

Machine Learning for Modern Artificial Intelligence

Hsuan-Tien Lin

htlin@csie.ntu.edu.tw

Dept. of Computer Science and Information Engineering
National Taiwan University



Guest Lecture for AI and Digitalization of Law
March 19, 2026

About Me

Professor
National Taiwan University



(former) Chief Data Scientist
Appier Inc.



The logo for Appier Inc., consisting of the word "Appier" in a blue, sans-serif font.



Co-author
Learning from Data



Instructor
NTU-Coursera MOOCs
ML Foundations/Techniques



research goal: making machine more realistic

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?

How Can Machines Learn?

Why Can Machines Learn?

How to Use Machine Learning?

Where Will AI Go?

From Intelligence to Artificial Intelligence

intelligence: thinking and acting **smartly**

- **humanly**
- **rationally**

artificial intelligence: **computers** (thinking and) acting **smartly**

- **humanly**
- **rationally**

humanly \approx **smartly** \approx **rationally**
—**are humans rational?** 😊

Humanly versus Rationally

What if your self-driving car decides one death is better than two—and that one is you? (The Washington Post <http://wpo.st/ZK-51>)

You're humming along in your self-driving car, chatting on your iPhone 37 while the machine navigates on its own. Then a swarm of people appears in the street, right in the path of the oncoming vehicle.

Car Acting **Humanly**

to save my (and passengers')
life, stay on track

Car Acting **Rationally**

avoid the crowd and crash the
owner for **minimum total loss**

which is **smarter**?
—depending on where I am, maybe? 😊

Traditional vs. Modern [My] Definition of AI

Traditional Definition

humanly \approx intelligently \approx rationally

My Definition

intelligently \approx easily

is your smart phone 'smart'? 😊

modern artificial intelligence
= **application** intelligence

Examples of Application Intelligence

Siri




By Bernard Goldbach [CC BY 2.0]

iRobot



By Yuan-Chou Lo [CC BY-NC-ND 2.0]

Amazon Recommendations




By Kelly Sims [CC BY 2.0]

Vivino



From nordic.businessinsider.com

AI Milestones



- first AI winter: AI cannot solve 'combinatorial explosion' problems
- second AI winter: expert system failed to scale

reason of winters: **expectation mismatch**

What's Different Now?

More Data

- cheaper storage
- Internet companies

Better Algorithms

- decades of research
- e.g. deep learning

Faster Computation


- cloud computing
- GPU computing

Healthier Mindset

- reasonable wishes
- key breakthroughs

data-enabled AI: mainstream nowadays

Bigger Data Enable Easier-to-use AI



By deepanker70 on <https://pixabay.com/>

past

best route by
shortest path

present

best route by
current traffic

future

best route by
predicted travel time

big data **can** make machine look smarter

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?

How Can Machines Learn?

Why Can Machines Learn?

How to Use Machine Learning?

Where Will AI Go?

From Learning to Machine Learning

learning: acquiring **skill**
with experience accumulated from **observations**



machine learning: acquiring **skill**
with experience accumulated/**computed** from **data**




What is **skill**?

A More Concrete Definition


skill

⇔ improve some performance measure (e.g. prediction accuracy)

machine learning: improving some performance measure with experience **computed** from data



Application 1: Computational Finance



Why use machine learning?

Machine Learning Connects Big Data and AI

From Big Data to Artificial Intelligence



ingredient



tools/steps



dish



Photos Licensed under CC BY 2.0 from Andrea Goh on Flickr

Chief Data Scientist for AI
≈ **Master Chef for Restaurant**

Why Use Machine Learning?



- 'define' trees and hand-program: **difficult**
- learn from data (observations) and recognize: a **3-year-old can do so**
- 'ML-based tree recognition system' can be **easier to build** than hand-programmed system

ML: a **promising alternative** to build complicated systems

Application 2: Student Performance Prediction



- **data**: students' records on quizzes on a Math tutoring system
- **AI**: predict whether a student can give a correct answer to another quiz question

A Possible ML Solution

answer correctly \approx \llbracket recent **strength** of student $>$ **difficulty** of question \rrbracket

- give ML **9 million records** from **3000 students**
- ML determines (**reverse-engineers**) **strength** and **difficulty** automatically

key part of the **world-champion** system from National Taiwan Univ. in KDDCup 2010


Good AI Needs Both ML and Non-ML Techniques



(Public Domain, from Wikipedia; used here for education purpose; all other rights still belong to Google DeepMind)

Non-ML Techniques


Monte C. Tree Search
≈ **move simulation** in brain



(CC-BY-SA 3.0 by Stannered on
Wikipedia)

ML Techniques

Deep Learning
≈ **board analysis** in human brain



(CC-BY-SA 2.0 by Frej Bjon on
Wikipedia)


Reinforcement Learn.
≈ **(self)-practice** in human training



(Public Domain, from Wikipedia)

good AI: important to use the **right**
techniques—ML & **others, including human**

Full Picture of ML for Modern AI



Human Learning

- subjective
- produce domain knowledge
- fast basic solution

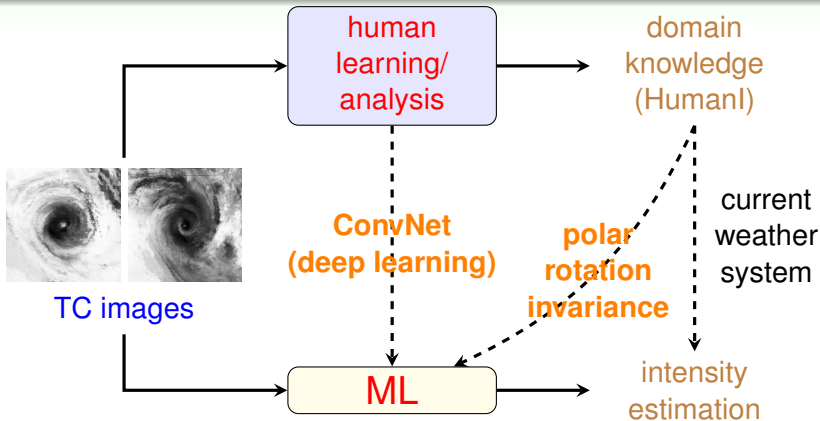
Machine Learning

- objective
- leverage computing power
- continuous improvement

tip: **use humans as much as possible first**
before going to machines

Application 3: Tropical Cyclone Intensity Estimation

meteorologists can 'feel' & estimate TC intensity from image



better than current system & **'production-ready'**
 (Chen et al., KDD '18; Chen et al., Weather & Forecasting '19)

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?

How Can Machines Learn?


Why Can Machines Learn?

How to Use Machine Learning?

Where Will AI Go?

Key Essence of Machine Learning

machine learning: improving some performance measure with experience **computed** from data



- 1 exists some 'underlying pattern' to be learned
—so 'performance measure' can be improved
- 2 but no programmable (easy) definition
—so 'ML' is needed
- 3 somehow there is data about the pattern
—so ML has some 'inputs' to learn from

key essence: help decide whether to use ML


Application 4: Credit Approval

applicant information

| | |
|-------------------|---------------|
| age | 23 years |
| gender | female |
| annual salary | NTD 1,000,000 |
| year in residence | 1 year |
| year in job | 0.5 year |
| current debt | 200,000 |




approve credit card or not?



(Picture Licensed under CC0 on Pixabay)

unknown pattern to be learned:
 'approve credit card good for bank?'

Learning Flow for Credit Approval




- target **unknown**
(i.e. no programmable definition)
- hypothesis hopefully similar to
but possibly **different** from target
(perfection ‘impossible’ when target
unknown)

What is a **learning model**?


Learning Model 1: Linear Model

debt



income

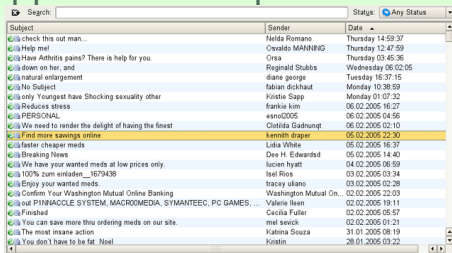
debt



income

linear model: find a 'line' (final hypothesis)
to separate ●(approve) and ●(disapprove)
(as much as possible) on **training data**

Application 5: Spam Filtering




(Licensed under GPL via Wikimedia Commons by Ascánder)

| | spam | non-spam |
|-------|-----------------------------------|---|
| words | free drug fantastic deal | machine learning statistics textbook |

simple **multi-dimensional** linear model:
checking if we have more **non-spam**-like words
than **spam**-like ones


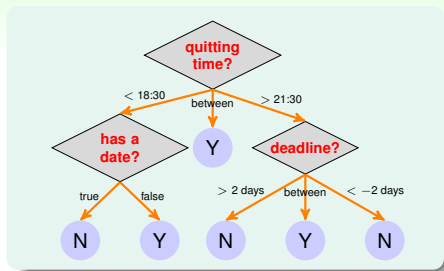
Historical Note: Perceptron Learning Algorithm (Rosenblatt, 1957)



- perceptron (linear model): mathematically simulating **single neuron** (for probabilistic logical inference)
- **'first' machine learning model**


perceptron learning alg.: its limitations (Minsky and Papert, 1969) partially caused **first winter**

Learning Model 2: Decision Tree



decision tree: finding **conditionally composite rules** automatically by machines

Historical Note: (C4.5) Decision Tree (Quinlan, 1986)



- mimicking **expert systems**
- still popular in **data mining, business decision making**

decision tree: finding its value after **second winter** in 'non'-AI applications


From Perceptrons to Neural Networks

perceptron (one neuron):

neural network (many neurons):


linear boundary

complicated boundary (and more)



by UC Regents Davis campus-brainmaps.org.

Licensed under CC BY 3.0 via Wikimedia Commons



by Lauris Rubenis.
Licensed under CC BY
2.0 via

<https://flic.kr/p/fkVuZX>





by Pedro Ribeiro
Simões. Licensed


under CC BY 2.0 via
<https://flic.kr/p/adiV7b>

neural network: **bio-inspired** model

Learning Model 3: Deep Learning



performance



deep learning:
leverage **huge data**+**huge computation**
to solve complicated ML tasks

Application 6: Face Recognition




original picture by F.U.S.I.A. assistant and derivative work by Sylenius via Wikimedia Commons

face recognition

- data: **faces and non-faces**
- AI: **locate faces**


mature traditional & **deep learning** models exist, but need **tuning** for **application needs**

Historical Note: deep learning (Hinton et al., 2006)




deep learning: rebranding of neural networks
after **second winter**

Historical Note: from Checkers to Go



(Samuel, 1959) Some studies in machine learning using the game of checkers



Picture extracted from the original paper of Samuel for educational purposes

(Silver et al., 2016) Mastering the game of Go with deep neural networks and tree search




Lee Sedol (B) vs AlphaGo (W) - Game 1


Picture by Wesalius, licensed under CC BY-SA 4.0 via Wikimedia Commons

Game AI witnessed
rising of **machine learning**

Historical Note: from Recognition to Generation




(Krizhevsky et al., 2012)
ImageNet
classification
with deep
convolutional
neural networks



| motor scooter | leopard |
|---------------|--------------|
| motor scooter | leopard |
| go-kart | jaguar |
| moped | cheetah |
| bumper car | snow leopard |
| golfcart | Egyptian cat |

Picture extracted from the original paper
of Krizhevsky et al. for educational purposes

(Ho et al., 2020)
Denoising
diffusion
probabilistic
models



Picture extracted from the original paper
of Ho et al. for educational purposes

deep learning facilitated **modern AI**

Generative AI/ML

recognition

listen/read/see

generation

speak/write/draw

specialty of generative AI

variation (creativity)



(Pictures Extracted from Ho et al. for educational purposes)

complication (structure)



(Pictures Licensed under CC0 on Wikipedia)

generative AI/ML: models that can make **varying** and **complicated** outputs

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?




How Can Machines Learn?

Why Can Machines Learn?




How to Use Machine Learning?

Where Will AI Go?


Machine/Human Learning Exercise



$$y_n = -1$$



$$y_n = +1$$




$$g(\mathbf{x}) = ?$$

What did your human learning output?

Two Controversial Answers

whatever you say about $g(\mathbf{x})$,


 $y_n = -1$

 $y_n = +1$

 $g(\mathbf{x}) = ?$

truth $f(\mathbf{x}) = +1$ because ...

- symmetry $\Leftrightarrow +1$
- (black or white count = 3) or (black count = 4 and middle-top black) $\Leftrightarrow +1$

truth $f(\mathbf{x}) = -1$ because ...

- left-top black $\Leftrightarrow -1$
- middle column contains at most 1 black and right-top white $\Leftrightarrow -1$

all valid reasons, your **adversarial teacher** can always call you 'didn't learn'. :-)

What is the Next Number?

1,4,1,5

What is the Next Number?

1,4,1,5

1,4,1,5,**0**,-1,1,6by $y_t = y_{t-4} - y_{t-2}$ 1,4,1,5,**1**,6,1,7by $y_t = y_{t-2} + \llbracket t \text{ is even} \rrbracket$ 1,4,1,5,**2**,9,3,14by $y_t = y_{t-4} + y_{t-2}$ **any number** can be the next!

Feasibility of Machine Learning



If **exam is similar to textbook**,
and you **understand textbook thoroughly**,
generalize to exam,
then you'll get good exam score



(Pictures Licensed under CC0 on Pixabay)

machine learning is feasible if

- 1 training data and test scenario **related**
- 2 **training** performance **sufficiently good**
- 3 **generalize** from training to testing

good training: **optimization**;
good testing: **generalization**
dilemma: two wishes often **contradicting**

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?

How Can Machines Learn?

Why Can Machines Learn?

How to Use Machine Learning?

Where Will AI Go?

Applications in This Talk



- Application 1: Computational Finance
- Application 2: Student Performance Prediction
- Application 3: Tropical Cyclone Intensity Estimation
- Application 4: Credit Approval
- Application 5: Spam Filtering
- Application 6: Face Recognition

machine learning has
facilitated **application intelligence**

Application 7: Image Generation



an astronaut riding a horse, by Hiroshige

(Picture in Public Domain on Wikipedia)

- new **photoshop**
- representative of modern generative AI

becoming **faster and more controllable**, with
new social issues for **copyright** and **faking**

Application 8: Chat System



(Snapshot from chat.openai.com in my class for educational purposes)

- machine learning technique: **attending** key context to do **next-word generation**
- with huge data/model/computation to reach **amazing performance**

evolving from chat systems
to application **agents**

The Next Topic Is

What is Artificial Intelligence?

What is Machine Learning?

How Can Machines Learn?

Why Can Machines Learn?

How to Use Machine Learning?

Where Will AI Go?

From Wright Flyer (1903)

to Commercial Airplanes (1919–)



by Wright Brothers.
Licensed under Public
Domain via
US Library of Congress



by Pedro Ribeiro
Simões. Licensed
under CC BY 2.0 via
[https://flic.kr/
p/adiV7b](https://flic.kr/p/adiV7b)

we are at **wright-flyer-age** of (generative) AI

What's Needed before Wider Acceptance



by Wright Brothers.
Licensed under Public
Domain via
US Library of Congress

- war? 😊
- technology advancements
—like lighter materials, more efficient engines, **better control**
- regulations
—like **laws**, licenses, etc.
- trials
—understanding success **and failure** cases

our work is not done!!

Civil Liability of Self-Driving Car (Courtesy of Attorney Yi-Hong Hsiao)

If Driver A is driving an autonomous vehicle and injures pedestrian B who ran a red light, who should be held responsible?

Tesla Autopilot Case 1: 2023

- Reuters reports the jury rejected claims from the plaintiff, Justine Hsu, that her Tesla Model S's Autopilot system didn't perform safely and awarded her zero damages.

Tesla Autopilot Case 2: 2026

- US judge upholds 243 million verdict against Tesla over fatal Autopilot crash.

AI will generate **new needs and liabilities**

Criminal Offense in AI Era

<https://www.youtube.com/watch?v=-a61zsRRONc>

Yaser Abu-Mostafa, Artificial Intelligence: The Good, the Bad, and the Ugly



1. Malicious use of AI



[Credit: iStock]




[Credit: Simplicam]



[Credit: adobe]

Cybercrime Expected To Skyrocket in the Coming Years

Estimated cost of cybercrime worldwide (in billion U.S. dollars)



[Credit: Statista]

doing that you use AI, your penalty will

26/30

artificial intelligence: like **guns**, may make it **easier** to conduct some crime

Conclusion



Photos Licensed under CC BY 2.0 from Andrea Goh on Flickr

- What is Artificial Intelligence? **easy-to-use application** intelligence
- What is Machine Learning? **data-driven** tools for achieving AI
- How Can Machines Learn? from linear to **deep** models
- Why Can Machines Learn? **generalize** with train/test relationship
- How to Use Machine Learning: prediction, recognition, **generation**
- Where Will AI Go? new **opportunities and challenges**

Thank you!