Adaptive Control — What Can We Learn?

Anders Rantzer

on the occasion of Malcolm Smith's 60th birthday

with best regards from Karl Johan Åström

Adaptive Control — What Can We Learn?

IEEE TRANSACTIONS ON AUTOMATIC CONTROL, VOL. AC-31, NO. 4, APRIL 1986

Stable Adaptive Regulation of Arbitrary nth-Order Plants

GERHARD KREISSELMEIER AND MALCOLM C. SMITH

Abstract

This paper presents an algorithm for adaptively stabilizing and asymptotically regulating an arbitrary single-input single-output linear time-invariant plant, which is controllable and observable, of known order n, and has unknown parameters. No further assumptions are made. No external probing signal is required.

Adaptive Control — What Can We Learn?

IEEE TRANSACTIONS ON AUTOMATIC CONTROL, VOL. AC-31, NO. 4, APRIL 1986

Stable Adaptive Regulation of Arbitrary nth-Order Plants

GERHARD KREISSELMEIER AND MALCOLM C. SMITH

Abstract

This paper presents an algorithm for adaptively stabilizing and asymptotically regulating an arbitrary single-input single-output linear time-invariant plant, which is controllable and observable, of known order n, and has unknown parameters. No further assumptions are made. No external probing signal is required.

Adaptive Control — What Can We Learn?

IEEE TRANSACTIONS ON AUTOMATIC CONTROL, VOL. AC-31, NO. 4, APRIL 1986

Stable Adaptive Regulation of Arbitrary nth-Order Plants

GERHARD KREISSELMEIER AND MALCOLM C. SMITH

Abstract

This paper presents an algorithm for adaptively stabilizing and asymptotically regulating an arbitrary single-input single-output linear time-invariant plant, which is controllable and observable, of known order n, and has unknown parameters. No further assumptions are made. No external probing signal is required.

Adaptive Control — What Can We Learn?

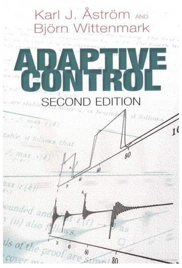

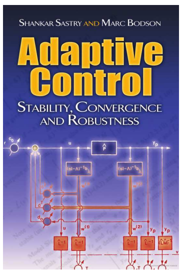

Adaptive Control — What Can We Learn?

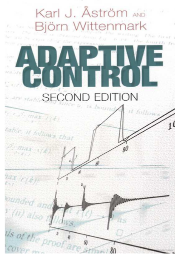

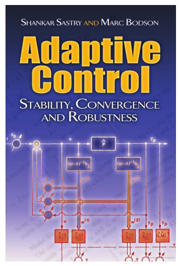

Åstrom & Wittenmark 1995:
"*Unfortunately, there is no collection of results that can be called a theory of adaptive control in the sense specified.*"

Global Stability and Performance?

Conclusions from rich literature in the 1970-80s:

- Adaptive controllers can be made to converge under ideal conditions and without forgetting factor in the estimator.
- Forgetting factors are desirable in practice!
- Converging controller gains give lack of excitation.
- External probing saves stability, but worsen performance.
- Dual control needed: Exploration/exploitation trade-off.

Outline

- **A "simple" adaptive problem**
- Using recent progress on concentration inequalites
- Multivariable extensions

A "Simple" Adaptive Control Problem

A scalar system, one input u_t, one disturbance w_t and one *unknown* parameter $a \in \mathbb{R}$:

$$x_{t+1} = ax_t + u_t + w_t$$

Problem:
What is the smallest ℓ_2-gain from w to x achievable by a causal (possibly adaptive) feedback law from x to u?

A "Simple" Adaptive Control Problem

A scalar system, one input u_t, one disturbance w_t and one *unknown* parameter $a \in \mathbb{R}$:

$$x_{t+1} = ax_t + u_t + w_t$$

Problem:
What is the smallest ℓ_2-gain from w to x achievable by a causal (possibly adaptive) feedback law from x to u?

Megretski 2004:
There exists *no* feedback law that gives finite ℓ_2-gain for all a!

A "Simple" Adaptive Control Problem

A scalar system, one input u_t, one disturbance w_t and one *unknown* parameter $a \in \mathbb{R}$:

$$x_{t+1} = ax_t + u_t + w_t$$

Problem:
What is the smallest ℓ_2-gain from w to x achievable by a causal (possibly adaptive) feedback law from x to u?

Megretski 2004:
There exists no feedback law that gives finite ℓ_2-gain for all a!

Conclusion:
Worst case gain not the right framework. Use stochastics!

A "Simple" Adaptive Control Problem

A scalar system, one input u_t, one disturbance w_t and one *unknown* parameter $a \in \mathbb{R}$:

$$x_{t+1} = ax_t + u_t + w_t$$

Problem:
What is the smallest ℓ_2-gain from w to x achievable by a causal (possibly adaptive) feedback law from x to u?

Conclusion:
Worst case gain not the right framework. Use stochastics!

Good News:
Powerful theory on stochastic tail and concentration bounds! (already exploited in statistical machine learning)

A "Simple" Adaptive Control Problem

A scalar system, one input u_t, one disturbance w_t and one unknown parameter $a \in \mathbb{R}$:

$$x_{t+1} = ax_t + u_t + w_t$$

Self-tuning controller:

$$\widehat{a}_t = \frac{\sum_{k=1}^{t-1}(x_{k+1} - u_k)x_k}{\sum_{k=1}^{t-1} x_k^2} \qquad u_t = -\widehat{a}_t x_t$$

Prove bounds on estimation error and regret!
Analyse interplay between exploration and exploitation.

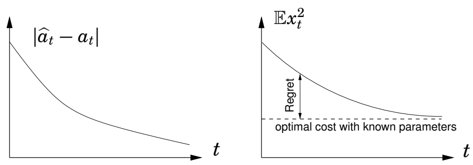

Outline

- A "simple" adaptive problem
- **Using recent progress on concentration inequalites**
- Multivariable extensions

Tail and Concentration Inequalites

- Mathematics (measure theory, combinatorics, analysis)
- Compressed sensing
- Statistical model selection
- Machine Learning
- Network Routing
- Pattern recognition
- ⋮

The Chernoff Bound

Let $\theta > 0$. Then the probability that the random variable X exceeds a is bounded above by the expected value of $e^{\theta(X-a)}$.

$$\mathbb{P}[X > a] = \mathbb{E}\chi_a(X) \leq \mathbb{E}e^{\theta(X-a)}$$

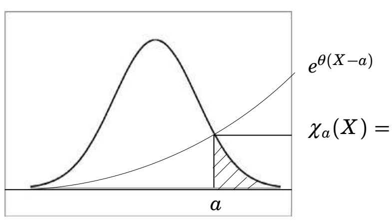

$$\chi_a(X) = \begin{cases} 1 & \text{if } X > a \\ 0 & \text{else} \end{cases}$$

Example: X Gaussian with unit variance.
$\mathbb{E}e^{\theta(X-10)} = e^{\theta^2/2-10\theta}$, so $\mathbb{P}[X > 10] \leq \min_\theta e^{\theta^2/2-10\theta} = e^{-50}$.

“Random Design” Linear Regression

$$x_{t+1} = ax_t + u_t + w_t$$

Least-squares esimate:

$$\widehat{a}_t = \frac{\sum_{k=1}^{t}(x_{k+1} - u_k)x_k}{\sum_{k=1}^{t} x_k^2} \qquad \widehat{a}_t - a_t = \frac{\sum_{k=1}^{t} w_k x_k}{\sum_{k=1}^{t} x_k^2}$$

Chernoff gives

$$\mathbb{P}\left[\widehat{a}_t - a_t \geq \rho\right] = \mathbb{P}\left[\sum_{k=1}^{t} w_k x_k \geq \rho \sum_{k=1}^{t} x_k^2\right] \leq \frac{1}{(1+\rho^2)^{t/2}}$$

independently of control law $u_k = \mu(x_k)$!

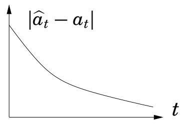

Decay of Regret

$x_{t+1} = ax_t + u_t + w_t$ with self-tuning controller $u_t = -\widehat{a}_{t-1}x_t$.
Define $X_t = \sum_{k=1}^{t} x_k^2$ and $Y_t = \sum_{k=1}^{t} w_k x_k$. Then

$$x_{t+1} = \frac{Y_{t-1}}{X_{t-1}} x_t + w_t$$

$$\frac{|x_{t+1}|}{\sqrt{X_t}} \leq \left|\frac{Y_{t-1}}{X_{t-1}}\right| + \frac{|w_t|}{\sqrt{X_{t-1}}}$$

$$|x_{t+2}| \leq \left(\frac{Y_{t-1}^2}{X_{t-1}} + w_t^2\right)\frac{|x_t|}{X_{t-1}} + \left|\frac{Y_{t-1}}{X_{t-1}} w_t\right| + |w_{t+1}|$$

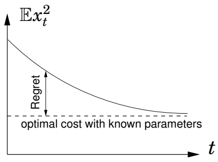

Outline

- A “simple” adaptive problem
- Using recent progress on concentration inequalites
- **Multivariable extensions**

A Matrix Chernoff Bound

Let X be a random matrix. The probability that the maximal eigenvalue of X exceeds a is bounded above as follows:

$$\mathbb{P}[\lambda_{\max(X)}(X) > a] \leq \operatorname{tr} \mathbb{E} e^{\theta(X - aI)}$$

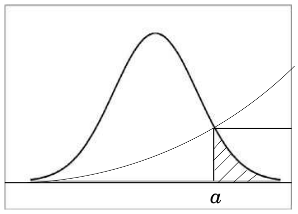

A Less Simple Adaptive Control Problem

For a MIMO system, with unknown $A, B \in \mathbb{R}^{n \times n}$:

$$x_{t+1} = Ax_t + Bu_t + w_t$$

Self-tuning controller:

$$\begin{bmatrix}\widehat{A}_t & \widehat{B}_t\end{bmatrix} = \sum_{k=1}^{t-1} x_{k+1} \begin{bmatrix} x_k \\ u_k \end{bmatrix} \left(\sum_{k=1}^{t-1} \begin{bmatrix} x_k \\ u_k \end{bmatrix} \begin{bmatrix} x_k \\ u_k \end{bmatrix}^T\right)^{-1} \qquad u_t = -\widehat{B}_t^{-1}\widehat{A}_t x_t$$

On-going work: Bounds on estimation error and regret!

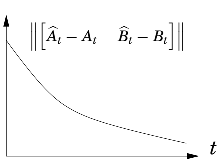

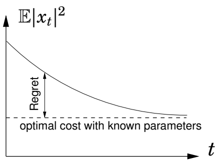

Summary

Conclusion:
Forgetting factor needs to be combined with external excitation. But when and how much?

Theory of the 1980s were lacking efficient tools for dual control.

Good News:
Powerful theory on stochastic tail and concentration bounds!

Supports exploration/exploitation trade-off analysis.

Congratulations Malcolm!