

Machine Learning for Modern Artificial Intelligence

Hsuan-Tien Lin
林軒田

Dept. of Computer Science and Information Engineering,
National Taiwan University
國立臺灣大學資訊工程學系

July 29, 2021
Winbond Electronics Corp.
華邦電子高階主管講座




About Me

Professor
National Taiwan University



Chief Data Science Consultant
(former Chief Data Scientist)
Appier Inc.



Appier



Co-author
Learning from Data



Instructor
NTU-Coursera MOOCs
ML Foundations/Techniques



Outline

ML for (Modern) AI

ML Research for Modern AI

ML for AI in Reality

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) Artificial Intelligence

Thinking Humanly

- cognitive modeling
—now closer to Psychology than AI

Thinking Rationally

- formal logic—now closer to Theoreticians than AI practitioners

Acting Humanly

- dialog systems
- humanoid robots
- computer vision

Acting Rationally

- recommendation systems
- cleaning robots
- character recognition

acting humanly or rationally:
more academia/industry attention nowadays

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**

AI: Now and Next

2010–2015: AI |

AI becomes
promising, e.g.

- initial success of **deep learning** on ImageNet
- mature tools for SVM (**LIBSVM**) and others

2016–2020: AI +

AI becomes
competitive, e.g.

- super-human performance of **alphaGo** and others
- all big technology companies become **AI-first**

2021–: AI ×

AI becomes
necessary

- “You’ll not be replaced by AI, but **by humans who know how to use AI**”
(Sun, Chief AI Scientist of Appier, 2018)

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

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

many possibilities when
using the right tools

ML-based AI Applications (1/6): Medicine



By DataBase Center for Life Science;
licensed under CC BY 4.0 via Wikimedia Commons

for computer-assisted diagnosis

- data:
 - **patient status**
 - **past diagnosis from doctors**
- **AI**: dialogue system that **efficiently identifies disease of patient**

my student's earlier work
as intern @ HTC DeepQ

ML-based AI Applications (2/6): Communication



By JulianVilla26;

licensed under CC BY-SA 4.0 via Wikimedia Commons

for 4G LTE communication

- **data**:
 - **channel information** (the channel matrix representing mutual information)
 - **configuration** (precoding, modulation, etc.) that reaches the highest throughput
- **AI**: predict **best configuration to the base station** in a new environment

my student's earlier work as intern @ MTK

ML-based AI Applications (3/6): Education



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

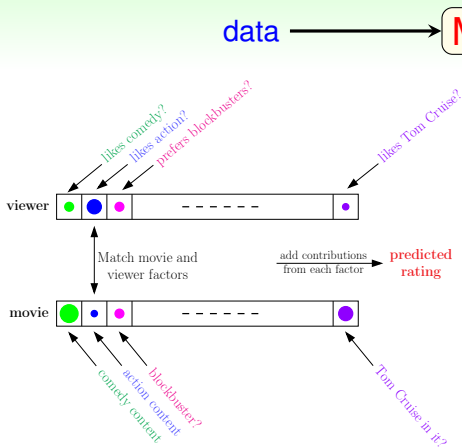
A Possible ML Solution

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


ML-based AI Applications (4/6): Entertainment



- **data**: how many users have rated some movies
- **AI**: predict how a user would rate an unrated movie

world-champion (again!) from
National Taiwan Univ. in KDDCup 2011

ML-based AI Applications (5/6): Manufacturing



By Raimond Spekking;
licensed under CC BY-SA 4.0 via Wikimedia Commons

for PCB fault detection

- data: **PCB images of normal and abnormal PCBs**
& maybe human-marked faulty locations
- AI: predict **which PCBs are faulty**

ongoing research for smart factory

ML-based AI Applications (6/6): Security



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: predict **which boxes contain faces**

mature **ML technique**, but often need **tuning**
for different **application intelligence** needs


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
 \approx **move simulation** in
 brain



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

ML Techniques

Deep Learning
 \approx **board analysis** in
 human brain



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

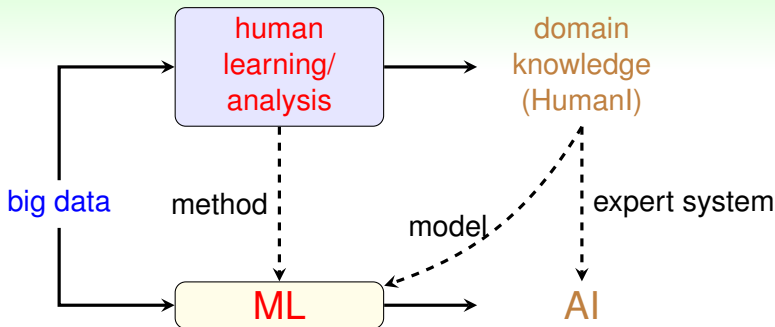
Reinforcement Learn.
 \approx **(self)-practice** in
 human training



(Public Domain, from Wikipedia)

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


Full Picture of ML for Modern AI



industry: **black plum is as sweet as white**

Example: Tropical Cyclone Intensity Estimation

meteorologists can 'feel' & estimate TC intensity from image



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

Outline

ML for (Modern) AI

ML Research for Modern AI

ML for AI in Reality

Cost-Sensitive Multiclass Classification

What is the Status of the Patient?



?

By DataBase Center for Life Science;
licensed under CC BY 4.0 via Wikimedia Commons



COVID19



cold



healthy

Pictures Licensed under CC BY-SA 3.0 from 1RadicalOne on Wikimedia Commons

- a **classification** problem
—grouping ‘patients’ into different ‘status’

are all mis-prediction costs equal?

Patient Status Prediction

error measure = society cost

predicted \ actual	COVID19	cold	healthy
COVID19	0	1000	100000
cold	100	0	3000
healthy	100	30	0

- COVID19 mis-predicted as healthy: **very high cost**
- cold mis-predicted as healthy: **high cost**
- cold correctly predicted as cold: **no cost**

human doctors consider costs of decision;
how about computer-aided diagnosis?

Our Works

	binary	multiclass
regular	well-studied	well-studied
cost-sensitive	known (Zadrozny et al., 2003)	ongoing (our works, among others)

selected works of ours

- cost-sensitive SVM (Tu and Lin, ICML 2010)
- cost-sensitive one-versus-one (Lin, ACML 2014)
- cost-sensitive deep learning (Chung et al., IJCAI 2016)

why are people **not**
using those **cool ML works for their AI?** 😊

Issue 1: Where Do Costs Come From?


A Real Medical Application: Classifying Bacteria

- by human doctors: **different treatments** \iff serious costs
- cost matrix averaged from two doctors:

	Ab	Ecoli	HI	KP	LM	Nm	Psa	Spn	Sa	GBS
Ab	0	1	10	7	9	9	5	8	9	1
Ecoli	3	0	10	8	10	10	5	10	10	2
HI	10	10	0	3	2	2	10	1	2	10
KP	7	7	3	0	4	4	6	3	3	8
LM	8	8	2	4	0	5	8	2	1	8
Nm	3	10	9	8	6	0	8	3	6	7
Psa	7	8	10	9	9	7	0	8	9	5
Spn	6	10	7	7	4	4	9	0	4	7
Sa	7	10	6	5	1	3	9	2	0	7
GBS	2	5	10	9	8	6	5	6	8	0

issue 2: is cost-sensitive classification
really useful?

Cost-Sensitive vs. Traditional on Bacteria Data




(Jan et al., BIBM 2011)

cost-sensitive better than **traditional**;
but why are people **still not**
using those cool ML works for their AI? 😊

Issue 3: Error Rate of Cost-Sensitive Classifiers

The Problem



- cost-sensitive classifier: **low cost but high error rate**
- traditional classifier: **low error rate but high cost**
- how can we get the **blue** classifiers?: **low error rate and low cost**

cost-and-error-sensitive:
more suitable for **real-world medical needs**

Improved Classifier for Both Cost and Error

(Jan et al., KDD 2012)

Cost

iris	≈
wine	≈
glass	≈
vehicle	≈
vowel	○
segment	○
dna	○
satimage	≈
usps	○
zoo	○
splice	≈
ecoli	≈
soybean	≈

Error

iris	○
wine	○
glass	○
vehicle	○
vowel	○
segment	○
dna	○
satimage	○
usps	○
zoo	○
splice	○
ecoli	○
soybean	○

now, **are people using those cool ML works for their AI?** 😊

Lessons Learned from Research on Cost-Sensitive Multiclass Classification



?



COVID19



cold



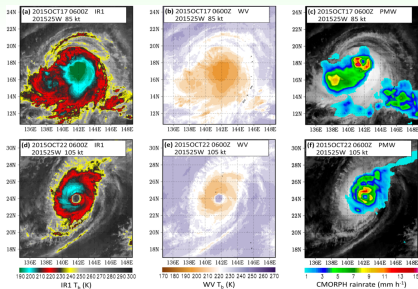
healthy

See Page 16 of the Slides for Sources of the Pictures

- 1 more realistic (generic) in academia
 \neq **more realistic (feasible) in application**
e.g. the 'cost' of **inputting a cost matrix?** 😊
- 2 **cross-domain collaboration** important
e.g. getting the 'cost matrix' from **domain experts**
- 3 not easy to win **human trust**
—humans are somewhat **multi-objective**

Tropical Cyclone Intensity Estimation

Experienced Meteorologists Can 'Feel' and Estimate Tropical Cyclone Intensity from Image



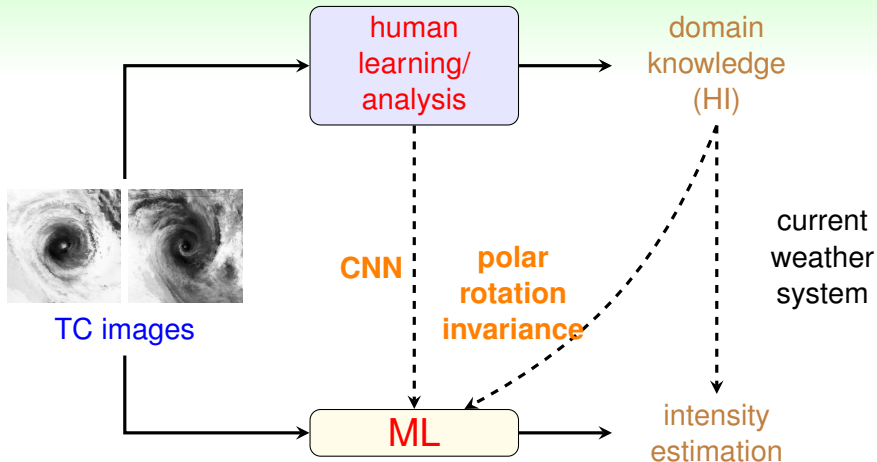
Can ML do the same/better?

- lack of **ML-ready datasets**
- lack of **model that properly utilizes domain knowledge**

issues addressed in our latest works

(Chen et al., KDD '18; Chen et al., Weather & Forecasting '19)

Recall: Flow behind Our Proposed Model



is proposed **CNN-TC** better than current weather system?

Results


RMS Error

ADT	11.75
AMSU	14.40
SATCON	9.66
CNN-TC	9.03

CNN-TC much better than current weather system (SATCON)

why are people **not**
using this **cool ML model?** 😊

Lessons Learned from Research on Tropical Cyclone Intensity Estimation



- 1 again, **cross-domain collaboration** important
e.g. even from ‘organizing data’ to be ML-ready
- 2 not easy to claim **production ready**
—can ML be used for ‘**unseenly-strong** TC’?
- 3 good AI system requires **both human and machine learning**
—still an ‘art’ to blend the two

Outline

ML for (Modern) AI

ML Research for Modern AI

ML for AI in Reality

Frequently Asked Questions of ML for AI (1/4)

What is the best AI project for
(my precious big) data?

My Polite Answer

good start already 😊, any more thoughts that you have in mind?

My Honest Answer

I don't know.

or a slightly longer answer:

if you don't know, I don't know.

A Similar Scenario

What is the best AI project for
(my precious big) data?
how to find a research topic for my thesis?

My Polite Answer

good start already 😊, any more thoughts that you have in mind?

My Honest Answer

I don't know.

or a slightly longer answer:
I don't know, but perhaps you can **start** by
thinking about **motivation** and **feasibility**.

Finding AI Projects \approx Finding Research Topics

- **motivation**: what are you interested in?
- **feasibility**: what can or cannot be done?

motivation

- something publishable?
oh, possibly **just for people in academia** 😊
- something that **improves xyz performance**
- something that inspires deeper study

—helps **generate** questions

feasibility

- **modeling**
- **computational**
- budget
- timeline
- ...

—helps **filter** questions

tip: important for **first AI project** to be
of high success possibility

Frequently Asked Questions of ML for AI (2/4)

Should I use ML (or my precious Deep Learning) for my AI project?

My Polite Answer

let's understand more about the constraints of your project, shall we 😊?

My Honest Answer

I don't know.

or a slightly longer answer:
if you don't know, I don't know.

Necessary Conditions for Using ML


machine learning: improving some **AI goal**
with experience **accumulated** from **data**



- 1 exists **some** “underlying pattern” to be learned
—so “AI goal” possible
- 2 but **no** programmable (easy) **definition**
—so “ML” is needed
- 3 somehow there is **enough data** about the pattern
—so ML has some “inputs” to learn from

necessary, but **not sufficient**, for using ML

Human Learning versus Machine Learning



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

Frequently Asked Questions of ML for AI (3/4)

What is the best machine learning model for
(my precious big) data and AI?

My Polite Answer

the best model is
data-dependent, let's **chat**
about your data first

My Honest Answer

I don't know.

or a slightly longer answer:
I don't know about **best**, but perhaps you can
start by thinking about **simple models**.

Sophisticated Model for AI

What is the best machine learning model for
(my precious big) data and AI?

What is the **most sophisticated** machine
learning model for (my precious big) data
and AI?

- myth: my AI works best with **most sophisticated** model
- sophisticated model:
 - time-consuming to **train** and **predict**
 - difficult to **tune** or **modify**
 - hard to “simplify” nor “analyze”

sophisticated model shouldn't be **first choice**

Simple First

What is the **first** machine learning model for
(my precious big) data and AI?

Taught in ML Foundations on NTU@Coursera

simple model first:

- efficient to **train** and **predict**
- easy to **tune** or **modify**
- somewhat “**analyzable**”
- little **risk**

My KISS Principle:
Keep It Simple, ~~Stupid~~ Safe

Frequently Asked Questions of ML for AI (4/4)

How to Get my AI Project Started?

Old Me


I don't know. 😊

New Me

I know one key factor!

let's see what the key factor is

Todos in AI Project



data

- gathering
- cleaning
- storing
- ...

techniques


- modeling
- computation
- non-ML tech.
- ...

usage

- **evaluation**
- deployment
- scalability
- ...

key first step: set up **evaluation criteria**

Evaluation Criteria Guide AI Project Planning



(free image by Manfred Steger from Pixabay)

suggest improvement opportunities



data

hint
preparation steps

techniques

assist
model/tech. choices

usage

define
acceptance goals

tip: always start with
reasonable & measurable criteria
to describe prioritized **AI goal**

Summary

- ML for (Modern) AI:
tools + human knowledge \Rightarrow **easy-to-use application**
- ML Research for Modern AI:
need to be **more open-minded**
—in methodology, in collaboration, in KPI
- ML for AI in Reality:
 - motivated/feasible project with **measurable criteria**
 - human and/or **simple** model first

Thank you! Questions?