xi) \xi^{\frac{r}{c_1}-1} d\xi$$

$$-Kr = g(S) - g(s)$$

From the first and second order conditions,

$$g'(x) = f'(x) + \frac{rk}{c_0} \frac{r}{1-x} + k \frac{c_1}{c_0} \frac{r}{(1-x)^2}$$

$$g''(x) = f''(x) + \frac{rk}{c_0} \frac{r}{(1-x)^2} + 2k \frac{c_1}{c_0} \frac{r}{(1-x)^3} > 0 \quad \text{if } f''(x) > 0$$

therefore g is a convex function.

Let a be the minimum of g . Since $g(S) < g(s)$, thus $s < a < S$. We must have

$$\int_s^S g'(\xi) \xi^{\frac{r}{c_1}} d\xi = 0$$

We can observe that the function $\Phi_s(S) = \int_s^S g'(\xi) \xi^{\frac{r}{c_1}} d\xi$ is increasing from $\int_s^a g'(\xi) \xi^{\frac{r}{c_1}} d\xi$ to $\int_s^1 g'(\xi) \xi^{\frac{r}{c_1}} d\xi = +\infty$.

Hence for any $s \leq a$ that exists, a unique $S(s)$ satisfies the first equation.

Differentiating in s the relation $\int_s^{S(s)} g'(\xi) \xi^{\frac{r}{c_1}} d\xi = 0$, we get

$$S'(s)g'(S)S^{\frac{r}{c_1}} - g'(s)s^{\frac{r}{c_1}} = 0.$$

Hence $S'(s) < 0$. The function $\Psi(s) = g(S(s)) - g(s)$ satisfies

$$\Psi'(s) = g'(S(s))S'(s) - g'(s)$$

$$= g'(s) \frac{s^{\frac{r}{c_1}}}{S^{\frac{r}{c_1}}} - g'(s) > 0$$

$\Psi(a) = 0$ assume $\int_0^a g'(\xi) \xi^{\frac{r}{c_1}} d\xi < \infty$ and $f(0)$ is sufficiently large, then $S(0) < 1$ and $g(S(0)) < \infty$. Hence

$\Psi(0) = -\infty$. Therefore, there exists a unique s solution of the second equation. ■

E.2. Equivalence of State Space and Time Space Policies

The maintenance policy described by Proposition 2 can be transformed to a periodic maintenance policy (Ω, U) , where the effort U is applied every Ω time units, given by

$$\Omega = -\frac{1}{c_1} \ln \frac{s}{S} \quad U = \frac{1}{c_0} \ln \frac{1-s}{1-S}$$

Proof: The discrimination ability before the impulse (U) is s and the discrimination ability after is S . Therefore,

$$S = 1 - (1-s)e^{-c_0 U}$$

which leads to

$$U = \frac{1}{c_0} \ln \frac{1-s}{1-S}$$

The discrimination ability drifts from S to s following the differential equation $\dot{x} = -c_1 x$. This can be solved as

$$\ln(x) = -c_1 t + \ln(S)$$

Therefore, the time taken Ω for the discrimination ability to drift to s is given by

$$\Omega = \frac{1}{c_1} \ln \frac{S}{s}$$

■

Table 5 Dataset Description.

Dataset	Number of Attributes	Number of Records
Intrusion Detection	41	6000
Breast Cancer	9	698
IPUMS Census Data	8	768
PIMA Diabetes	60	2100

Appendix F: Applications of ROC Curves

We used four datasets to test our assumptions about the ROC curve. The datasets came from four different domains, namely, intrusion detection, breast cancer diagnosis, diabetes diagnosis, and IPUMS census data and were obtained from the UCI machine-learning repository. We randomly selected a subset of the data from the IDS dataset and the IPUMS Census Data dataset. In all other cases, we used the complete dataset. All datasets had two classes, malicious and benign. In the IDS dataset, the two classes were obtained by combining different types of malicious classes into one malicious class. Table 5 summarizes the datasets used in this experiment.

The IDS dataset with 41 attributes is pre-split into training data consisting of 5 million records and testing data consisting of 1 million records. The training data was collected from 7 weeks of simulated network traffic and the testing data was collected in the following two-week period. Though there are 24 attack classes in the training set and 14 attack classes in the testing set we combined them into a single “attack” class.

The breast cancer dataset consists of clinical cases reported by the University of Wisconsin Hospital and contains 698 records with 9 attributes and two classes (benign and malignant).

The Pima Indians Diabetes dataset consists of 768 instances with 8 attributes and 2 classes (positive or negative). The data was collected from adult females of Pima Indian heritage by the Applied Physics Laboratory of The John Hopkins University.

The Integrated Public Use Microdata Series (IPUMS) census data consists of 60 attributes and describes households and individuals in the Los Angeles - Long Beach area. The data is provided by University of Minnesota. The data was collected for the years 1970, 1980, and 1990 and contains approximately 7000 records per year. We used one of the attributes, the “ownership of dwelling,” to be the class variable (rent or own).

Procedure

All four datasets were split into training and testing data using three different training to test ratios (50-50,

65-35, and 80-20). 10 randomly selected samples were generated for each split resulting in a total of 3×10^4 experiments.

We used the C4.5 decision tree induction technique (Quinlan 1993) to be the classifier. Default settings were used for the command to generate the decision tree as well as the command to generate the production rules. There are $i = 1 \dots k$ number of rules generated, where each rule results in a positive or negative classification. We then revisit the training data to identify the number of positive and negative events for each rule. We denote the number of negative events and positive events associated with a negative rule to be n_i and m_i ; and the number of negative events and positive events associated with a positive rule to be f_i and d_i . The true positive (detection) and false positive (false alarm) rates can then be computed by

$$P_d = \frac{\sum_{i=1}^k d_i}{\sum_{i=1}^k d_i + \sum_{i=1}^k m_i} \quad \text{and} \quad P_f = \frac{\sum_{i=1}^k f_i}{\sum_{i=1}^k f_i + \sum_{i=1}^k n_i}$$

To generate the ROC curve, which consists of multiple (P_d, P_f) pairs, we perturb the costs from the C4.5 classification results and reevaluate the classification decision with each perturbed cost. To do so, we first define the expected cost for an event in the dataset to be

$$(1 - p_b)(c_m + c_d) + p_b(c_f + c_n)$$

For an event to be classified as negative, the expected cost of a negative (benign) event has to be greater than the expected cost of a positive (malicious) event. Let us denote p_b^c to be the threshold at which the expected benign costs are larger. From the above equation, we derive the condition that a rule gives a negative classification when

$$p_b^c \geq \frac{1}{1 + \psi},$$

where $\psi = \frac{c_f + c_n}{c_m + c_d}$. We first randomize c_f , c_n , c_d , and then select different c_m values such that the value of ψ varies over a wide range of values. For each perturbation, we record the P_d and P_f values. We then rearrange the ROC equation to $P_d - P_f = x(P_f^a - P_f)$. Because there are two unknown variables a and x , we first select an a value by trial and error. The a value is chosen such that it results in good R^2 values for the different training and testing data splits in a domain. Once the a value is fixed, we run linear regression to obtain the x values.

All four datasets were found to fit well with (7) (see Table 6). The R^2 values ranges from 0.88 to 0.94. It has also observed that a is less than 0.5 in all four datasets.

Table 6 ROC Applications.

Domain	Average R^2 value	Average Accuracy	Domain Parameter a
Intrusion Detection	0.89	0.908	0.02
Breast Cancer	0.94	0.904	0.01
IPUMS Census Data	0.89	0.994	0.001
PIMA Diabetes	0.88	0.778	0.2

Appendix G: Proof of Lemma 2

The expected error costs function is given by

$$g(x, p_f) = a_1 + a_2(\rho\psi P_f - R(x, P_f)).$$

The first derivative with respect to x yields

$$\frac{dg(x, P_f)}{dx} = a_2 \left(\rho\psi \frac{dP_f}{dx} - \frac{\partial R}{\partial P_f} \frac{dP_f}{dx} - \frac{\partial R}{\partial x} \right) \quad (\text{A-17})$$

$$\frac{\partial g(x, P_f)}{\partial P_f} = a_2 \left(\rho\psi - \frac{\partial R}{\partial P_f} \right) = 0 \quad (\text{A-18})$$

Deriving P_f^* from (A-18), we get

$$P_f^* = \left[\frac{ax^a}{c_4 + x^a} \right]^{1/1-a}$$

where $c_4 = \rho\psi - 1$.

Using the optimality condition, the second order condition simplifies to

$$\frac{d^2g(x, P_f)}{dx^2} = -a_2 \left[\left(\frac{\partial^2 R}{\partial P_f^2} \frac{dP_f}{dx} + \frac{\partial^2 R}{\partial x \partial P_f} \right) \frac{dP_f}{dx} + \frac{\partial^2 R}{\partial x \partial P_f} \frac{dP_f}{dx} + \frac{\partial^2 R}{\partial x^2} \right] \quad (\text{A-19})$$

Let

$$S(x, P_f) = \frac{\partial R}{\partial P_f} \quad (\text{A-20})$$

Hence,

$$\frac{dP_f}{dx} = -\frac{\partial S / \partial x}{\partial S / \partial P_f} = -\frac{\partial^2 R / \partial P_f \partial x}{\partial^2 R / \partial P_f^2} \quad (\text{A-21})$$

Substituting the above in the second order condition becomes

$$\frac{d^2g(x, P_f)}{dx^2} = a_2 \left[\frac{(\partial^2 R / \partial x \partial P_f)^2}{\partial^2 R / \partial P_f^2} - \frac{\partial^2 R}{\partial x^2} \right] \quad (\text{A-22})$$

$$= a_2 \left[\frac{P_f^{2(a-1)}(2a-1) + P_f^{a-1}(a^2 - 4a + 1) + 1}{(a-1)P_f^{a-2}} \right] \quad (\text{A-23})$$

Substituting $P_f = P_f^*$, we can conclude that the RHS of the above equation is positive.

■

Appendix H: Solving the Specific Control Problem

Here, we solve the control problem with the specific classification cost function defined in (7). Specifically, we show how the optimal decision variables (\hat{u} , τ , θ) are found.

H.1. Singular Region Exists

When a singular solution exists, the initial effort can be either full effort or zero effort depending on whether the initial discrimination ability is lower or higher than the discrimination ability in the singular region. We follow the steps discussed in Section 2 to obtain singular region solutions. The procedure to solve for τ (time at which the singular region begins) is discussed below

Case 1 ($x_0 < \hat{x}$) : The state equation in this region is given by

$$\dot{x} = c_0(1 - x) - c_1x. \quad (\text{A-24})$$

Solving (A-24), we get

$$x(t) = \frac{c_0}{c_0 + c_1} - \left(\frac{c_0}{c_0 + c_1} - x_0 \right) e^{-(c_0 + c_1)t}. \quad (\text{A-25})$$

Because $x(\tau) = \hat{x}$, using (A-25) we obtain

$$\tau = \frac{1}{c_0 + c_1} \ln \left[\frac{\frac{c_0}{c_0 + c_1} - x_0}{\frac{c_0}{c_0 + c_1} - \hat{x}} \right] \quad (\text{A-26})$$

Case 2 ($x_0 > \hat{x}$) : The state equation in this region is given by

$$\dot{x} = -c_1x, \quad (\text{A-27})$$

which can be solved as

$$x(t) = x_0 e^{-c_1 t}. \quad (\text{A-28})$$

Therefore,

$$\tau = \frac{1}{c_1} \ln \left[\frac{x_0}{\hat{x}} \right] \quad (\text{A-29})$$

To find λ and θ , we know that $u = 0$ beyond $t = \theta$. Thus the state variable in the region $[\theta, T]$ is given by $x = \hat{x} e^{-c_1(t-\theta)}$, $t \geq \theta$.

We need to solve for the adjoint variable in this region. Since $u = 0$, the Hamiltonian is

$$H = f(x) - \lambda c_1 x$$

and the adjoint equation is

$$\dot{\lambda} = -\frac{\partial H}{\partial x} = -f'(x) + \lambda c_1.$$

The solution of the above differential equation is of the form

$$\lambda e^{-c_1 t} = \int -f'(x) e^{-c_1 t} dt + C$$

Letting $y = \hat{x} e^{-c_1(t-\theta)}$ and $\lambda(T) = 0$, the adjoint variable solves as

$$\lambda(t) = \frac{e^{c_1(t-\theta)}}{c_1 \hat{x}} [f(\hat{x} e^{-c_1(t-\theta)}) - f(\hat{x} e^{-c_1(T-\theta)})] \quad (\text{A-30})$$

Using the above, the value of θ can be found by setting $\lambda(\theta) = \hat{\lambda}$. Since the effort exerted in the third segment of the solution is the same for the two cases ($x_0 > \hat{x}$, and $x_0 < \hat{x}$), the solution for λ and θ holds for both cases.

H.2. No Singular Region

When a singular solution does not exist, the optimal solution consists of full effort until $t = \tau$ and zero effort beyond τ . In the first segment, the state equation is given by (A-25), and in the second segment, the state equation is given by (A-28). Following the logic used to derive (A-30), we can write

$$\lambda(\tau) = \frac{1}{c_1 \bar{x}} [f(\bar{x}) - f(\bar{x} e^{-c_1(T-\tau)})]$$

where $\bar{x} = \frac{c_0}{c_0+c_1} - \left(\frac{c_0}{c_0+c_1} - x_0\right) e^{-(c_0+c_1)\tau}$. Given that

$$H_u = k + \lambda(\tau)(1 - x(\tau))c_0,$$

we can find the value of τ where H_u switches from a negative value to 0.

For the infinite horizon case with a continuous maintenance policy, the system requires τ time units to reach steady-state. Thus, there is no second switching point θ . For the discrete policy, starting with any initial discrimination ability x_0 , the discrimination ability level will be either increased to S ($x_0 \leq s$) or allowed to degrade to s ($x_0 > s$). After that the system will reach steady-state as described by Proposition 2.

Appendix I: A Logistic Model of Concept Drift


It is possible that system discrimination ability degrades slowly with high discrimination ability and low discrimination ability levels, but degrades rapidly with intermediate discrimination ability levels. In this section, we examine the optimal model with a state equation that possess such properties.

Consider a state equation in which discrimination ability drifts in the following form.

$$\dot{x} = c_0 u(t)(1 - x(t)) - c_1(1 - x(t))x(t). \quad (\text{A-31})$$

For different values of the initial discrimination ability x_0 , Figure A-2 plots the drift in system discrimination ability for the state equation in (A-31). It can be seen that when the initial discrimination ability is large ($x_0 \cong 0.95$), system discrimination ability degrades slowly at high and low levels of discrimination ability. When the initial discrimination ability is low ($x_0 = 0.3$), the new state equation exhibits a drift pattern similar to the one exhibited by (1). A near-linear drift can also occur when discrimination ability is moderate ($x_0 \cong 0.5$). Thus it can be seen that (A-31) is more versatile than (1) and can represent a variety of different situations.

Figure A-2 Three Curves with Different Drift Patterns.



The revised model yields the same Hamiltonian ($H_u = k + \lambda c_0(1 - x)$) and hence the same solution structure (3) as before: typically starting with zero or full effort, followed by singular effort, and ending with zero effort. The singular solution values, however, change given the modified state equation. Following the solution approach in Section 2, the singular region solutions, denoted by \tilde{x} , \tilde{u} and $\tilde{\lambda}$ are

$$f'(\tilde{x}) = -\frac{k c_1}{c_0}, \quad \tilde{u} = \frac{c_1 \tilde{x}}{c_0}, \quad \text{and} \quad \tilde{\lambda} = -\frac{k}{c_0(1 - \tilde{x})}. \quad (\text{A-32})$$

With the modified deterioration model in (24), it can be seen that the singular level of effort is always at least as high as the level with the deterioration model in (1); i.e., $\tilde{x} \geq \hat{x}$. However \tilde{u} can be larger or smaller than \hat{u} . For the specific ROC curve defined in (7), optimal values for the decision variables $(\tilde{u}, \tau, \theta)$ can be derived following the approach described in Appendix H.

We can also apply the logistic form of the state equation to the discrete model. The stationary results obtained using the original state equation carry over to the revised discrete model. It can be shown that the optimal policy has the same structure as the policy derived for the original state equation. When system discrimination ability is greater than s' , no effort is applied. When system discrimination ability deteriorates to s' , enough effort is applied to raise discrimination ability to S' . The values for S' and s' can be obtained by solving the following two equations:

$$g(S')S'^{\frac{r}{c_1}} = g(s')\frac{s'^{\frac{r}{c_1}}}{1-s'} + \frac{r}{c_1} \int_{s'}^S \frac{g(\xi)\xi^{\frac{r}{c_1}-1}}{(1-\xi)^{\frac{r}{c_1}+1}} d\xi$$

$$-Kr = g(S') - g(s')$$

where $g(x) = f(x) - r\frac{k}{c_0} \log(1-x) + \frac{kc_1x}{c_0}$.

■