Learning, Dynamics and Control

ESE6180

Ingvar Ziemann, Wednesday August 28 2024

ESE6180: Learning, Dynamics and Control

Instructor: Ingvar Ziemann

Teaching Assistants: Bruce Lee, Thomas Zhang

Class: Monday, Wednesday both at noon-1:30pm in GLAB101

Office Hours: Monday 2:30-4pm, Wednesday 10am-11:30am in M70 (Towne)

Website: https://ingvarziemann.com/teaching (browse your way to ESE6180)

Mail: ingvarz@seas.upenn.edu (start subject line with ESE6180)

Introductions

Course Overview

Overview: This course will provide students an introduction to the emerging area at the intersection of machine learning, dynamics, and control.

We investigate machine learning algorithms that interact with the physical world, with an emphasis on a holistic understanding of the interplay between concepts from machine learning (e.g., generalization, sample complexity), probability and statistics (e.g., concentration, information-theoretic lower bounds) and dynamical systems and control theory (e.g., feedback, stability, observability).

Intended Audience: advanced graduate students who are interested in applying novel research concepts to their own work. By the end of this course, students will be ready to start doing research in the learning, dynamics and control space.

Course Overview
Tentative Topics

Part 1: Foundations

- IID Mean Estimation, Linear Regression and Concentration Inequalities
- Covering numbers in learning in $\mathbb{R}^d$
- The Hanson-Wright Inequality (concentration with quadratic dependence)
- Linear Regression with Dependent Data, Linear System Identification

Part 2: Control

- LQR Recap
- Offline Learning of LQR
- Policy Gradient Methods for LQR

Part 3: Fundamental Limits and Active Learning

- Information-Theoretic Lower Bounds (AKA fundamental limits)
- Active Learning for LQR

Part 4: Further Topics

- Learning in nonlinear time-series/dynamics
- Martingale Methods
- More TBD if time permits (rep learning?)

Lecture notes loosely follow this structure

Course Overview

Official Prerequisites: ESE500 (Linear Systems) and ESE530 (Probability & Random Processes)

Unofficial Prerequisites: and most importantly **mathematical maturity.**

This is an advanced theory intensive course: Our focus will be on proving strong theoretical guarantees (and corresponding fundamental limits to) about the sample efficiency, stability and performance of learning algorithms.

What this course is not: This is not a (deep) reinforcement learning class or an applied ML class. There may be some programming elements, but these will be minimal and mostly used to verify and support **theory**.

Grading

Homework (60%):

- there will five (5) homework assignments.
- An initial homework assignment, Homework 1, will be handed out on the first day of class, and will be worth 12%.
- Homework 1 is mandatory, and must be passed to a satisfactory level: it is used to check your knowledge of prerequisites.
- The remaining four homework assignments will also each be worth 12%.

Course Project (40%):

- Students will be expected to work on a theory-focused project (in groups of up to 2 students)

Homework Policy

Each hand-in must be written up in LaTeX in single column style in the article document class.

We ask that you write out detailed and rigorous solutions.

You get 6 free late days: **Beyond that no late assignments will be graded**.

You are allowed, even encouraged, to work on homework in small groups, but you must write up your own homework solutions and code to hand in – please indicate who you collaborated with on your assignments.

Each homework problem will be graded on a scale of 0-4.

Homeworks are submitted on Canvas.

Course Project

In groups of up to two students

Report Format: Latex single column in the article document class with options letterpaper and 11pt

Project Proposal: Your proposal should be 2 pages maximum (not including references), and should include title, team members, abstract, related works, problem formulation and goals

Midterm Report: Your report should be 4 pages maximum (not including references). Your midterm report should build on your project proposal, and outline your solution approach, current progress and preliminary results, as well as highlight challenges that you are facing.

Final Report: Your report should be 10 pages maximum (not including references and supplementary material).

https://ingvarziemann.com/ese6180/projects

Course Project

Your final report will be evaluated by the following criteria:

Merit: Is your problem formulation and solution strategy well-motivated? Can you justify the complexity-level of your approach?

Technical depth: Is your project technically challenging? Did you write your own code, or did you use available software packages? While it is ok for a project to lean more towards theory or implementation, the sum of theoretical + implementation efforts should remain consant (i.e., if you use existing software packages rather than write your own code, the theoretical component of your project should be more ambitious).

Presentation: Are your solution approach, assumptions, results, and interpretations of experimental/theoretical outcomes clearly explained and/or justified? Is the report clearly and written? Are the mathematical arguments rigorous and easy to follow? Are graphs/visualizations clear?

Course Intro

What led us here?

What is ML?

What is control?

Learning with dependent data… and is this relevant?

How do we study their synthesis?

What led us here?

Ambition

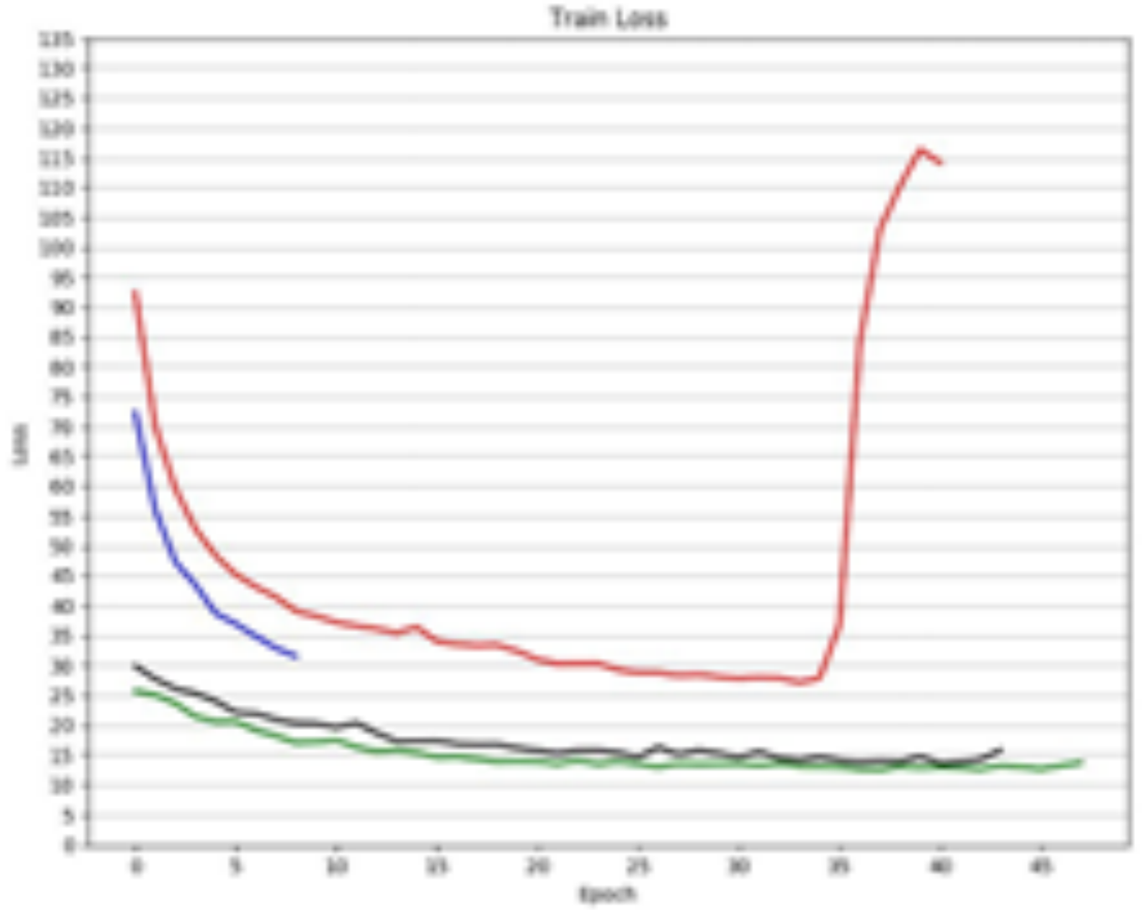

Reality

Not just in sim

What is ML?

Why do we need it?

Using Past data to learn about/and or act upon the world

Too Complex Environments

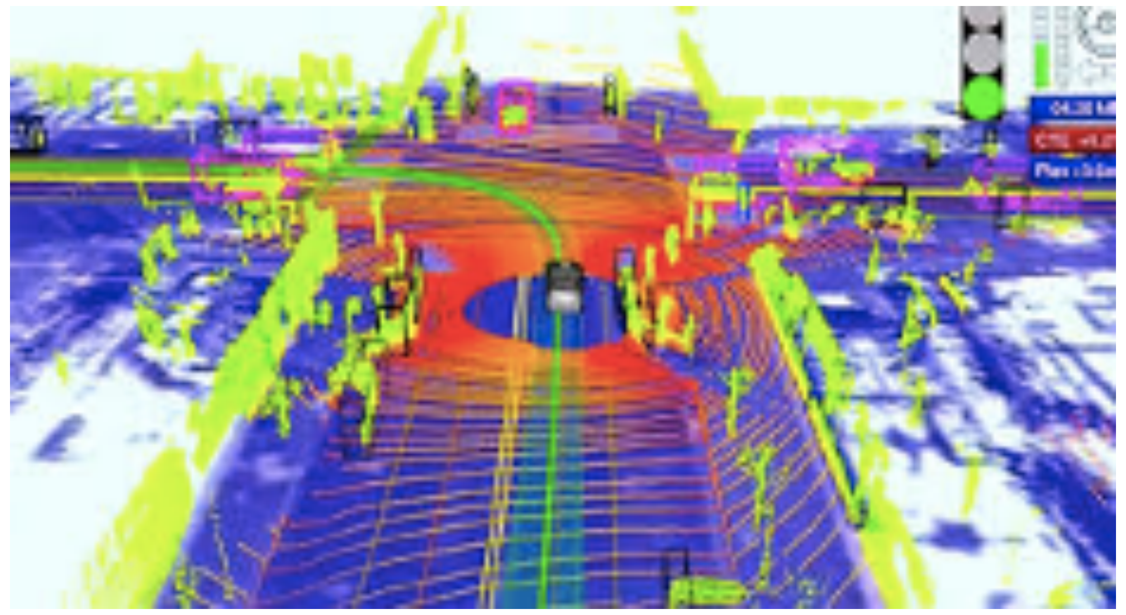

Too Complex Sensing

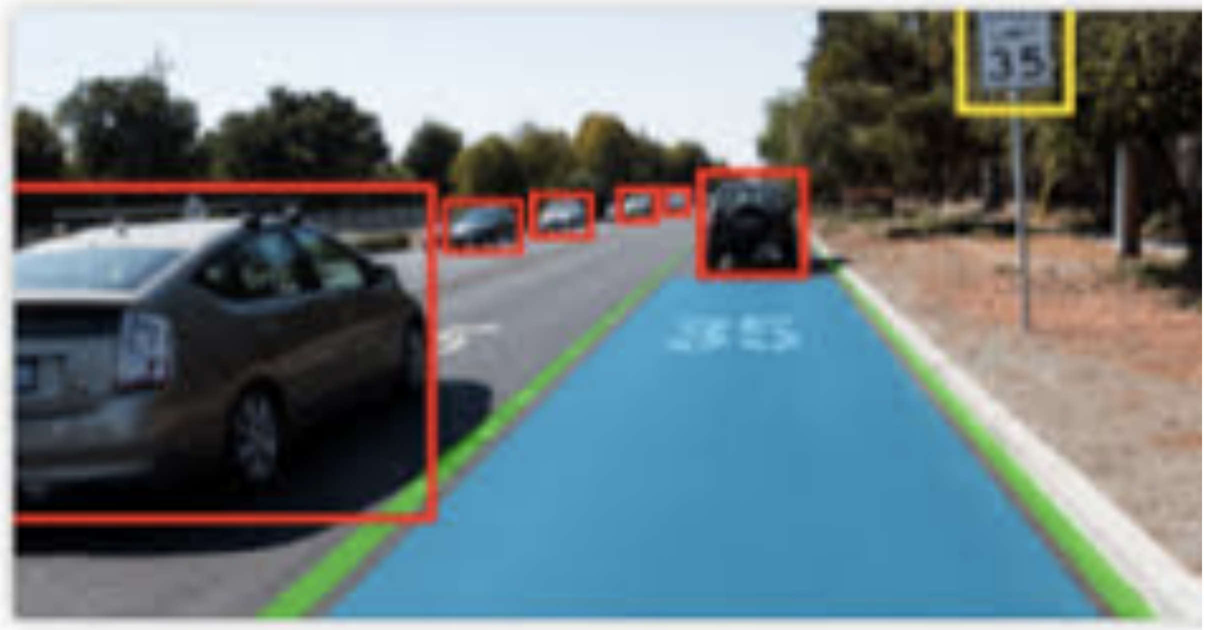

No Known Models

What is control?

Why do we need it?

Using feedback to mitigate dynamic uncertainty

Uncertain Environments

Uncertain Sensing Components

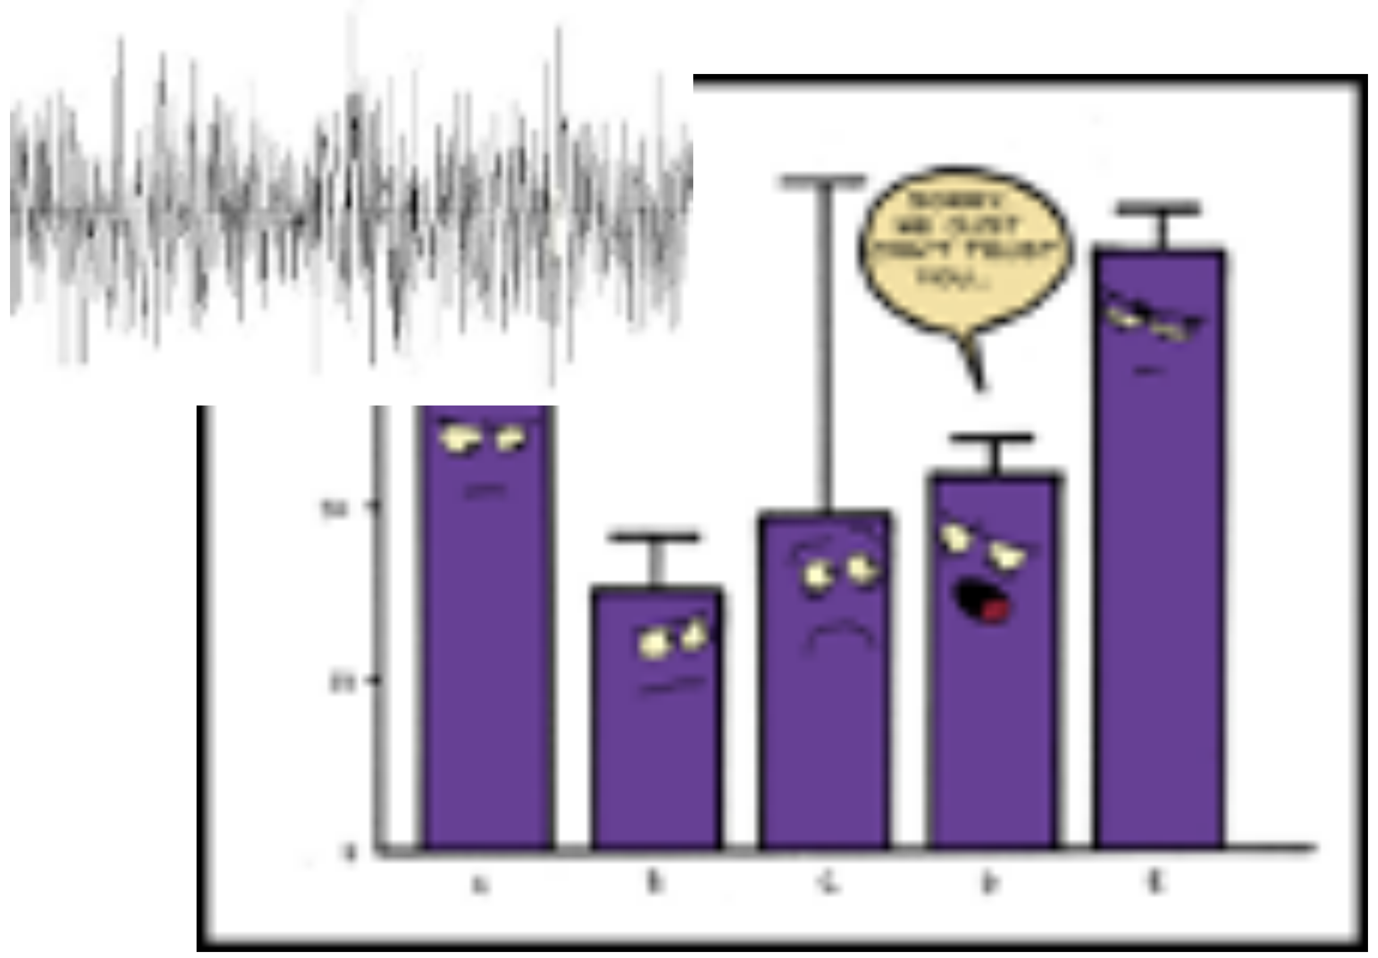

Uncertain Models

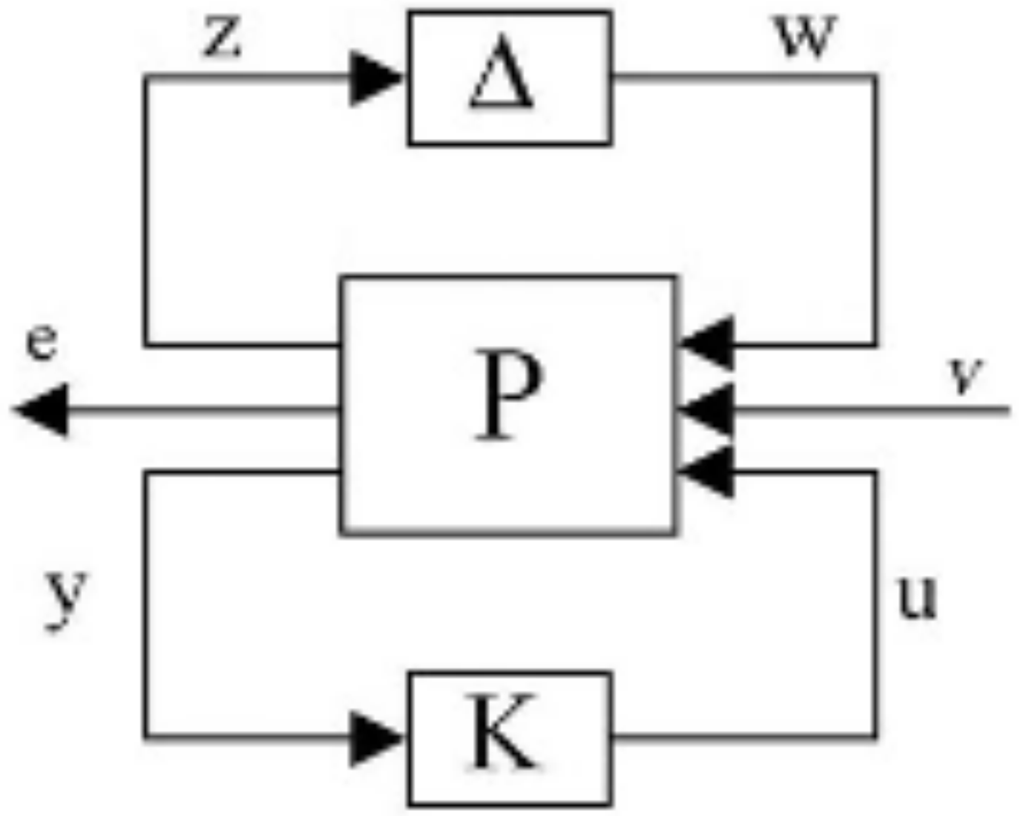

Learning and Control?

Machine Learning

Estimate and Predict

Uses data to reduce uncertainty

More data
⇒ Better Models

Decision-Making under Uncertainty

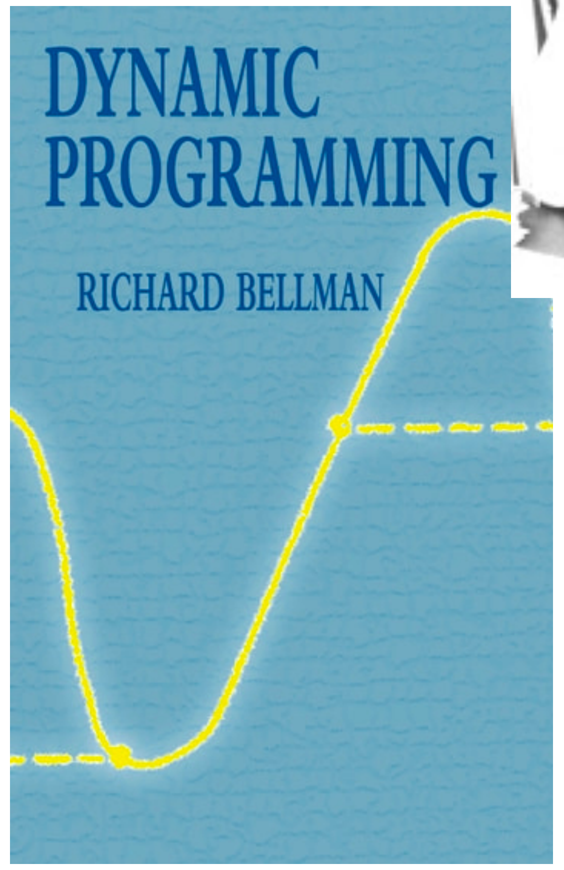

RL & Control

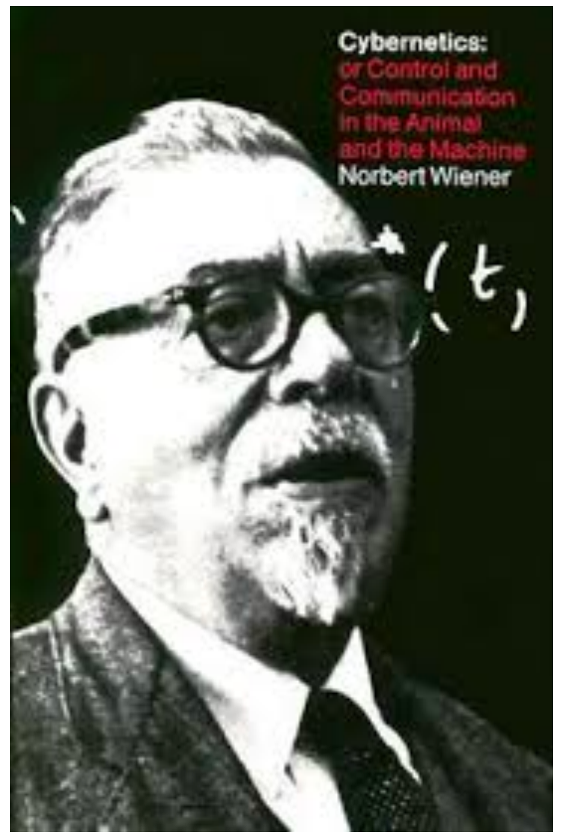

Control

Regulate and Control

Uses feedback to mitigate uncertainty

Better Models/Predictions ⇒
Better Performance

Decision-Making under Uncertainty

In fact lots of common history!

Another advantage: can reason about fundamental limits

Guaranteed Margins for LQG Regulators

JOHN C. DOYLE

Abstract—There are none.

INTRODUCTION

Considerable attention has been given lately to the issue of robustness of linear–quadratic (LQ) regulators. The recent work by Safonov and Athans [1] has extended to the multivariable case the now well-known guarantee of 60° phase and 6 dB gain margin for such controllers. However, for even the single-input, single-output case there has remained the question of whether there exist any guaranteed margins for the full LQG (Kalman filter in the loop) regulator. By counterexample, this note answers that question; there are none.

Bode's sensitivity integral

$$S(s) = \frac{1}{1 + G(s)C(s)}$$

$$\int_0^\infty \ln|S(j\omega)| d\omega = 0$$

Temporally dependent data is everywhere

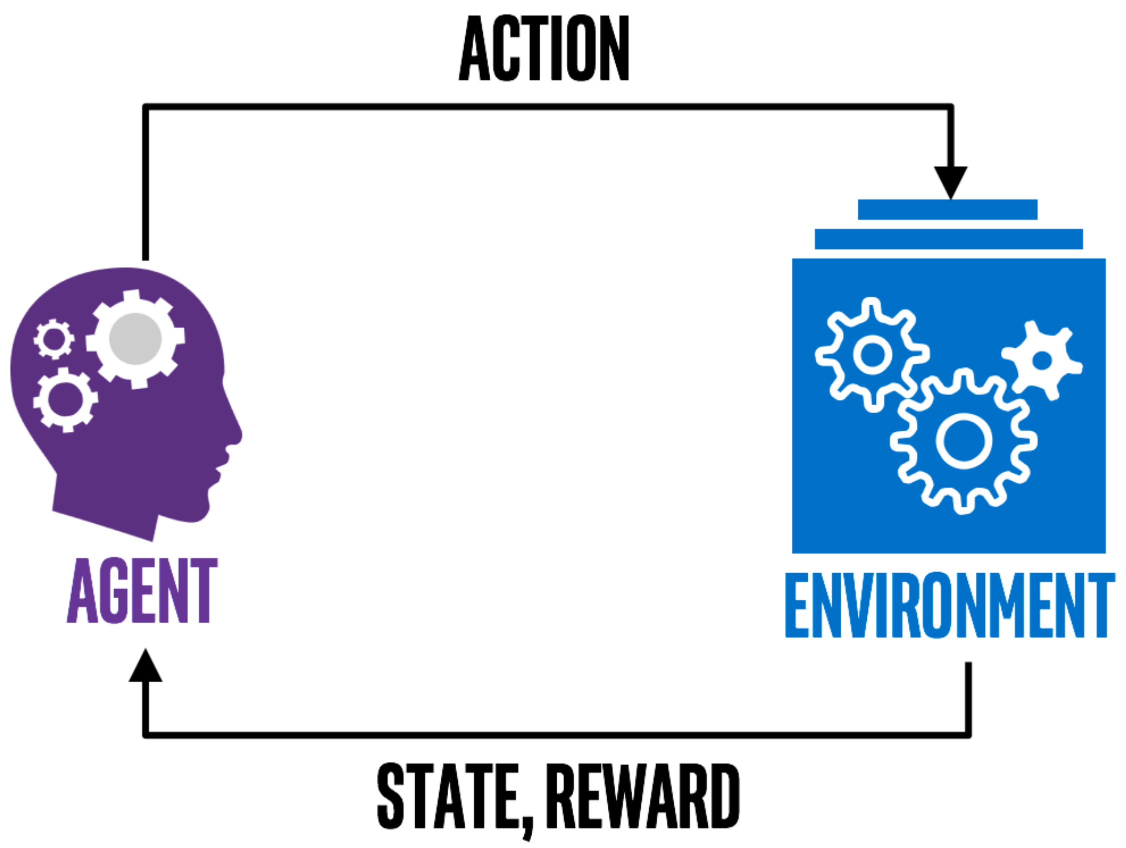

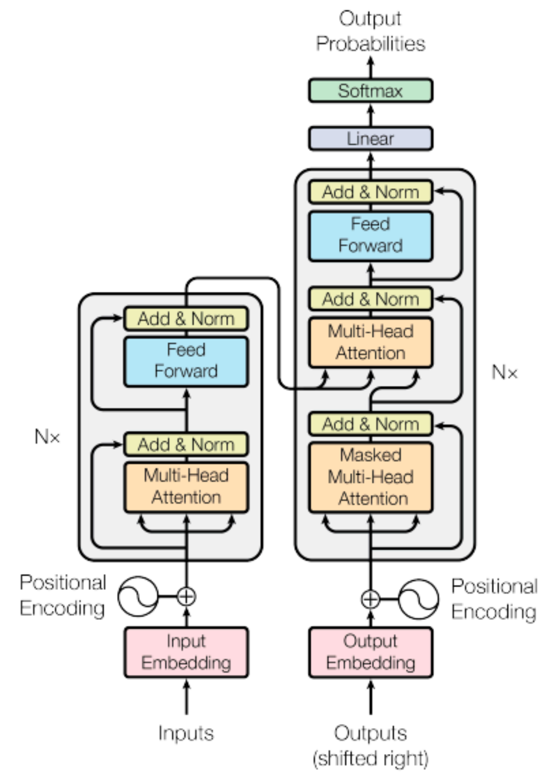

We understand iid learning very well

Uniform convergence, PAC, etc

Instance optimal (non-)asymptotics

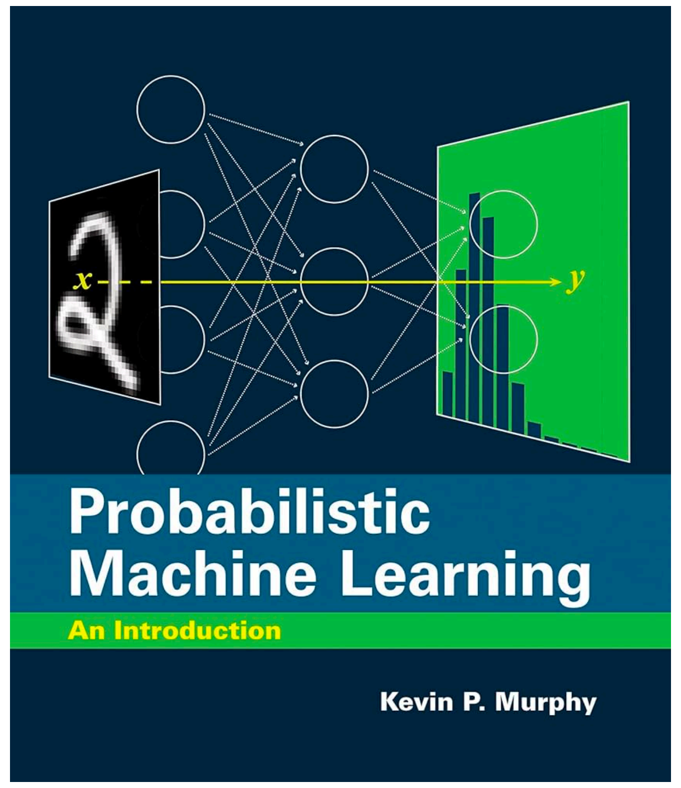

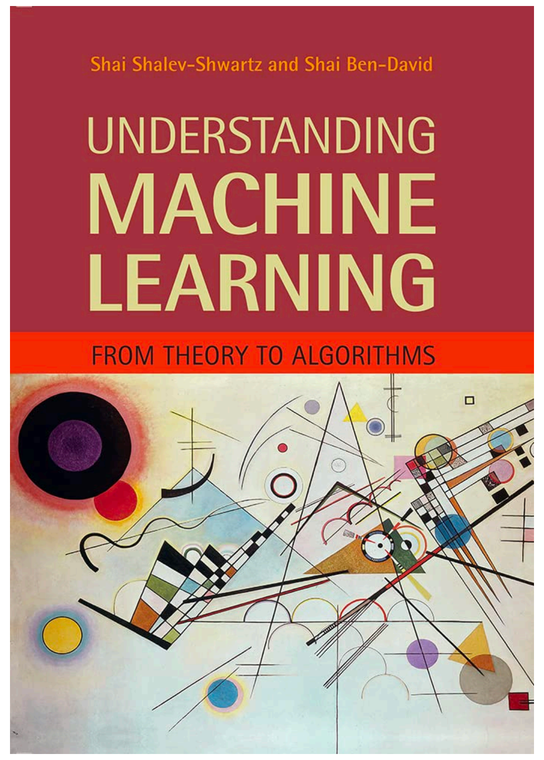

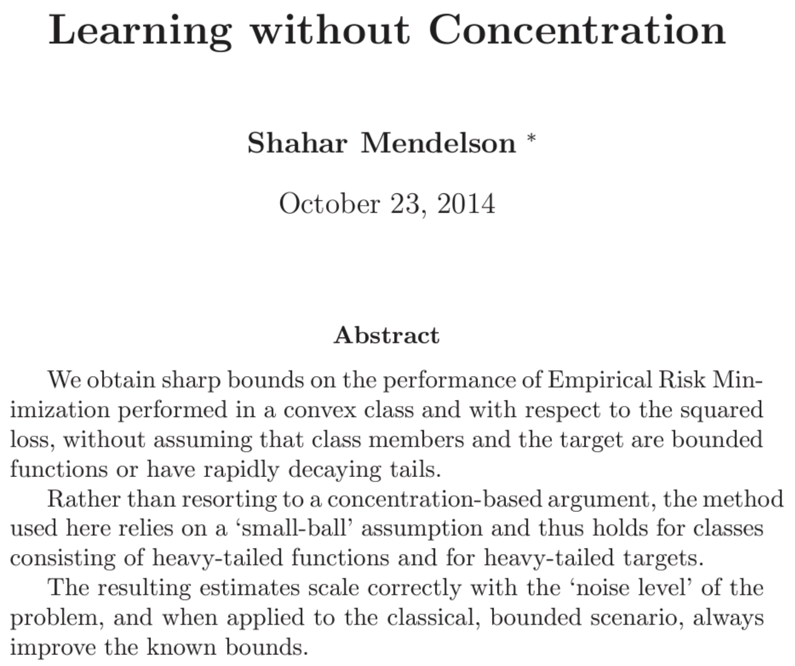

Learning without Concentration

Shahar Mendelson *

October 23, 2014

Abstract

We obtain sharp bounds on the performance of Empirical Risk Minimization performed in a convex class and with respect to the squared loss, without assuming that class members and the target are bounded functions or have rapidly decaying tails.

Rather than resorting to a concentration-based argument, the method used here relies on a 'small-ball' assumption and thus holds for classes consisting of heavy-tailed functions and for heavy-tailed targets.

The resulting estimates scale correctly with the 'noise level' of the problem, and when applied to the classical, bounded scenario, always improve the known bounds.

Dependent data is less well understood

1: Correct notion of dependence?

2: Optimal rates for some reasonable notion of dependence?

Focus on supervised learning with square loss $l_{sq}(f, x, y) = \|y - f(x)\|^2$

Learning from dependent data

Consider a time-series

$$Y_i = f_\star(X_i) + W_i, \qquad i = 1,\ldots,n$$

Y_i - the target or output

X_i - the covariates

W_i - the noise variables

(Y_i, X_i) depend on (Y_j, X_j) for $j < i$

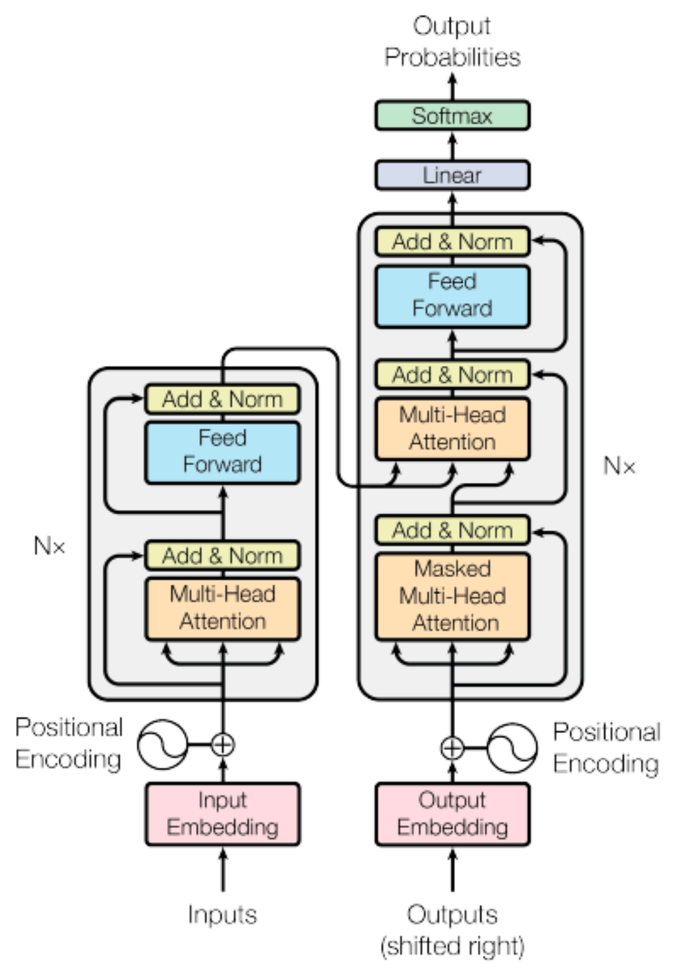

Examples

Models with context length

$$Y_i = f_\star(Y_{i-k:i-1}) + W_i, \quad i = 1,\ldots,n$$

Linear autoregressions:

$$X_{i+1} = \theta_\star X_i + W_i, \qquad i = 1,\ldots,n$$

$\theta_\star$ dynamics matrix

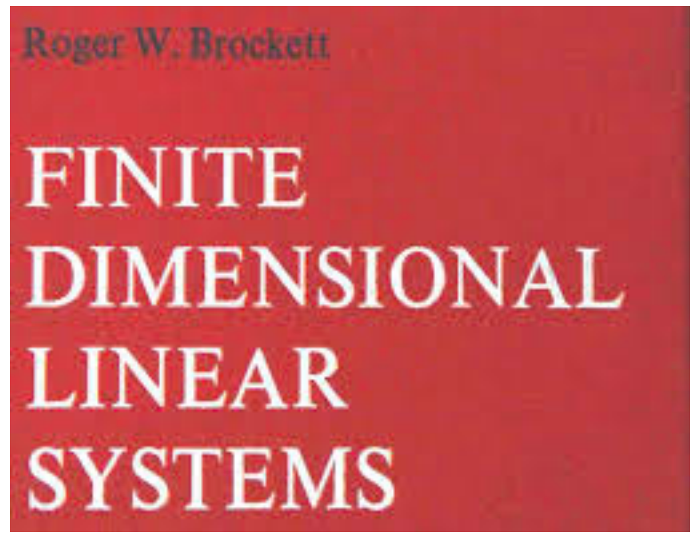

Learning from dependent data

How do we learn?

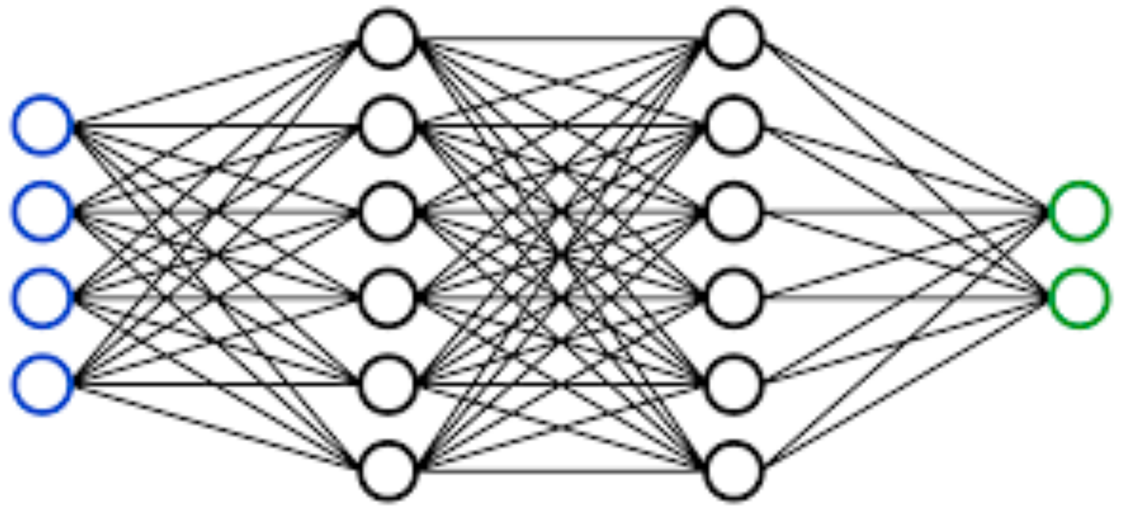

Often use empirical risk minimization (ERM), search over hypothesis class $\mathscr{F}$:

$$\hat{f} \in \operatorname{argmin}_{f \in \mathscr{F}} \frac{1}{n} \sum_{i=1}^{n} \mathtt{L}(f(X_i), Y_i)$$

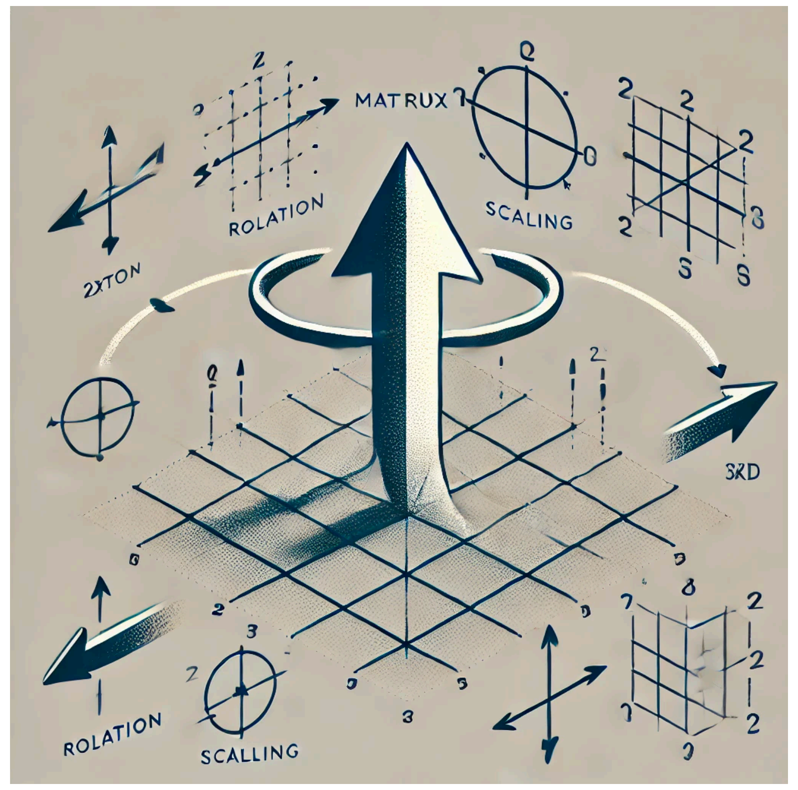

$\mathtt{L}$ a loss function, e.g., square loss $\mathtt{L}_{\mathrm{sq}}(y', y) \triangleq \|y' - y\|^2$

If $\mathscr{F}$ a finite dimensional linear space, this is just ordinary least squares:

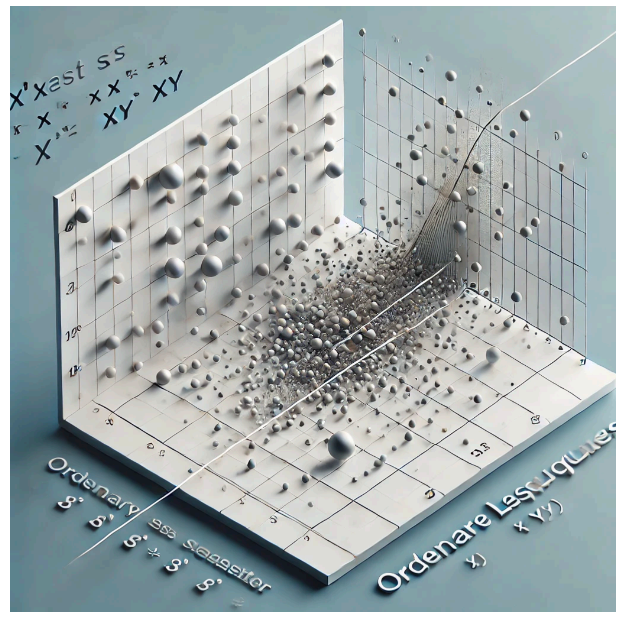

$$\hat{\theta} \in \operatorname{argmin}_{\mathbb{R}^{d_y \times d_x}} \frac{1}{n} \sum_{i=1}^{n} \|Y_i - \theta_\star X_i\|^2$$

Nontrivial with

Dependent data!

Control

X_i : state of your system

U_i : control inputs to your system

Will study learning to control the Linear Quadratic Regulator (LQR)

$$X_{i+1} = A_\star X_i + B_\star U_i + W_{i+1}, \qquad X_1 = W_0, \qquad i = 1, \ldots, n$$

Which consists of the above linear dynamics with quadratic costs:

$$\mathbb{V}_n^\pi \triangleq \mathbf{E}^\pi \left[X_n^\top Q_n X_n + \sum_{i=1}^{n-1} X_i^\top Q X_i + U_i^\top R U_i \right], \qquad Q, Q_n \succeq 0, R \succ 0$$

π : policy, the optimization variable we would like to learn from data

data : past observations of the X and U

Co

Ten

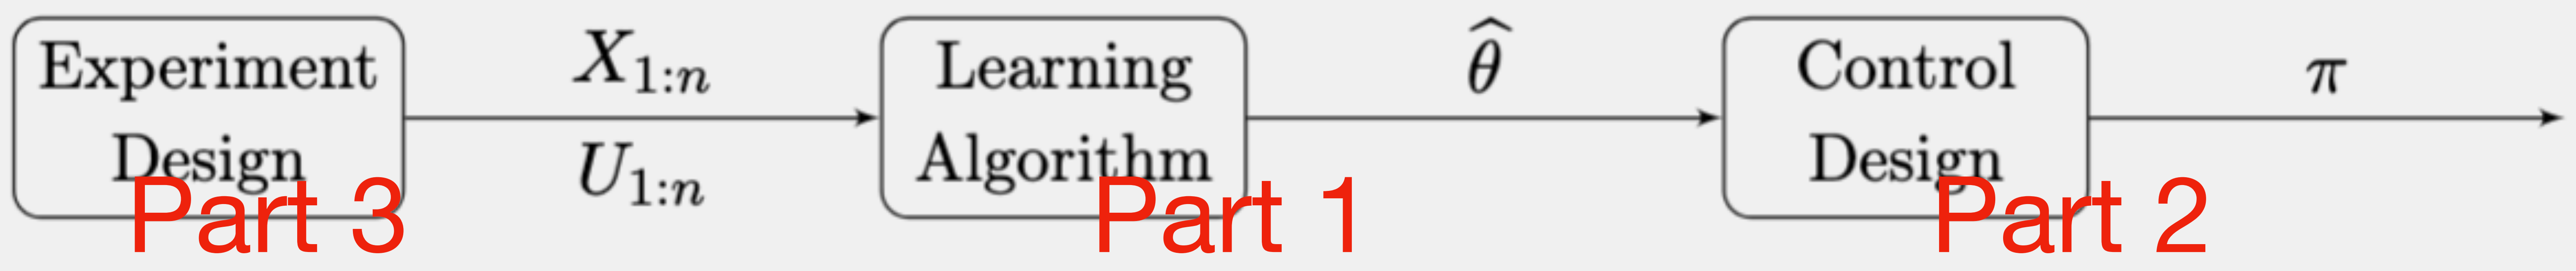

Figure 1.1.: A schematic overview of a model-based learning-to-control pipeline. Experimental data is fed into a learning algorithm, outputting an approximate model $\widehat{\theta}$, which is then used to design a control policy π.

Part 1: Foundations

- IID Mean Estimation, Linear Regression and Concentration Inequalities
- Covering numbers in learning in $\mathbb{R}^d$
- The Hanson-Wright Inequality (concentration with quadratic dependence)
- Linear Regression with Dependent Data, Linear System Identification

Part 2: Control

- LQR Recap
- Offline Learning of LQR
- Policy Gradient Methods for LQR

Part 3: Fundamental Limits and Active Learning

- Information-Theoretic Lower Bounds (AKA fundamental limits)
- Active Learning for LQR

Part 4: Further Topics

- Learning in nonlinear time-series/dynamics
- Martingale Methods
- More TBD if time permits (rep learning?)

Lecture notes loosely follow this structure

Part 1: Linear Regression with Dependence

Statistical Setup

Consider a time series model

$$Y_t = \theta^\star X_t + V_t, \quad t = 1, \ldots, T$$

benign noise

Where:

Y_t - Outputs in $\mathbb{R}^{d_Y}$

X_t - Covariates in $\mathbb{R}^{d_X}$

V_t - Noise in $\mathbb{R}^{d_Y}$

$\theta^\star$ - Unknown Parameter in $\mathbb{R}^{d_Y \times d_X}$

Example ARX(p,q):

$$Y_t = \sum_{i=1}^{p} A_i^\star Y_{t-i} + \sum_{j=1}^{q} B_i^\star U_{t-j} + W_t$$

In other words...

$$X_t = \begin{bmatrix} Y_{t-1:t-p}^\top & U_{t-1:t-q}^\top \end{bmatrix}^\top$$

$$\theta^\star = \begin{bmatrix} A_{1:p}^\star & B_{1:q}^\star \end{bmatrix}$$

$$V_t = W_t$$

Least Squares Estimation (LSE)

Consider a time-series model:

$$Y_t = \theta^\star X_t + V_t, \quad t = 1,\ldots,T$$

Least Squares Estimator:

$$\widehat{\theta} \in \operatorname{argmin}_{\theta \in \mathbb{R}^{d_Y \times d_X}} \left\{ \frac{1}{T} \sum_{t=1}^{T} \| Y_t - \theta X_t \|_2^2 \right\}$$

$\Rightarrow$

$$\widehat{\theta} = \left(\sum_{t=1}^{T} Y_t X_t^\top \right) \left(\sum_{t=1}^{T} X_t X_t^\top \right)^{-1}$$

Interested in:

$$\widehat{\theta} - \theta^\star = \left(\sum_{t=1}^{T} V_t X_t^\top \right) \left(\sum_{t=1}^{T} X_t X_t^\top \right)^{-1}$$

Part 1: Modern perspective on LSE

Draw on tools from:

- Machine Learning Theory
- High-Dimensional Statistics
- High-Dimensional Probability

Problem

Fix:

accuracy $\epsilon > 0$

failure probability $\delta \in (0,1)$

a norm $\|\cdot\|$

and a 'reasonable' estimator $\widehat{\theta}$

Establish finite sample guarantees:

$\|\widehat{\theta} - \theta^\star\| \leq \epsilon$ wpal. $1-\delta$

Typically we can prove:

$$\epsilon \propto (\text{noise scale}) \times \sqrt{\frac{\text{dimension} + \log(1/\delta)}{\text{sample size}}}$$

As long as:

$$\text{sample size} \gtrsim \text{dimension} + \log(1/\delta)$$

Persistence of Excitation

The Path Ahead

Covariance Concentration

At LLN Scale ~ 1

$$\hat{\theta} - \theta^\star = \left(\sum_{t=1}^{T} V_t X_t^\top\right)\left(\sum_{t=1}^{T} X_t X_t^\top\right)^{-1}$$

$$(\hat{\theta} - \theta^\star)\sqrt{\Gamma} = \frac{1}{T}\left[\left(\sum_{t=1}^{T} V_t X_t^\top\right)\Gamma^{-1/2}\right]\left(\frac{1}{T}\sum_{t=1}^{T} \Gamma^{-1/2} X_t X_t^\top \Gamma^{-1/2}\right)^{-1/2}$$

Random Matrix!

Random Walk at CLT Scale $\sim \sqrt{T}$

$$\Gamma \triangleq \frac{1}{T}\sum_{t=1}^{T} \mathbf{E}\left[X_t X_t^\top\right]$$

Todo:

CLT-analogue: Hanson-Wright

LLN-analogue: Also Hanson-Wright

Next Wednesday

We will look at something even simpler in more detail

Consider an iid model

$$Y_i = \theta^\star + V_i, \quad i = 1,\ldots,n$$

V_i - iid mean zero

What can say about estimating the mean $\theta^\star$?

$$\hat{\theta} \in \text{argmin}_{\theta \in \mathbb{R}^{d_Y}} \left\{ \frac{1}{n} \sum_{i=1}^{n} \|Y_i - \theta\|_2^2 \right\} \quad \Rightarrow \quad \hat{\theta} = \frac{1}{n} \sum_{i=1}^{n} Y_i \quad \Rightarrow \quad \hat{\theta} - \theta_\star = \frac{1}{n} \sum_{i=1}^{n} V_i$$

Understanding the non asymptotic behavior of $\frac{1}{n} \sum_{i=1}^{n} V_i$ requires **concentration inequalities**

i.e. we want to understand $\mathbf{P}\left(\frac{1}{n} \sum_{i=1}^{n} V_i > t \right) \leq ???$

Think of these as nonasymptotic CLT/LLN

More next week!