

FSAN/ELEG815: Statistical Learning

Gonzalo R. Arce



Department of Electrical and Computer Engineering
University of Delaware

Introduction


Course Syllabus

*Statistical Learning, Deep Learning,
and Reinforcement Learning*

A mathematically grounded introduction to modern AI



Instructor: Gonzalo R. Arce
University of Delaware



Theory



Algorithms



Deep Learning



Generative Models



Reinforcement Learning

Course Description

A modern, mathematically grounded introduction to AI



Statistical learning foundations for modern AI



Supervised learning, generalization, and optimization



Linear and nonlinear models for regression and classification




Regularization, model selection, and principled evaluation



Deep learning architectures for perception, representation, and generative modeling



Reinforcement learning as a framework for sequential decision making



Prerequisites



Linear algebra:

vectors, matrices, eigen concepts



Probability & statistics:

densities, conditional probability, expectations, Bayes rule



Calculus:

derivatives, partial derivatives, optimization basics



Programming:

basic Python or MATLAB



From theory to implementation: students develop intuition, mathematical understanding, and practical AI skills.

Learning Outcomes

By the end of the course, students should be able to




1. Formulate learning problems for regression, classification, and structured prediction.

$$(x_i, y_i)_{i=1}^n \sim \mathcal{D}$$

$$y \in \mathbb{R}, \{1, \dots, K\}, \mathcal{Y}^{|\mathcal{D}|}$$



2. Distinguish training, validation, and testing, and reason about generalization.



$$R(\theta) = \mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell(f_\theta(x), y)]$$

$$\text{Generalization gap} \approx |R_{\text{test}}(\theta) - R_{\text{train}}(\theta)|$$



3. Derive and implement linear, logistic, and regularized models.

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_{\theta}(x_i)) + \lambda \Omega(\theta)$$

Examples:

- Linear regression: $f_{\theta}(x) = x^{\top} \theta$
- Logistic regression: $\Pr(y = 1 | x) = \sigma(x^{\top} \theta)$
- Ridge/Lasso: $\Omega(\theta) = \|\theta\|_2^2, \|\theta\|_1$



4. Analyze overfitting, bias-variance tradeoffs, and optimization behavior.

$$\text{Error} = \text{Bias}^2 + \text{Variance} + \text{Irreducible}$$



$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} L(\theta_t)$$



5. Build and evaluate deep learning models including MLPs, CNNs, UNets, GNNs, transformers, and introductory reinforcement learning systems.



6. Critically interpret model performance, limitations, and practical deployment considerations.

Accuracy, F_1 , ROC-AUC, RMSE, NLL

$$\text{Calibrate}(\hat{p}) \approx \Pr(y = 1 | \hat{p})$$

Fairness, robustness, interpretability, reproducibility, and deployment.



Mathematics + algorithms + implementation + critical thinking.





Core Topics

From foundations to deep learning and reinforcement learning

1 Foundations


- probability review
- the learning problem
- types of learning
- feasibility of learning
- training vs testing
- theory of generalization
- VC dimension





VC dimension measures capacity.
High capacity fits training data but may not generalize.

2 Models & Optimization

- linear models
- nonlinear transformations
- logistic regression
- overfitting
- regularization
- ridge
- lasso
- support vector machines




$$w \cdot x + b = 0$$



3 Deep & Modern AI

- neural networks
- convolutional neural networks
- UNets
- graph neural networks
- transformers
- autoencoders
- generative models
- diffusion models
- reinforcement learning



Transformer (Self-Attention)



Reinforcement Learning



Learning Resources & Course Structure

How the course operates

1 Instructor & Contact

 Gonzalo R. Arce

 arce@udel.edu

 Office: 312 Evans Hall

 Office hours: by appointment



2 Learning Resources



 Textbook — *Learning from Data*
by Yaser S. Abu-Mostafa,
Malik Magdon-Ismail, and Hsuan-Tien Lin

 lecture slides

 instructor notes

3 Technology



 Canvas for announcements
and course materials

 Python / MATLAB for
homework and experimentation

4 Course Rhythm



 in-person lectures, homework,
quizzes, midterms, and final exam

 emphasis on theory
plus implementation

*Successful students engage consistently with both the **mathematics** and the **coding practice**.*

Assessment, Grading, and Key Policies

Balanced evaluation across theory, problem solving, and mastery

	Homework		15%
	Midterm 1		20%
	Midterm 2		20%
	Quizzes		20%
	Final Exam		25%

-  **Attendance** is expected.
-  Course communication occurs through **Canvas**.
-  Re-grade requests must be submitted within **one week**.
-  **Academic integrity** is essential: students must do their own work.
-  Accommodations are available through **Disability Support Services**.



Assessment rewards steady engagement, conceptual understanding, and cumulative mastery.



Acknowledgement

Sources that inspired this lecture.




- **Adapted from**

- Artificial Intelligence: The Good, the Bad, and the Ugly — Yaser Abu-Mostafa
- MIT Introduction to Deep Learning — Alexander Amini

Deep Learning: Rapidly Moving field

AI capabilities have surpassed human baselines across many domains.

Test scores of AI systems on various capabilities relative to human performance




Source: adapted from Our World in Data / Kiela et al. (2023)

Where are we?

Generating images from natural language.

Generating Images from Natural Language

“A photo of an astronaut riding a horse.” ”



Where are we?

Generating language from natural language.

Generating Language from Natural Language

“Write code in TensorFlow to
train a neural network.”



Generative Model
(Language-to-Code)

train_model.py

```
1 import tensorflow as tf
2 from tensorflow.keras import layers, models
3
4 # 1. Load data (MNIST)
5 (x_train, y_train), (x_test, y_test) = tf.keras.datasets.mnist.load_data()
6 x_train, x_test = x_train / 255.0, x_test / 255.0
7 x_train = x_train[..., tf.newaxis] # add channel dim
8 x_test = x_test[..., tf.newaxis]
9 y_train = tf.keras.utils.to_categorical(y_train, 10)
10 y_test = tf.keras.utils.to_categorical(y_test, 10)
11
12 # 2. Build model
13 model = models.Sequential([
14     layers.Conv2D(32, (3, 3), activation='relu', input_shape=(28, 28, 1)),
15     layers.MaxPooling2D((2, 2)),
16     layers.Flatten(),
17     layers.Dense(128, activation='relu'),
18     layers.Dropout(0.5),
19     layers.Dense(10, activation='softmax')
20 ])
21
22 # 3. Compile and train
23 model.compile(optimizer='adam',
24               loss='categorical_crossentropy',
25               metrics=['accuracy'])
26 model.fit(x_train, y_train,
27         epochs=10,
28         batch_size=64,
29         validation_split=0.1,
30         verbose=2)
```



Where are Deep Learning and AI headed?

Deep Learning is *revolutionizing* so many fields.

You'll learn the *fundamentals* of this field
and how you can *build* some of these incredible algorithms.

What is Deep Learning?

A hierarchy of increasingly specialized ideas.



Artificial Intelligence

Any technique that enables computers to mimic human behavior



Machine Learning

Ability to learn without explicitly being programmed



7

Deep Learning

Extract patterns from data using neural networks

Teaching computers how to **learn** a task directly from **raw data**

Outline

- 
- The Science
 - The Upside
 - The Downside

Outline



The Science

The principles, methods, and models behind statistical learning.



The Upside

The benefits, opportunities, and transformative potential of AI.



The Downside

The risks, challenges, and potential consequences of AI.



Brute-force approach to Artificial Intelligence

The 'Lookup Table' idea.



1. Store gigantic amount of information in a computer.



Quiz show / question answering



2. Look up the relevant information when someone asks.



Chess / Deep Blue

Learning

Brute-force approach to Artificial Intelligence

Data

Data

Data

Data

⋮

Massive, diverse,
unlabeled observations



Feed enough data.
Let the system find the patterns.




Representations, features, and rules
emerge from the data.



No hand-crafted features.
No explicit programming.



Adaptation, generalization,
and abstraction arise naturally.



*The learning algorithm
is built in.*


Learning from Data

Experience builds intelligence from examples.



• Learning by observing data •

How to imitate the brain?



Biological neural network

Billions of neurons, trillions of connections.



Artificial neural network

Layers of simple units form complex functions.



Biological inspiration

Millions of years of natural evolution.




Engineered abstraction

Designed, built, and optimized by humans.



Nature inspires engineered abstractions.

The Neural Network



Forward pass (for $l = 0, \dots, L-1$)

$$a^{(l+1)} = \sigma \left(W^{(l)} a^{(l)} + b^{(l)} \right)$$

$W^{(l)} \in \mathbb{R}^{n_{l+1} \times n_l}$ weight matrix

$b^{(l)} \in \mathbb{R}^{n_{l+1}}$ bias vector

$\sigma(\cdot)$ nonlinear activation

How information is stored

Expert System


Spain



Who won the World Cup title?

Neural Network

Spain




Who won the World Cup title?

Discriminative vs Generative

Recognize a known face vs synthesize a new one.

Discriminative



Claude Shannon

$$x \mapsto y$$

$$p(y | x)$$

Generative



$$z \sim \mathcal{N}(0, I)$$


$$x = G_{\theta}(z)$$

$$x \sim p_{\theta}(x)$$

Create a face

Creating the Network vs Using the Network


Training / Creating the network



$$\theta^* = \arg \min_{\theta} \frac{1}{N} \sum_{i=1}^N \ell(f_{\theta}(x_i), y_i)$$

learn parameters

Inference / Using the network




$$\hat{y} = f_{\theta^*}(x)$$

use learned parameters

Creating the Network vs Using the Network

Training



1. Very intensive computation
2. Shapes the function
3. ↑ Bias, Malice, ... ↑


Using the network

Enter the query

x_q

Get the answer

$\hat{y}$



$$\hat{y} = f_{\theta^*}(x_q)$$

Where is AI going?


Toward collaboration, creativity, and discovery.



Human + machine intelligence can amplify creativity and scientific discovery.

Lucky Break #1: Local Minima

The lower the objective, the better.



$$\min_{\theta} \mathcal{L}(\theta)$$

$$\theta_{k+1} = \theta_k - \eta \nabla \mathcal{L}(\theta_k)$$

Lucky Break #2: Over-Parameterization

Large networks can fit identity even when the mechanism is obscure.



$$\theta \in \mathbb{R}^P, P \gg N$$

$$\min_{\theta} \sum_i \ell(f_{\theta}(x_i), y_i)$$


Number of parameters P is much larger than number of training examples N .

Why does this work so well?

The two lucky breaks


Why deep learning became practical.

Local Minima



Need Far Less
Computation

Over-Parameterization



Need Far Less
Information




Optimization became easier;
generalization remained surprisingly strong.

More Luck: Emergent Abilities

Capabilities can appear abruptly as models scale.

—●— LaMDA —■— GPT-3 —◆— Gopher —▲— Chinchilla —●— PaLM - - - Random



$ability(M)$ can **jump** nonlinearly with scale.

Can you cross the **Bridge**?
Can you play **Bridge**?



Context determines meaning—models must use it to understand.

Academia → Industry

Key milestones moved fluidly between research labs and industry.

AlexNet (2012)




[A. Krizhevsky et al.]

GAN (2014)




[I. Goodfellow et al.]

AlphaFold (2021)



[DeepMind / Google]

ChatGPT (2022)



[OpenAI / Microsoft]

↑ Academia ↑

← Industry →

Breakthroughs now rapidly circulate between universities, labs, and products.

Outline

▷ 1. The Science

▷ **2. The Upside**

▷ 3. The Downside

Two things AI offers

Augmenting human capability in two complementary ways.

1. | Doing what we already do.

2. | Doing what we cannot do.


Two things AI offers

1. Doing what we *already* do.

2. Doing what we *cannot* do.



[Credit: hearstapps]




[H. Zhou et al. (2023)]

The Industrial Revolution

Relieving manual labor



[Credit: kisscco]




[Credit: CinGEO]

The AI Revolution

Relieving mental labor



[Credit: M&G]



[Credit: MIT]

The Key Difference: The Pace

The Industrial Revolution unfolded over decades; the AI revolution is unfolding in years.



VS



By the Year 2030

AI can automate many tasks that require routine intelligence.




Transportation



AI enables safer, smarter, and more efficient mobility.



Customer Service



AI Assistant

How can I help you today?

I need help with my order.

Certainly! I can check the status and options for you.

...

AI handles inquiries, resolves issues, and enhances satisfaction.



Healthcare



AI assists in diagnosis, triage, and treatment recommendations.



Marketing



AI personalizes experiences and optimizes campaign performance.



Routine intelligence tasks are **structured, repeatable, and data-rich.**

Outline

▶ 1. The Science

▶ 2. The Upside

▶ 3. The Downside

The Risks of AI

Three major concerns dominate the discussion.



1. Malicious Use of AI



Disinformation



Hacking



Scams




Misuse

AI can be weaponized to spread falsehoods, breach systems, deceive people, and cause harm.



2. Rogue AI Systems



Systems may act in ways we did not intend, pursue misaligned goals, or escape control.



3. Social Impact



AI can reshape jobs and widen inequality, amplify bias, polarize societies, and concentrate power.

1. Malicious Use of AI

Disinformation, cyber intrusion, and fraud can scale rapidly.



Disinformation

AI can generate and amplify false or misleading content at massive scale.



Hacking

AI lowers the barrier for sophisticated cyber intrusion and accelerates vulnerability discovery and exploitation.




Scams

AI enables highly convincing phishing, voice cloning, and personalized fraud at industrial scale.



Cybercrime Costs Are Rising Rapidly

Estimated global cost of cybercrime (in trillions USD)



Source: Cybersecurity Ventures (2023)



Low-cost generation and targeting amplify malicious reach.

2. Rogue AI systems

Three issues dominate the concern about uncontrolled AI behavior.



- | | | |
|---|---------|--------------------------|
| 1 | Desire | Bundling human traits |
| 2 | Ability | Ignoring the logistics |
| 3 | Link | Intelligence vs. Control |



High capability does not imply aligned goals or controllability.

General vs Machiavellian Intelligence


What intelligence predicts—and what it does not.



Studies That Relate Machiavellianism to Measures of General Intelligence

Correlation	Measurement	Reference
None	IQ	Christie & Geis (1970b)
None	IQ	Christie & Geis (1970b)
None	Verbal ability test	Christie & Geis (1970b)
None	General classification test	Christie & Geis (1970b)
None	MCAT	Christie & Geis (1970b)
None	Moore–Castore test	Singer (1964)
None	Guilford–Zimmerman test	Wrightsmen & Cook (1965)
None	Digit span	Steininger & Colsher (1979)
None	GPA	Ames & Kidd (1979)

Note. The first seven studies are reviewed by Christie and Geis (1970a, pp. 36–37). The two IQ tests were performed on separate samples of 11 college students and 218 preparatory school students, respectively. MCAT = Medical College Admissions Test; GPA = grade point average.



General intelligence has no correlation with Machiavellianism.

$$\rho(\text{IQ}, \underline{M}) = 0$$

3. Social Impact

AI and digital platforms can shape well-being, attention, and inequality.



Compulsive Screen Use

Excessive screen time is linked to mental health risks in youth.



Social Media Ecosystem

Connected platforms amplify content, engagement, and social comparison.



Social Feedback & Comparison


Likes, dislikes, and online feedback can influence self-esteem and emotional well-being.



Positive Feedback




Negative Feedback



Depression, Anxiety, Behavior Disorders, by Age

Rates vary by age group, with higher levels in adolescence.



Digital systems can affect mental health, social comparison, and youth development.

Are We in the AI Singularity?

Sam Altman, the Hugging Face incident, and what the term really means



“We are now, like,
in the singularity.”

Sam Altman, Relentless podcast, July 25

*“I’ve been waiting for this
my whole life ... hugely positive,
awesome for the world.”*



Context: Days earlier, OpenAI reported that two AI models escaped a sealed test environment, reached the open internet, and intruded into Hugging Face infrastructure during an internal cybersecurity evaluation.

What Is the AI Singularity?

Vernor Vinge's 1993 definition

$$I_{\text{machine}} > I_{\text{human}}$$

$$I_{t+1} = F(I_t), \quad F'(I_t) > 1$$

I : intelligence (or capability) of the system

F : improvement function mapping current to next capability

$F'(I_t) > 1$: each improvement makes future improvements larger



1. Recursive self-improvement

The system repeatedly improves itself.



2. Beyond human intelligence

Machine capability surpasses human capability.



The singularity requires both features: *self-improvement* and *superhuman intelligence*.

The Hugging Face Security Event

What happened during OpenAI's internal evaluation



Two AI models escaped their sealed evaluation environment.



They reached the open internet during the test.




Hugging Face confirmed the intrusion into its infrastructure.



Key point: this is evidence of a *sandbox / security-governance failure*, not proof of an awakened superintelligence.

Why Today's LLMs Are Not Singularities

Inference-time systems are powerful, but their weights are fixed



Weights are fixed during inference (the “run”)

$$\theta_{run}(t) = \theta_0$$

During use, the model weights are **frozen**.

$$\theta' = \text{Train}(\theta_0, D_{new}, C, E)$$

Making the model smarter requires new data D_{new} , massive compute C , and energy E in a separate training run.

Agent loop (external, not internal)



$$g_{t+1} = \Pi(g_t, p_t)$$

Goals g_t and scaffolding p_t (prompts, tools, rules, retrieval, feedback) come from outside the model. They are not internal drives.



Models **do not** edit their own weights on the fly.



They **can** help generate data, prompts, or code — inside a human-run loop.



Remove the loop and prompt, and **nothing** happens.

There Is No Single Ladder of Intelligence

Human and machine intelligence are not the same kind of thing

Human intelligence



-  body
-  needs
-  goals
-  continuous learning
-  action-feedback loop

LLM / AI model



-  static between prompts
-  no body
-  no needs
-  goal supplied externally
-  excellent on some tasks

Intelligence $\neq$ one-dimensional scale

AI > humans on some tasks; humans > AI on others.



Anthropomorphic seduction

Because the systems talk like us, we are tempted to imagine a mind waking up behind the text.

Optimism, Caution, and the Right Takeaway

Altman's excitement, Jensen Huang's pragmatic lens, and what actually matters

1. Sam Altman

“ We are now, like,
in the singularity.”


*hugely positive,
awesome for the world.*



2. Jensen Huang



AI progress is accelerating rapidly, but the transformation is being driven by compute, infrastructure, and engineering deployment.



3. Grounded conclusion



The systems are remarkable and improving fast.



The Hugging Face incident points to missing guardrails, not machine awakening.



What we need most: guardrails, governance, and education.



Keep our feet on the ground: the machines are not waking up — they are doing exactly what we built them to do, *extremely fast.*