

Combining learned and highly-reactive management

Alva L. Couch and Marc Chiarini
Tufts University

couch@cs.tufts.edu,
mchiar01@cs.tufts.edu

Context of this paper

- This paper is Part 3 of a series
- Part 1 (AIMS 2009): Can ignore external influences and still manage systems in which cost and value are simply increasing.
- Part 2 (ATC 2009): Can ignore external influences and still manage SLA-based systems.
- Part 3: (this paper) Can integrate these strategies with more conventional management strategies and reap “the best of both worlds”.

The inductive step

- In fact, one might think of the first two steps as the basis case of an induction proof.
- Now we proceed to the inductive step, in which we
 - “assume true for n ”
 - “show true for $n+1$ ”.
- Where n is the number of management paradigms we wish to apply!

The basis step

- Just because we can manage without detailed models, doesn't mean we should.
- If we have precise models, we also have accurate measures of efficiency.
- But the capability to manage without details is a fallback position that allows less robust models to recover from **catastrophic changes**.

The big picture

- In a truly open world, the structure of the applicable model of behavior may change over time.
- A truly open strategy should cope with such changes.
- Key is to consider each potential model of behavior as a **hypothesis to be tested** rather than a **fact to be trusted**.

Good news and bad news

- The upside of machine learning is that it creates usable models of previously unexplained behaviors.
- The downside is that these models react poorly to catastrophic changes and mis-predict behavior until retrained to the new behavior of the system.
- Can we have the best of both worlds?

Best of both worlds?

- Highly-reactive model: tuned to short-term behavior.
- Historical model: tuned to long-term history.
- If the system changes unexpectedly, then the historical model is invalidated, but the highly-reactive model continues to manage the system **until the long-term model can recover.**

A simple demonstration

- Basis model: highly reactive, utilizes 10 steps of history.
- Historical model: based upon 200 steps worth of history.

Our simulation parameters

- R = resource utilization.
- L = known (measurable) load.
- X = unknown load.
- P = performance = $a R/(L+X) + b$
- $V(P)$ is the value of P (a step function).
- $C(R)$ is the cost of R (a step function).
- Attempt to learn $P \sim c R/L + d$ and maximize $V(P(R,L)) - C(R)$.

What is acceptable accuracy?

- Some statistical notion of whether a model should be believed.
- Best characterized as a **hypothesis test**.
- **Null hypothesis:** the model is correct.
- **Accept the null hypothesis** unless there is **evidence to the contrary**.
- **Else reject the null hypothesis** and **don't use the model**.

A demon called independence

- Many statistical tests require **independence of samples**.
- We **almost never** have that.
- Our training tuples (P_i, R_i, L_i) are measured **close together in time**, and in realistic systems, nearby measurements in time are usually **dependent**.
- So many statistical tests of model correctness **fail to apply**.

Coefficient of determination

- Coefficient of determination (r^2) is a measure of how accurate a model is.
- $r^2=1$ → model **precisely reflects measurements.**
- $r^2=0$ → **model is useless** in describing measurements.

Why r^2 ?

- Doesn't require independence.
- Can test models determined by other means.
- Unitless.
- A good comparison statistic for **relative correctness** of models.

Coefficient of determination

- For samples $\{(X_i, Y_i)\}$ where $Y_i \sim f(X_i)$,
$$r^2 = 1 - \frac{\sum (Y_i - f(X_i))^2}{\sum (Y_i - \bar{Y})^2}$$
where \bar{Y} is the mean of $\{Y_i\}$
- In our case,
$$r^2 = 1 - \frac{\sum (P(R_i, L_i) - P_i)^2}{\sum (P_i - P)^2}$$
where
 - P_i is **measured** performance, $P = \text{mean}(P_i)$
 - $P(R_i, L_i)$ is **model-predicted** performance

Using r^2


- If $r^2 \geq 0.9$, **accept the hypothesis that the learned model is correct** and obey its predictions to the letter.
- If $r^2 < 0.9$, **reject the hypothesis that the learned model is correct** and manage via the reactive model.

A novel visualization


- Learned data with $r^2 \geq 0.9$ is green.
- Learned data with $r^2 < 0.0$ is yellow-green.
- Reactive data that is used is red.
- Reactive data that is unused is orange.
- Target areas of maximum V-C are gray.

Learned model
 $r^2 \geq 0.9$ is green
 $r^2 < 0.9$ is yellow
-green


Goodness of fit




Regions of maximum V-C



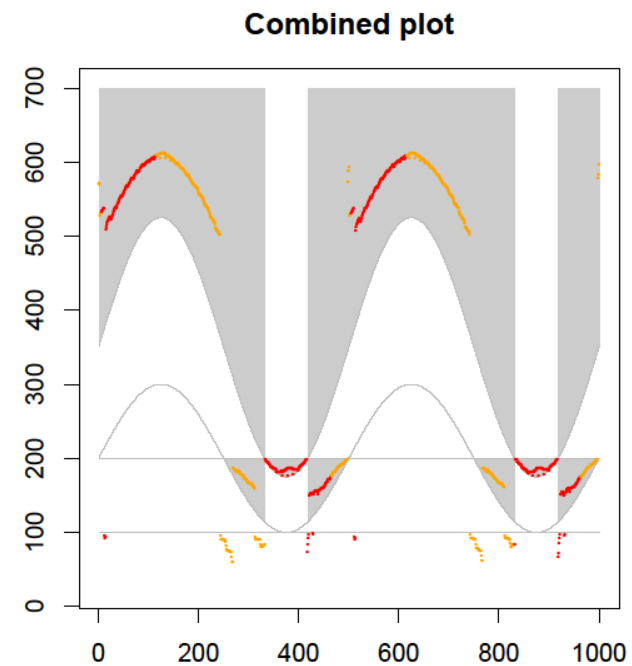
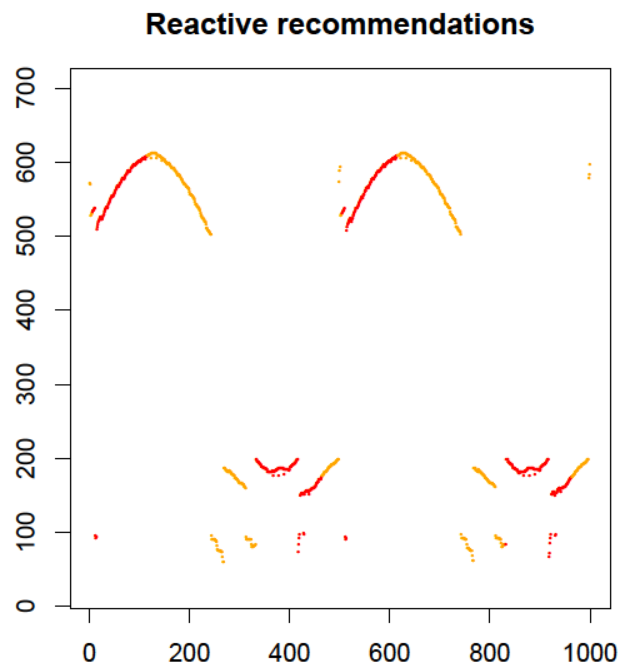
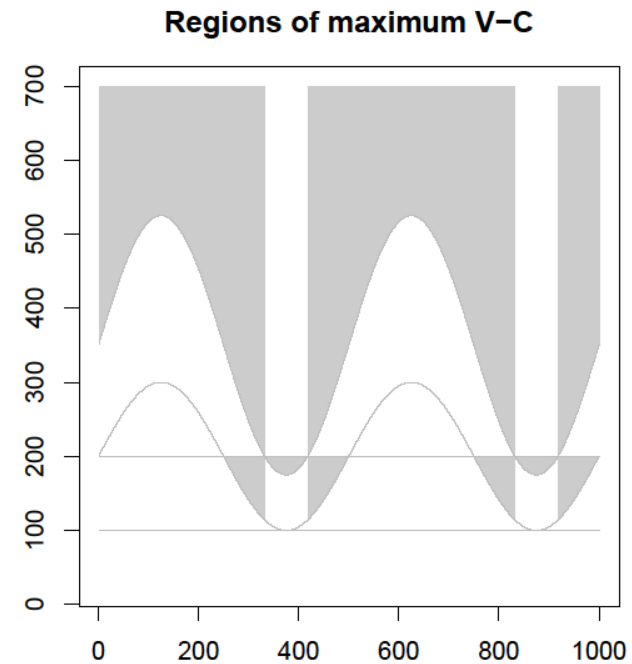
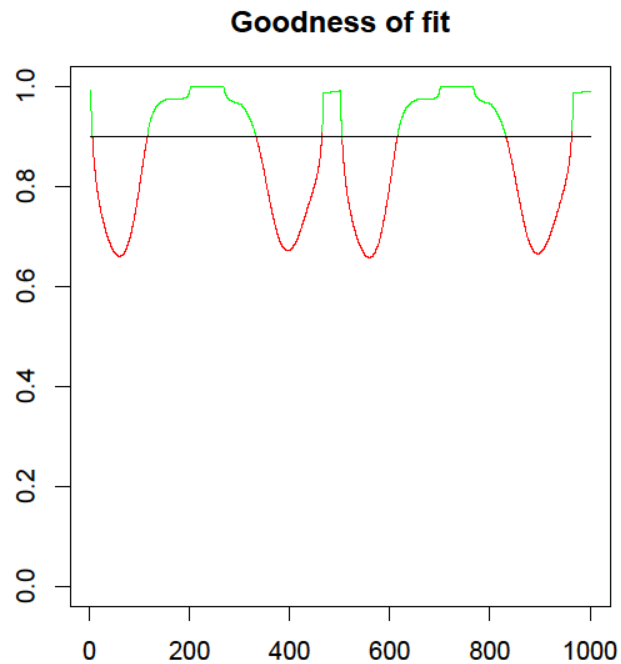
Learned recommendations



Combined plot



Reactive model
Active when red
Inactive when orange



In the diagrams



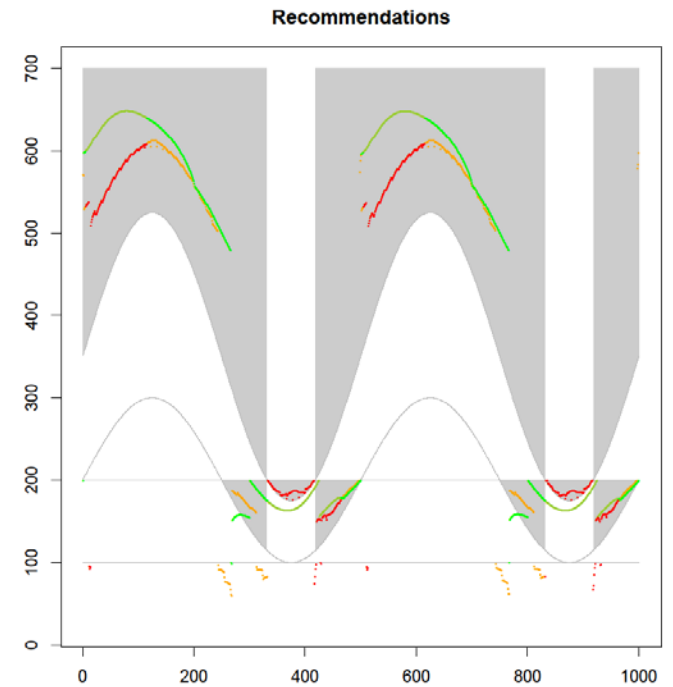

- X axis is time, Y axis is resources
- Gray areas represent theoretical optima for V-C.
- Gray curves depict changes in V.
- Gray horizontal lines depict changes in C.

Composite performance of the two models compared.

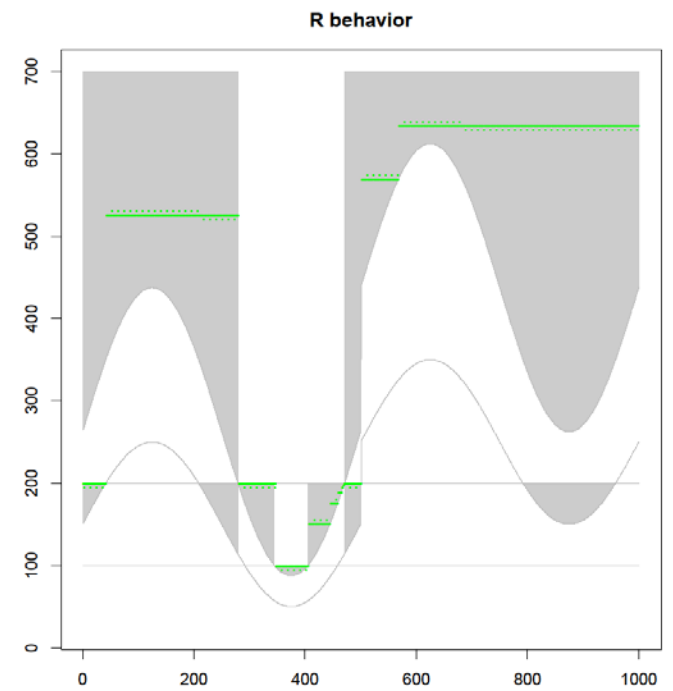



Cutoffs are models' ideas of where boundaries lie.

Recommendations are what the model suggests to do.

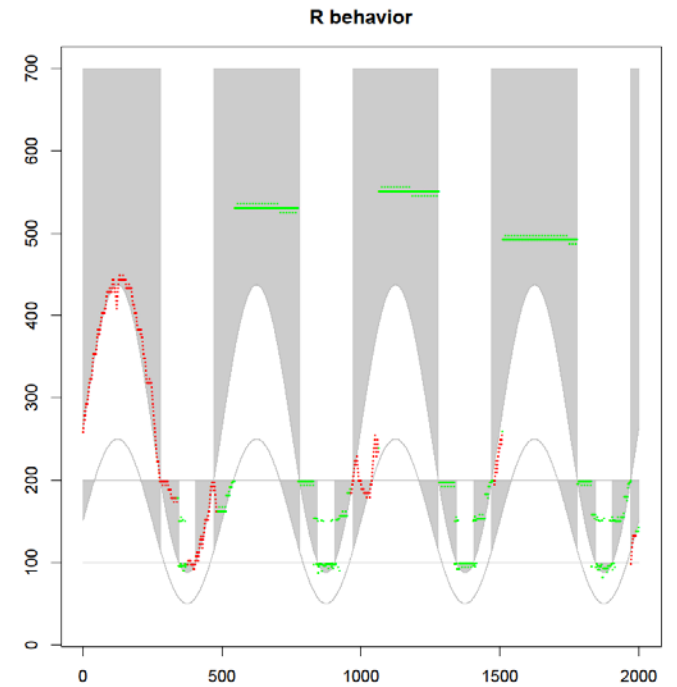

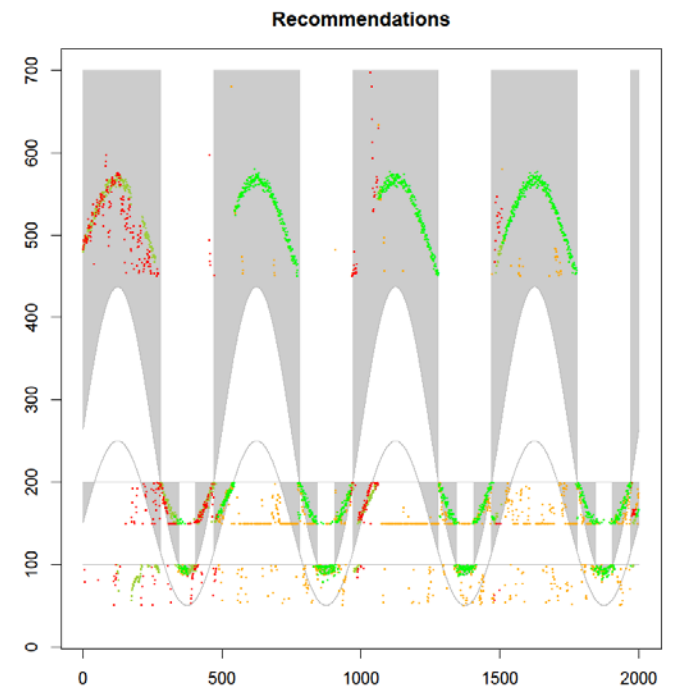

Behavior is what happens.





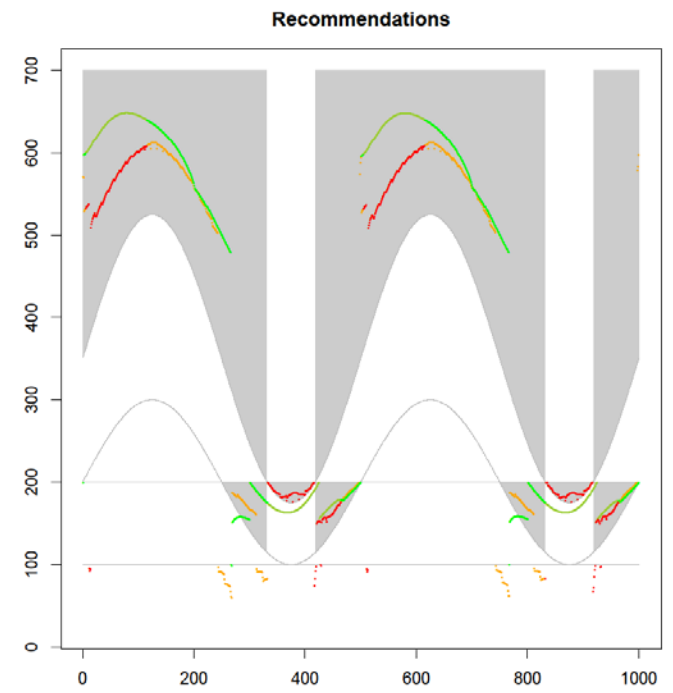

Learned model handles load discontinuities easily







Noise in measuring L leads to rejecting model validity



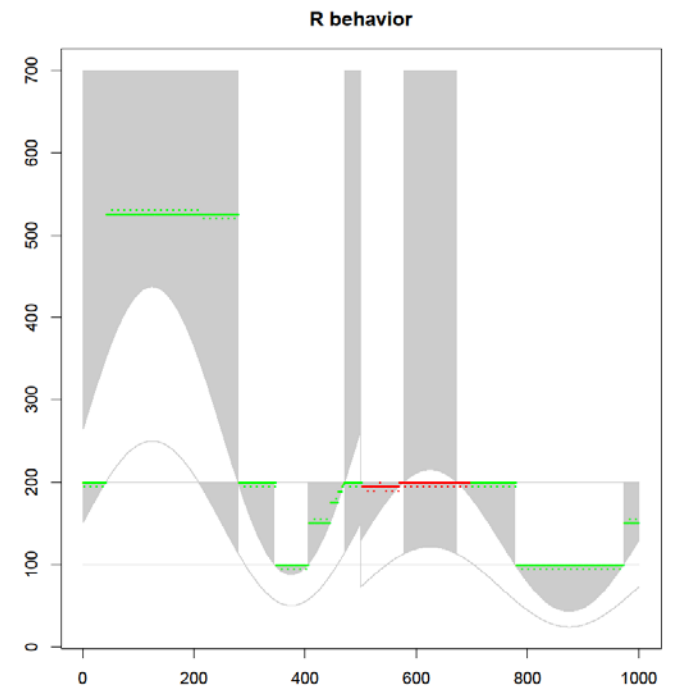



Even a constant unknown factor X periodically invalidates the learned model.



Periodic variation in the unknown X causes lack of belief in the learned model.



Catastrophe in which learned model fails is mitigated by reactive model.



The r^2 challenge

- At this point you may think I'm **crazy**, and it is only fair to return the favor. I ask:
- **Do *your* models pass the r^2 test?**
- **Or do you simply “believe in them”?**
- My conjecture: no commonly used model does!
- Passing an r^2 test is very tricky in practice:
 - Time skews must be eliminated.
 - Time dependences must be considered.

Conclusions

- We have shown that learned and reactive strategies can be combined to handle **even catastrophic changes** in the managed system.
- Key to this is to **validate the model being used** for the system.
- If all goes well, that model is **valid**.
- If the worst happens, that model is **rejected** and a fallback plan activates.
- Result is that the system can handle **open-world changes**.

Questions?

Combining learned and highly-reactive
management

Alva L. Couch and Marc Chiarini

Tufts University

couch@cs.tufts.edu, mchiar01@cs.tufts.edu