Impact of Deep Learning on our Society

2/14, 2018

Hiroshi Maruyama

Preferred Networks

Agenda

1. What is Deep Learning
2. Implications to Software Development
3. Implications to us Individuals

Agenda

1. What is Deep Learning
2. Implications to Software Development
3. Implications to us Individuals

What is Deep Learning? – A (Stateless) Function

Low-dimensional for classification, very high-dimensional for generation

X

$Y = f(X)$

Y

Very high-dimensional, any combination of continuous and categorial variables

Example: Converting Celsius to Fahrenheit

Requirements

Input: C
Output: F
Where F is Fahrenheit
equivalent of C in Celsius

Model

$F = 1.8 * C + 32$

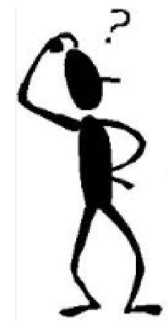

A Priori
Knowledge

Implementation

```
double c2f(double c) {
   return 1.8*c + 32.0;
}
```

Alternative Approach – Data-Driven, Inductive Programming

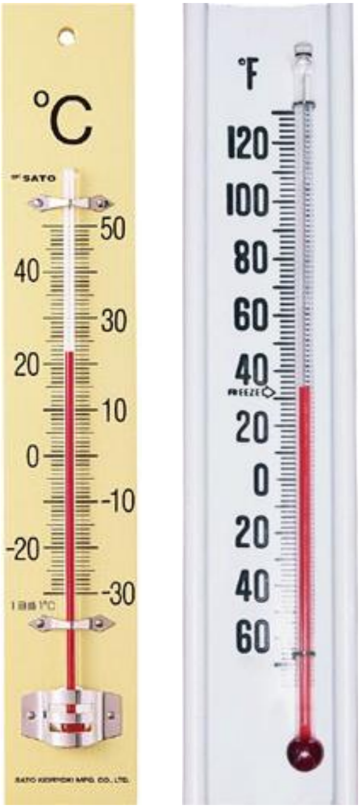

Find a model that represents this data set

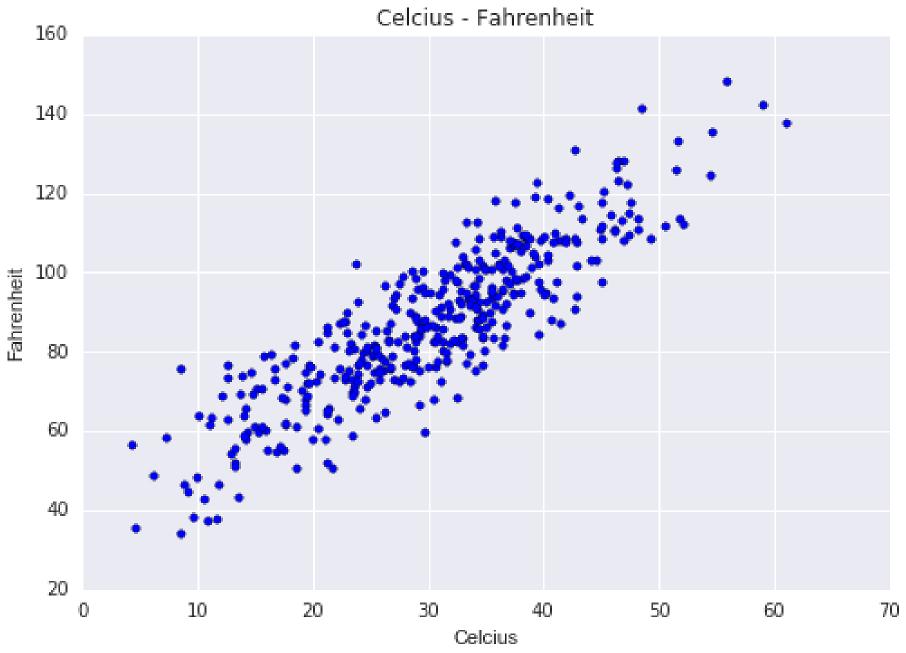

Machine Learning (aka Statistical Modeling) does this!

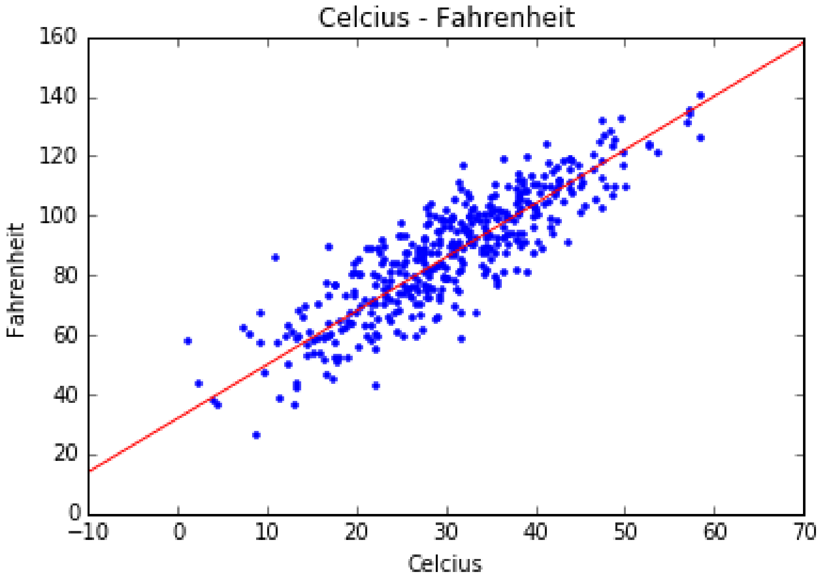

Estimated Model

F = 1.8 * C + 32 + e
e ~ N(0,10)

A Catch: the Model must be fixed in advance

What is the function that represents this data?

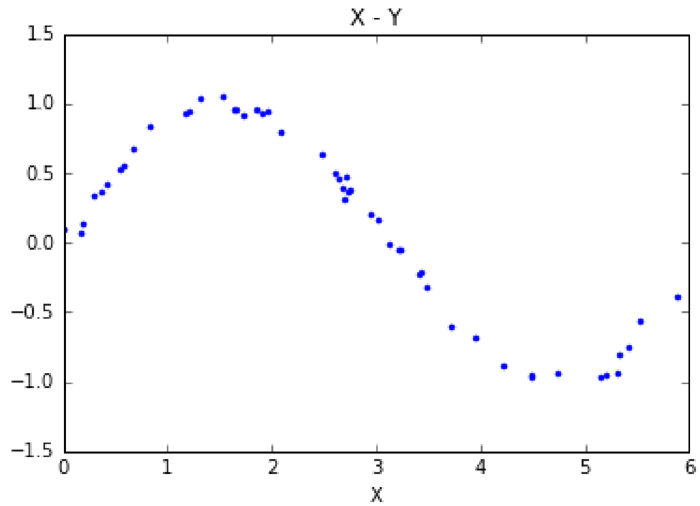

Too many parameters result in overfitting!

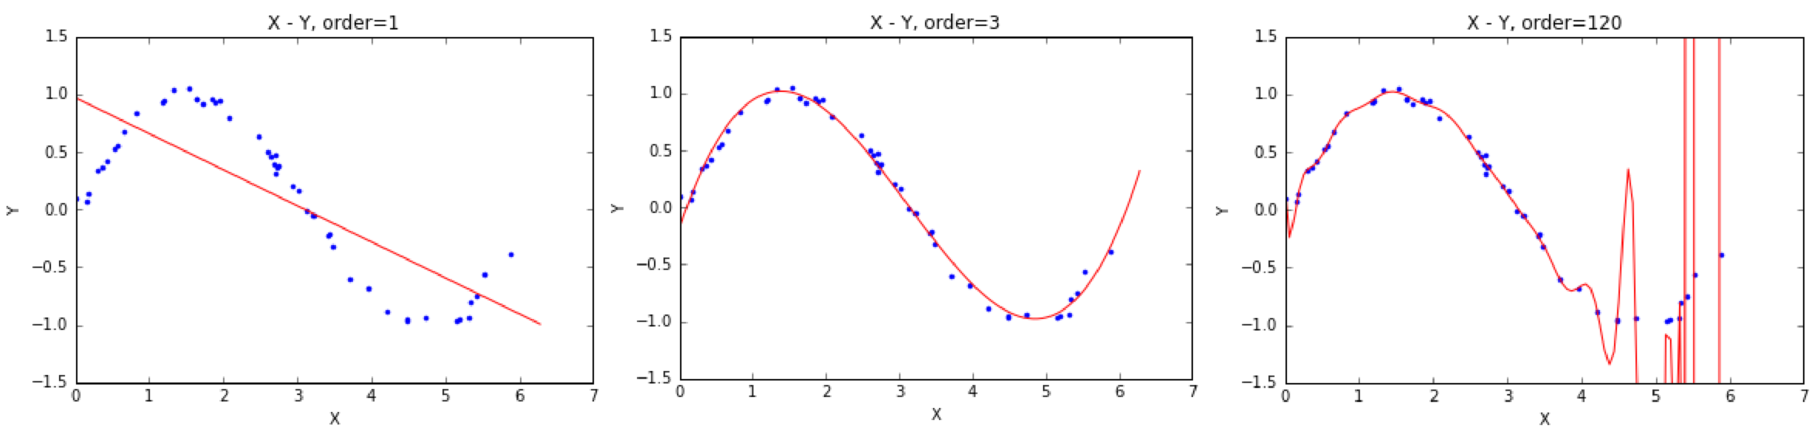

Choosing the right model is difficult

Deep Learning can approximate a function without too much overfitting (in many cases)

Mean-square error is back-propagated

Total 141 parameters

Ramp function as activation function

2-layers of 10-node hidden layers

Output

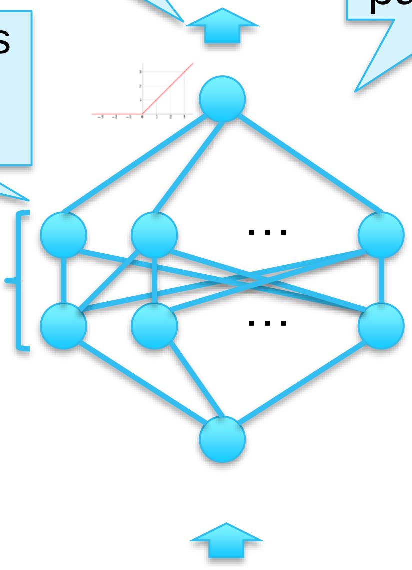

Input

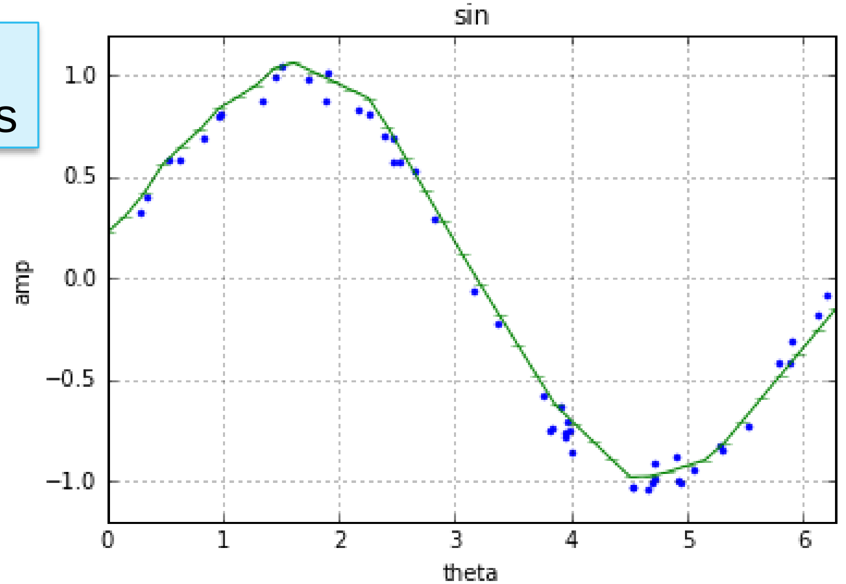

Good approximation without too much overfitting

Deep Learning as a Universal Computing Mechanism

Output

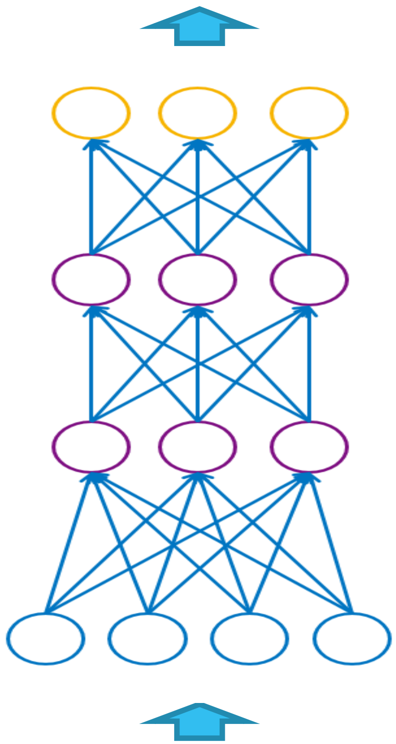

Input

- Very large number of parameters
- Can approximate ANY high-dimensional function*

➔ **Pseudo Turing Complete!**

* G. Cybenko. Approximations by superpositions of sigmoidal functions. Mathematics of Control, Signals, and Systems, 2(4):303–314, 1989.

How Deep Learning Works by Yann LeCun

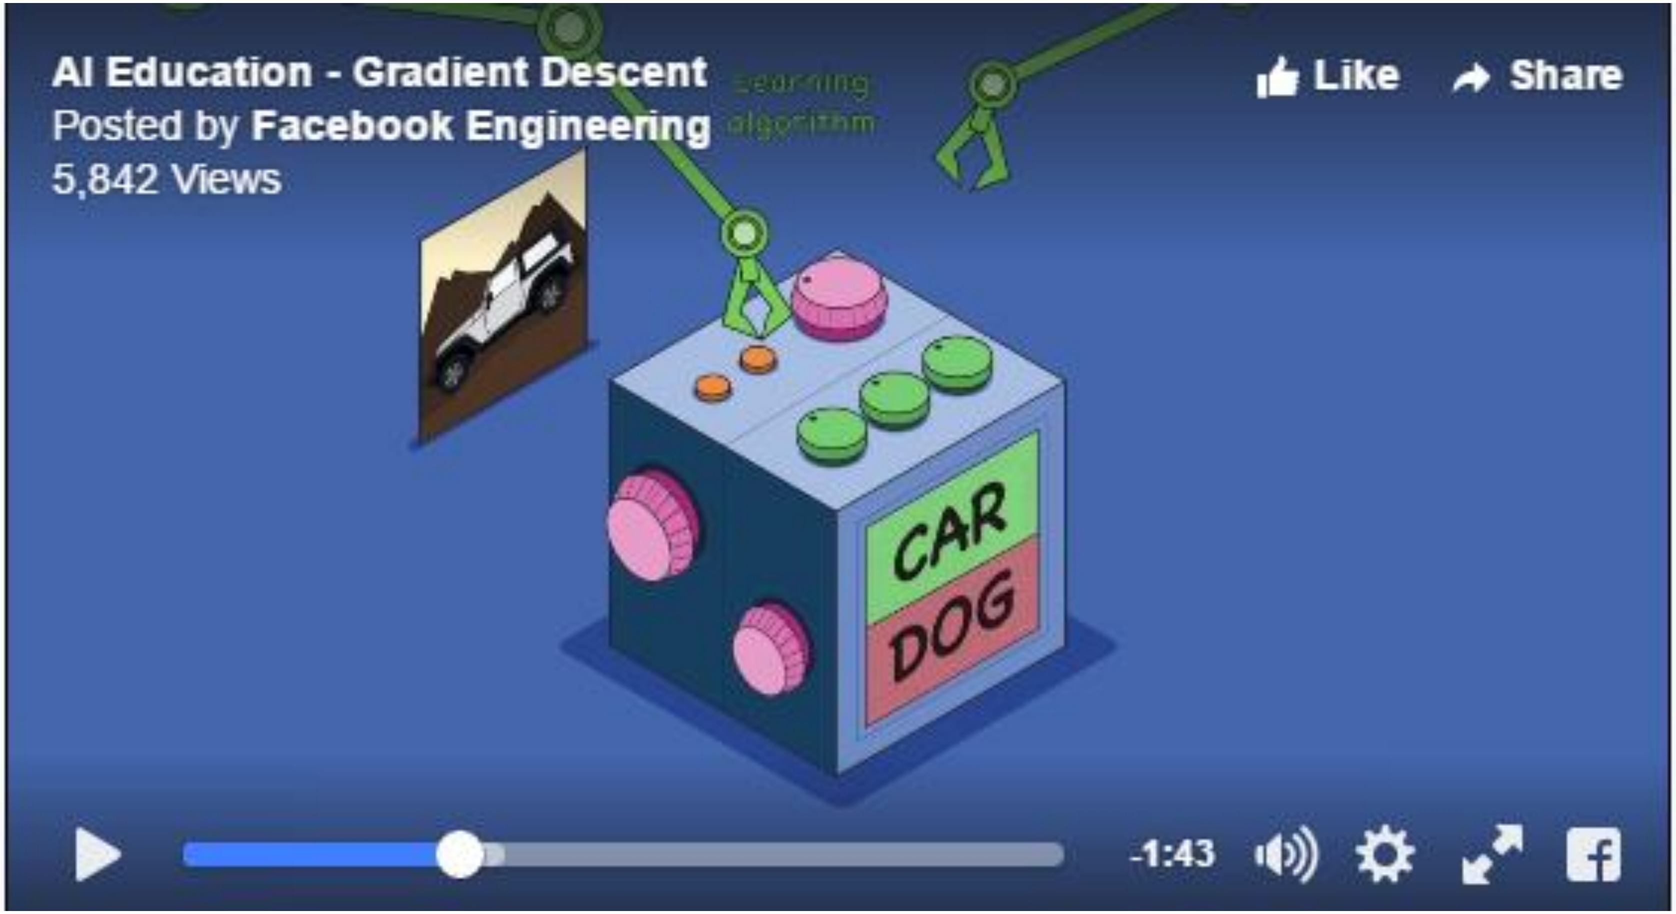

https://code.facebook.com/pages/1902086376686983

Image Segmentation for autonomous driving

https://www.youtube.com/watch?v=lGOjchGdVQs

Image Segmentation for Picking Robot

2nd Place in the Picking Task in Amazon Picking Challenge

Image + Natural Language for Controlling Robot

Auto Coloring Line Drawings

Reinforcement Learning for Autonomous Driving

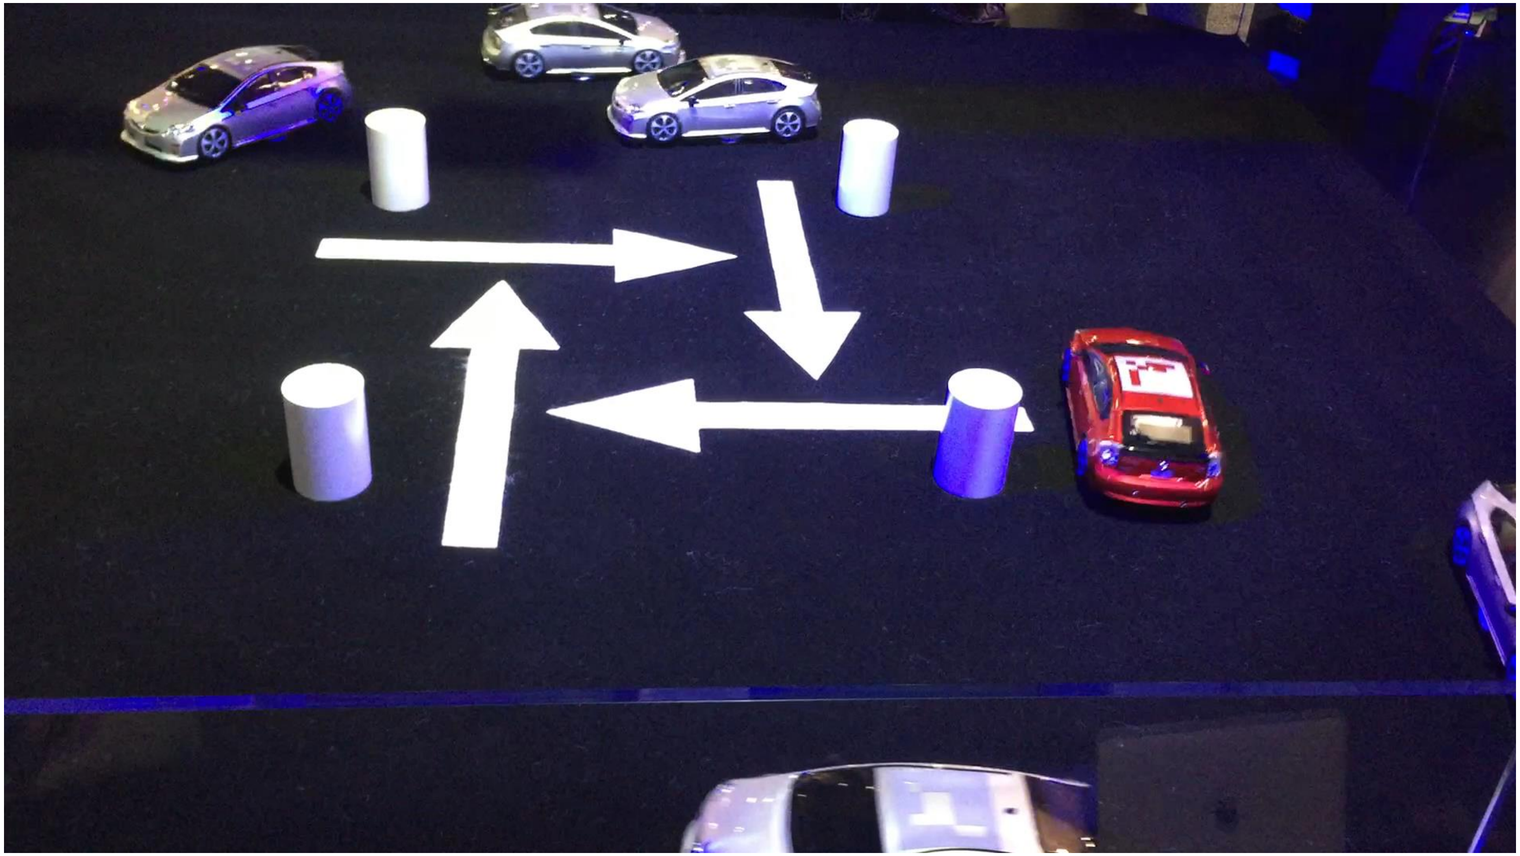

Consumer Electronics Show (CES) 2016

Development with Reinforcement Learning

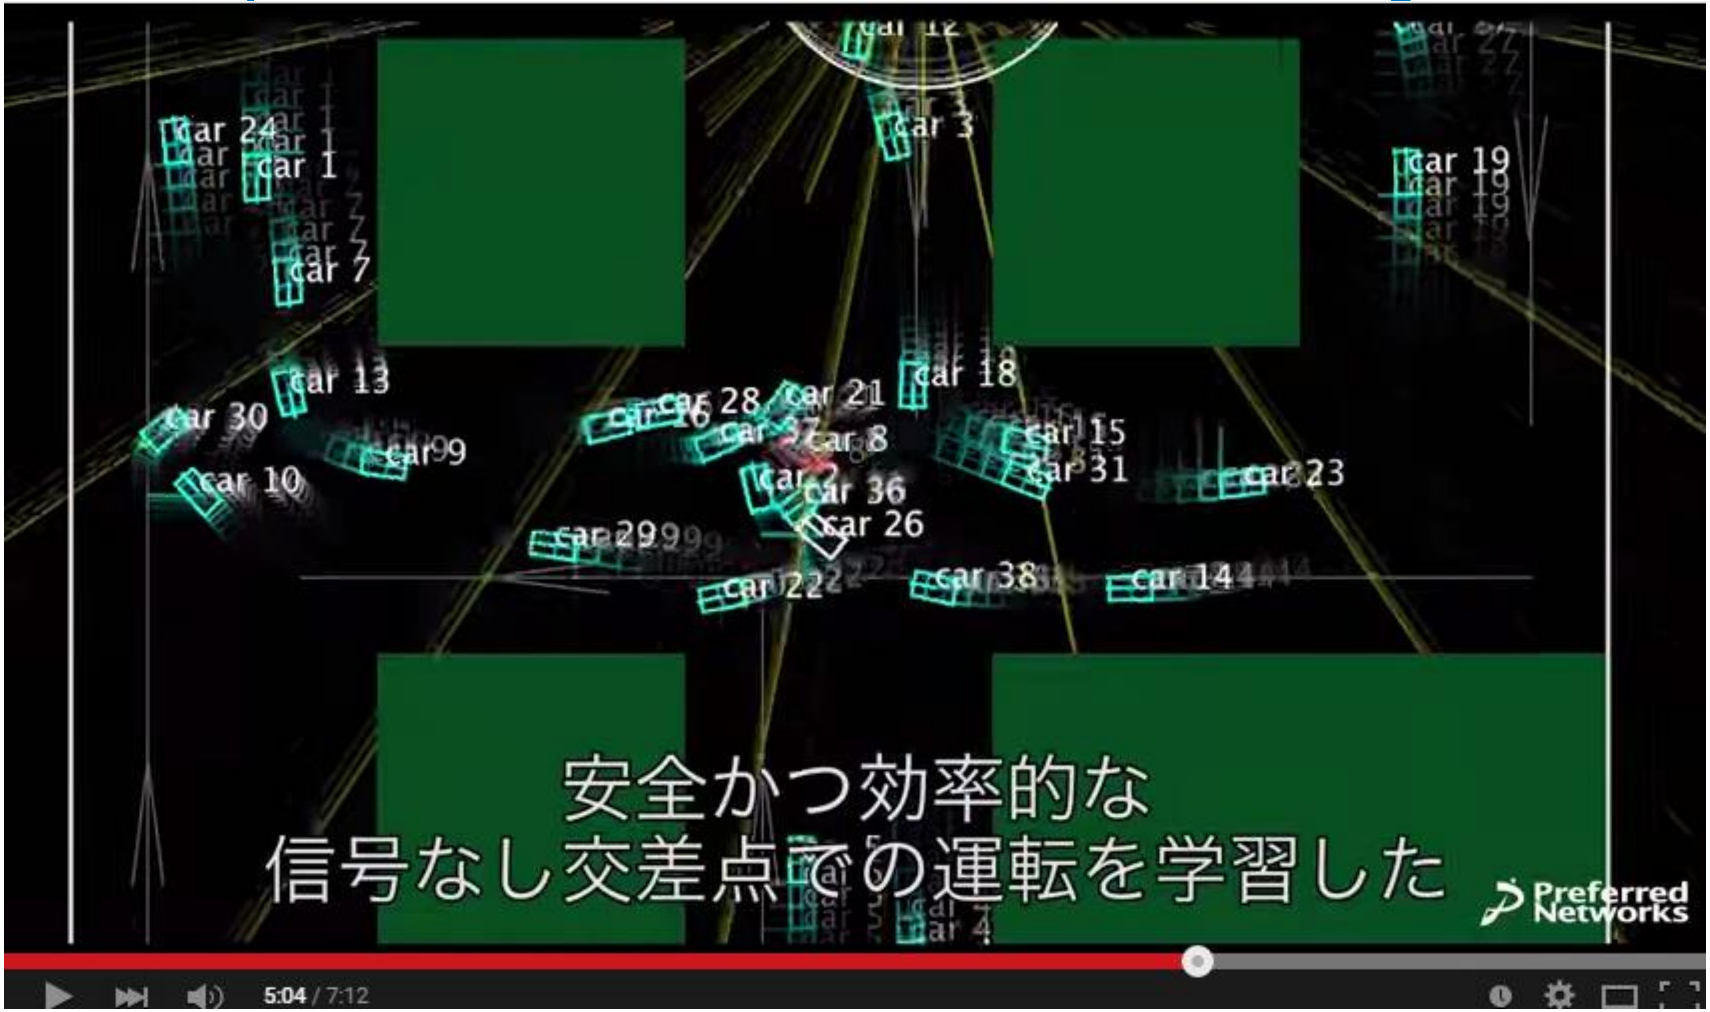

https://research.preferred.jp/2015/06/distributed-deep-reinforcement-learning/

Generative Model Captures Statistical Essence of Training Data

Deep Learning Requires lots of Computation

P:Peta
E:Exa
F:Flops

Image Video Rec.

10P (Image) ～ 10E (Video) Flops
100million images,

Life Science

100P ～ 1E Flops
10M SNPs per person. 100PF for 1million, 1EF for100 million.

Speech Rec.

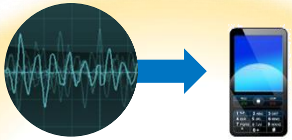

10P～ Flops
5K hours of 10K people of audio data
100K hours of synthetic audio data
for training [Baidu 2015]

Autonomous Driving

1E～100E Flops
1TB/day/autonomous cars
10~1000 cars, 100 days of data

Robotics/Drone

1E～100E Flops
1TB/car/year
Data from 1~100M cars

Machine generated data is much bigger than human generated data

These estimation is based on;
To finish training using 1GB within 1day require 1Tflops

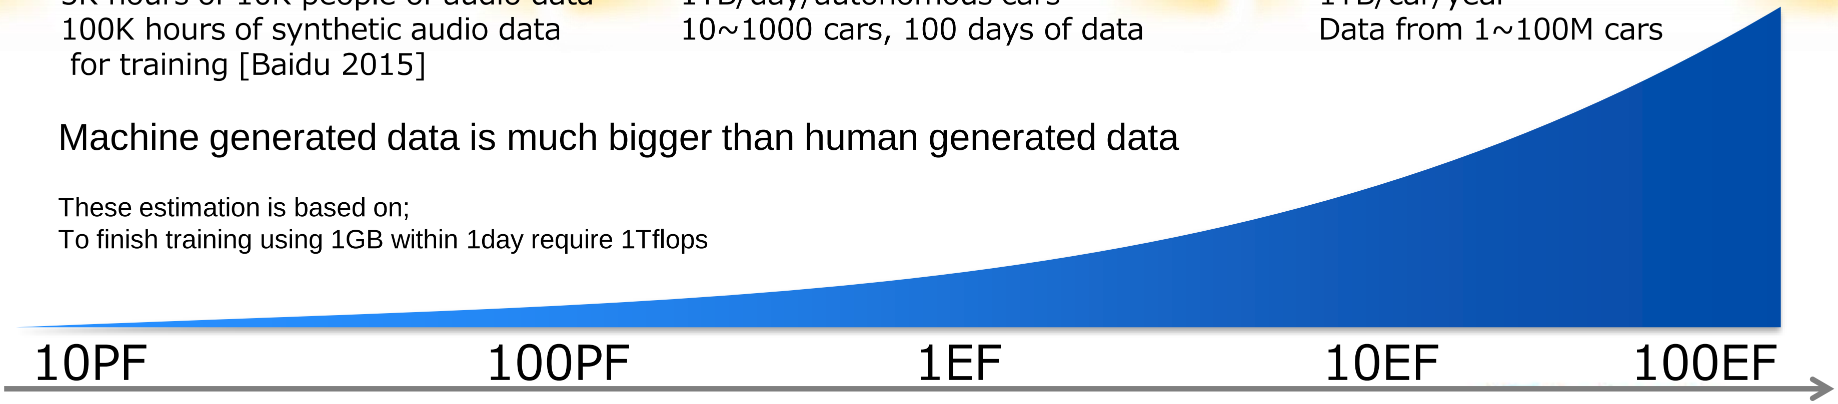

Preferred Networks

Fundamental Limitation of ML (1)

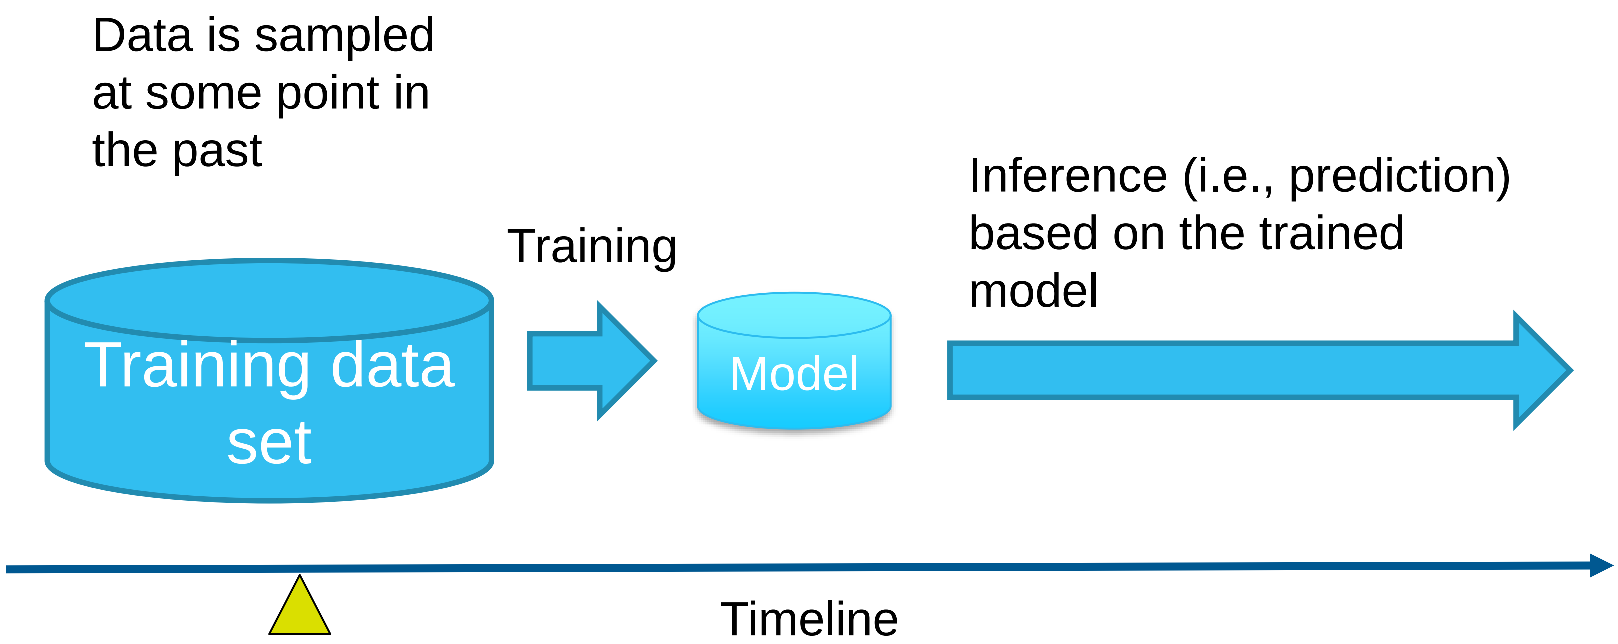

Statistical Machine Learning works only if the future is similar to the past

Fundamental Limitation of ML (2)

- Powerless on data in unseen regions

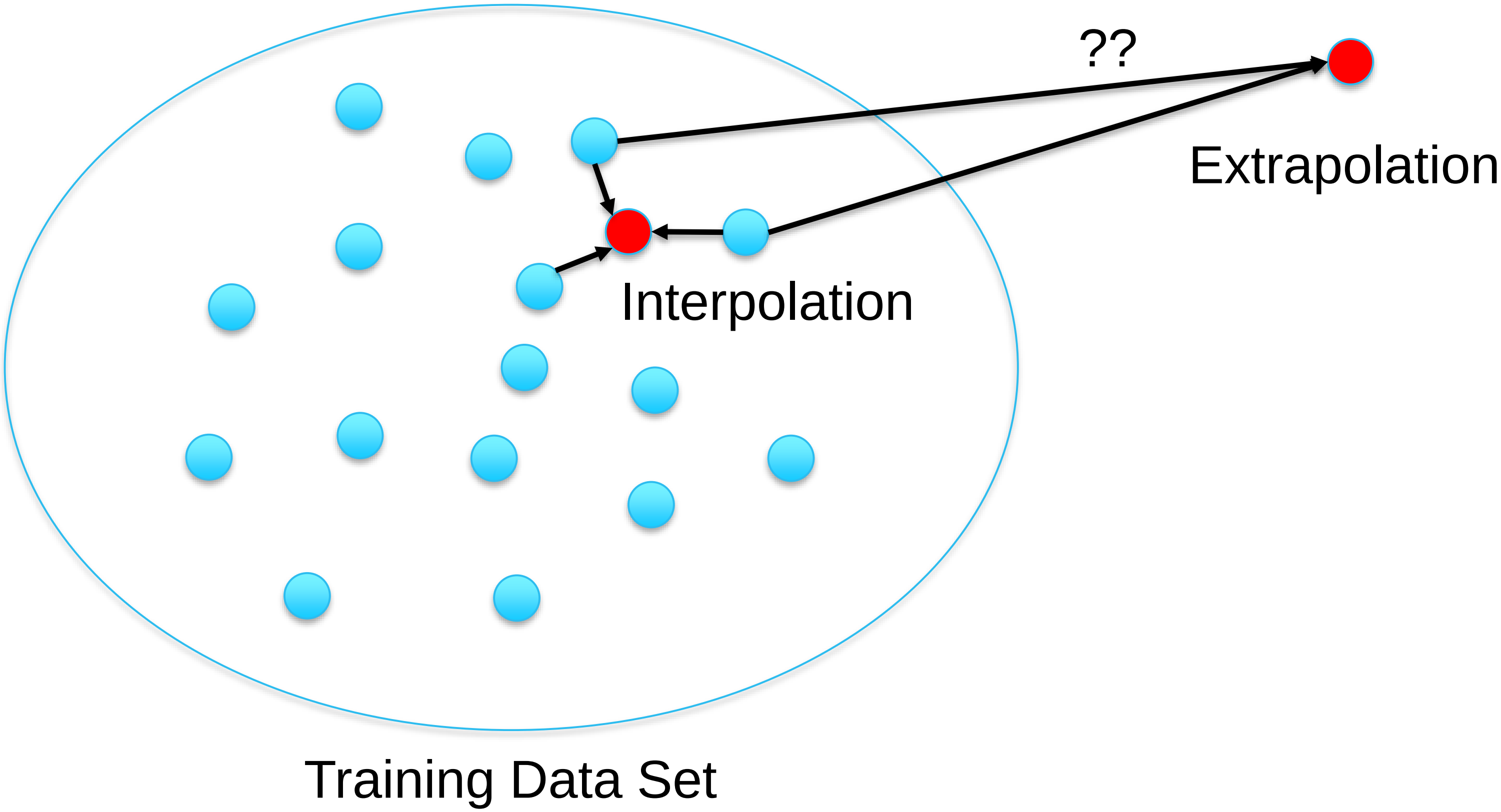

Fundamental Limitation of ML (3)

- Always works statistically

Random Sampling !!

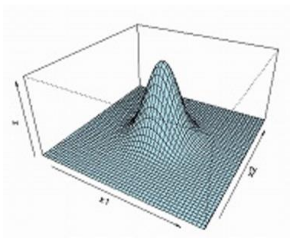

Original Distribution

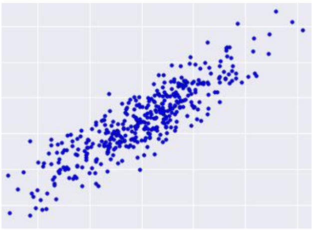

Training Data Set

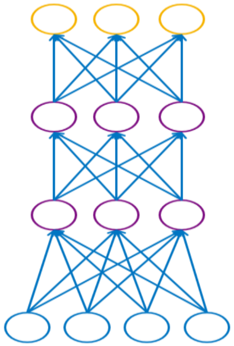

Trained Model

No guarantee of "100% correctness"

Agenda

1. What is Deep Learning
2. Implications to Software Development
3. Implications to us Individuals

Typical ML-Based System

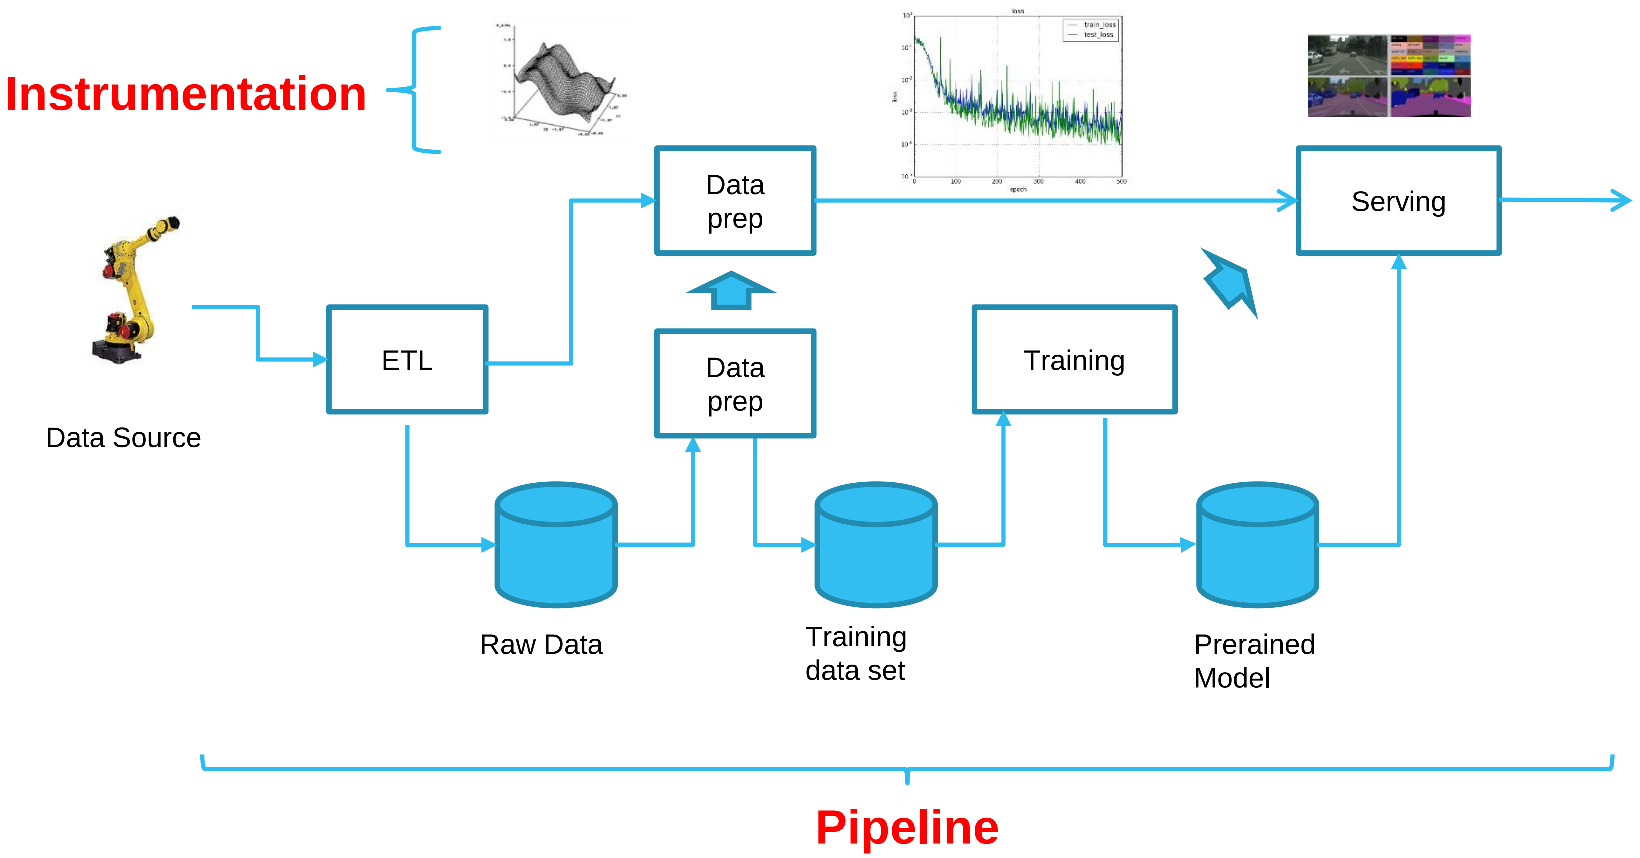

ML-Based System Lifecycle

Assessment → PoC → Pilot → Deployment → Maintenance

CP1
- Customer Readiness
- Clear Business KPI
- Fitness to DL
- Data Availability
- ：

CP2
- Appropriateness of Data (Volume + Quality)
- Appropriateness of Model (accuracy)

CP3
- Business KPI achievable?
- Business process readiness
- ：

CP4
- KPI review
- Trend of data drift
- ：

The Process is more exploratory

Conventional Deductive Development (Manufacturing-Inspired)

Requirements

Model-Driven Stepwise Refinement

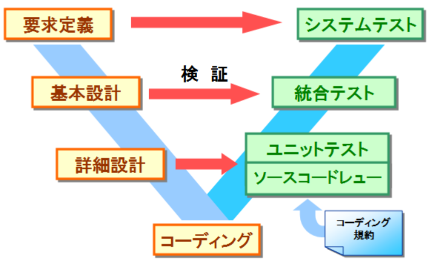

Input

System with modularized components (e.g., recognition, planning, feedback control, …)

Output

Inductive Development (=Machine Learning)

Training Data
(input/output pairs)

Machine Learning

Input

Monolithic end-to-end system (deep neural net)

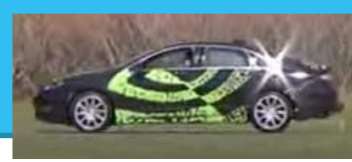

Output

Requirements = Training Data Sets

Autonomous Vehicle Demo at Toyota's Booth, Consumer Electronics Show 2016

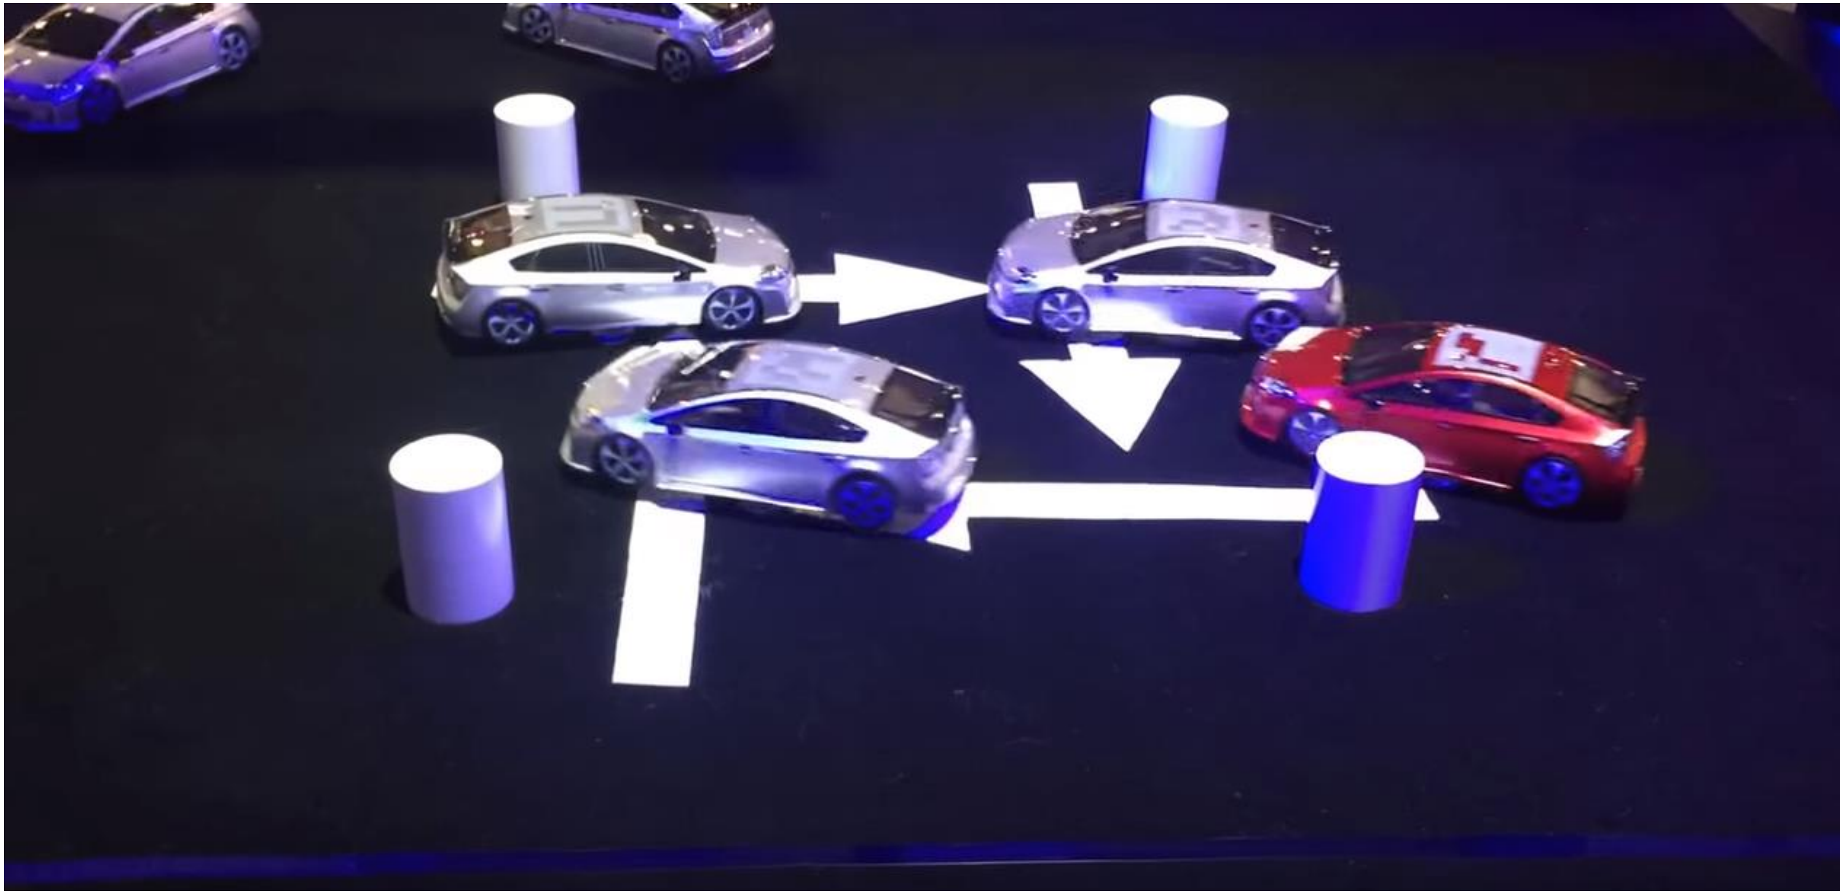

https://www.youtube.com/watch?v=7A9UwxvgcV0

Trial-and-Error Development (a.k.a. Reinforcement Learning)

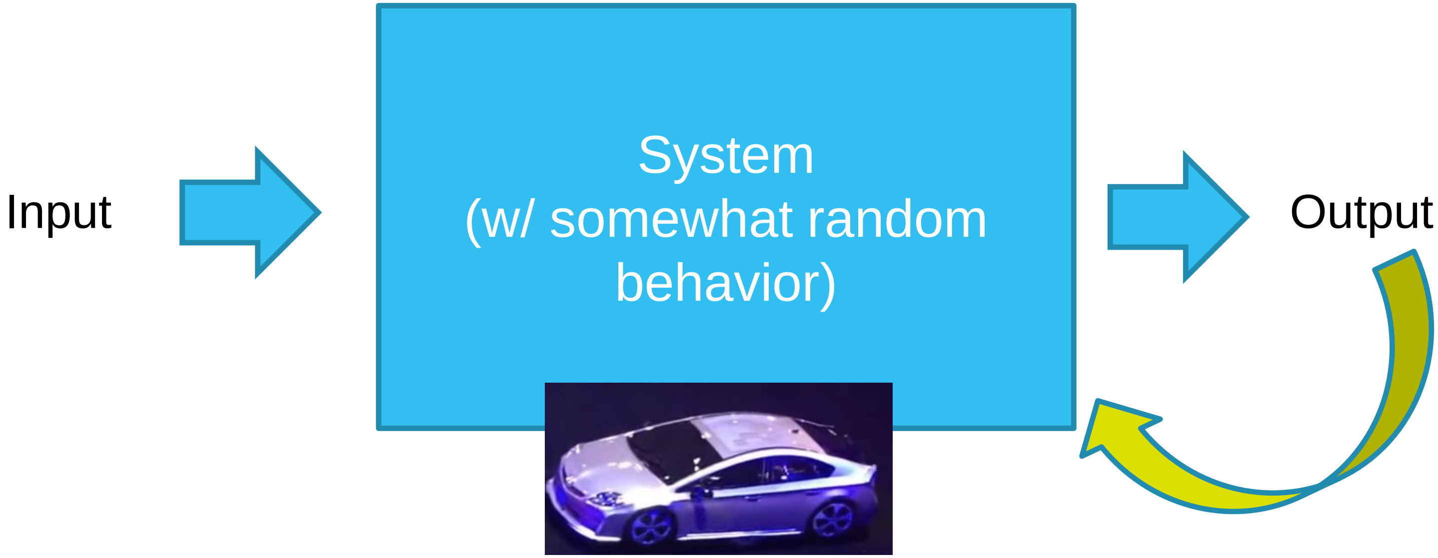

Requirements = After-the-fact Feedback

New Paradigm of System Development

- Manufacturing-Inspired
  - Huge cost of modification after development
  - Requirement Definition in advance

Stable and well-organized environments

- Machine Learning-Inspired
  - System modification cost is virtually zero
  - After-the-fact Requirement Definition

Dynamic and less-understood environments

Cf. Waterfall -> Agile -> DevOps -> …

Software Crisis in 1960's

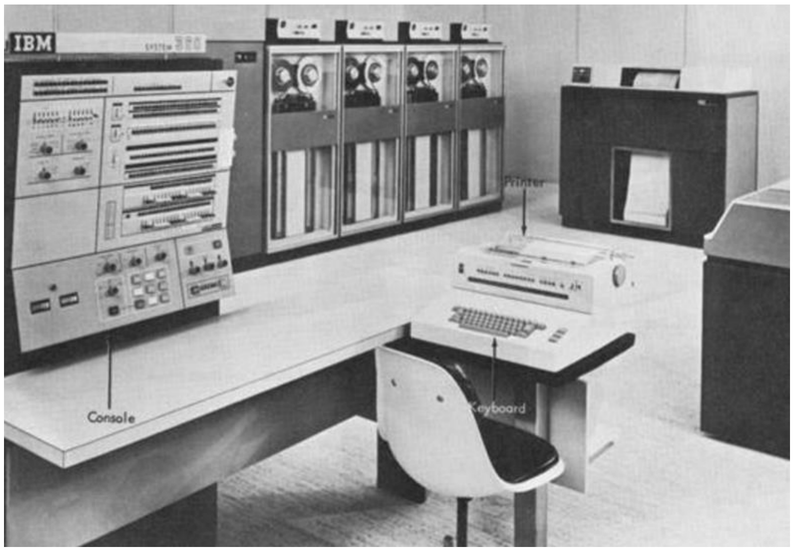


```
                                   TABLE 1
                      EXTENDED FLOATING POINT ADD ROUTINE
                         Copyright IBM Corporation 1984

* Special extended precision addition of FPRs 0 and 4
*
IHOXAR04 CSECT
RHOA     EQU      0
RLOA     EQU      2
RHOB     EQU      4
RLOB     EQU      6
         USING    *,15         Set base register
START    B        BEGIN        Branch around identifier
         DC       AL1(8)
         DC       CL8'IHOXAR04'
*
* Before touching GPR's, save them
*
BEGIN    STM      14,12,12(13) Save caller's registers
         LR       2,13
         ST       2,SAVE+4
         LA       13,SAVE
         ST       13,8(2)
         LTDR     RHOB,RHOB    Check high order part of B
         BZ       RETURN       B is zero so done
*
BNOTO    LTDR     RHOA,RHOA    Check high order part of A
         BNZ      ANOTO
         LDR      RHOA,RHOB    A is zero so set A to B
         LDR      RLOA,RLOB
         B        RETURN       Finished
*
* Neither A nor B is zero -- save A and B and continue
*
ANOTO    STD      RHOA,HOA     Save high order part of A
         STD      RLOA,LOA     Save low order part of A
         STD      RHOB,HOB     Save high order part of B
         STD      RLOB,LOB     Save low order part of B
*
* Check for result in range so standard add can be done
*
         L        2,HOA        Load high order A
         X        2,LOA        Exclusive or to check signs match
         L        3,HOB        Load high order B
         X        3,LOB        Exclusive or to check signs match
         LTR      2,2          Result > = 0 they match -- A in range
         BM       NOSTD        No match--out of range
         LTR      3,3          Result > = 0 they match -- B in range
         BM       NOSTD        No match--out of range
```

➔ Dawn of Software Engineering

Similarly, we should start "Machine Learning Engineering" now!

- Requirements engineering
  - How to translate requirements into training data sets?
  - How to communicate pros & cons of ML with stakeholders?
- Testing
  - How to do "Regression Test" in ML?
  - What can we learn from the practice of Test-Driven Development (TDD)?
- Tools
  - What should ML IDE look like?
  - How should we manage ML work products, especially training data sets and pretrained models?
  - How to avoid overfitting caused by overly automated tools?
- Operation
  - How to detect "concept drift" in operations?
- :

Agenda

1. What is Deep Learning
2. Implications to Software Development
3. Implications to us Individuals

Please make clear distinction between generalized AI and specialized AI

"There's a distinction, which is probably familiar to a lot of your readers, between **generalized AI** and **specialized AI**. "

https://www.wired.com/2016/10/president-obama-mit-joi-ito-interview/

(Specialized) A. I. is not a technology per se – It is an activity to extend horizon of computer science

1st Wave of A. I. (1956-1974)
- Symbol Processing (LISP)
- Means-End Analysis
- Language Parsing

→

- Garbage Collection
- Search Algorithms
- Formal Language Theory
- :

2nd Wave of A. I. (1980-1987)
- Knowledge Representation
- Expert System
- Ontology

→

- Object-Oriented Language
- Modeling
- Semantic Web
- :

3rd Wave of A. I. (2008-)
- Statistical Machine Learning
- Deep Learning

→

Inductive Programming

What is Intelligence? – Two Types of Thinking

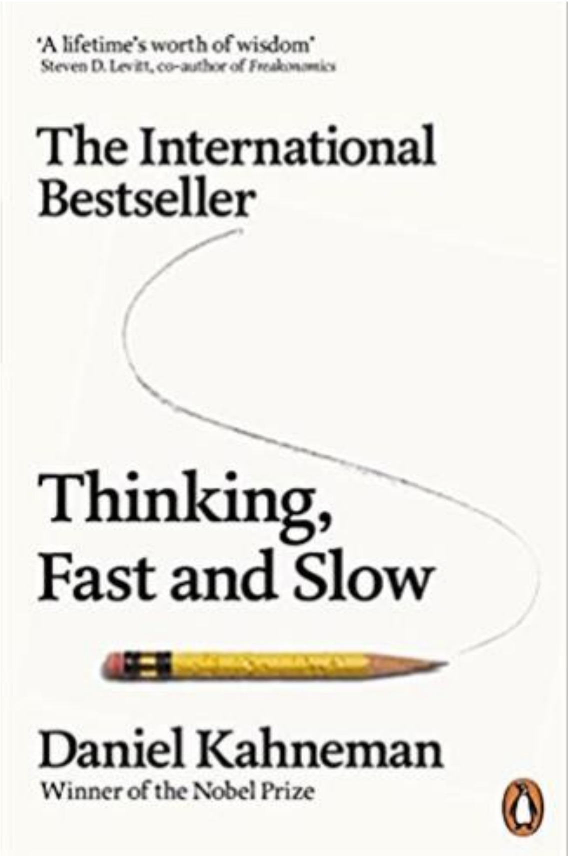

ISBN-13: 978-0141033570

- System 1 Thinking
  - Fast
  - Automatic
  - Intuitive
  - E.g., 2 x 2 = ?
- System 2 Thinking
  - Slow
  - Consumes attention
  - Logical
  - E.g., 27 x 31 = ?

Preferred Networks

How AI Research Evolved

1st and 2nd waves of AI

- System 2 Thinking
  - Logical
  - Reason
  - Symbolic
  - Deductive

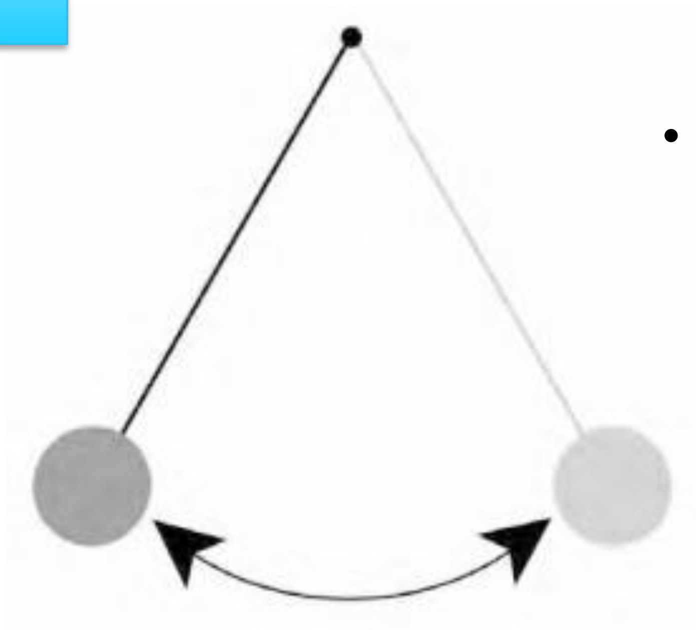

3rd wave of AI

- System 1 Thinking
  - Perception
  - Intuition
  - Distributed
  - Inductive

Which is more "human-like" intelligence?

I asked a carpenter’s son, “what would you like to be in the future?”

He said, “I want to be a great tofu maker like my father.”

Big Data (and ML) can manipulate you via your cognitive bias

“Echo-chamber effect” – everybody you see has the same opinion with you

Remember, our society is based on the fundamental value assumption that people are rational

- Democracy, capitalism, science, etc. are all strongly influenced by "enlightenment (啓蒙思想)"

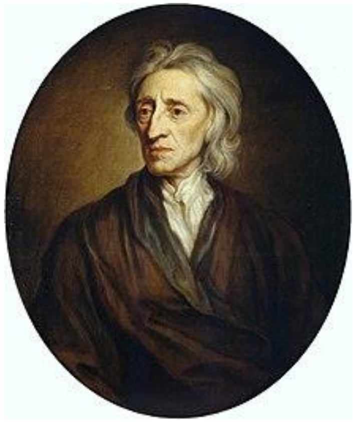

John Locke

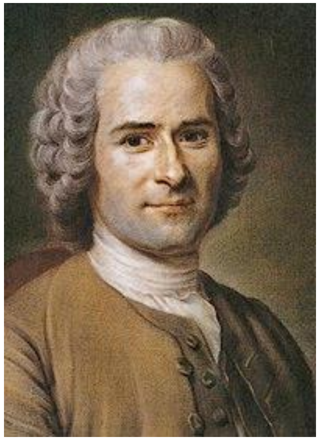

Jean-Jacques Rousseau

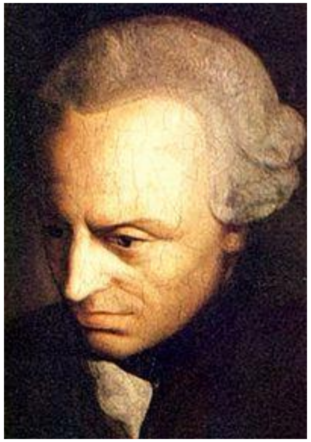

Immanuel Kant

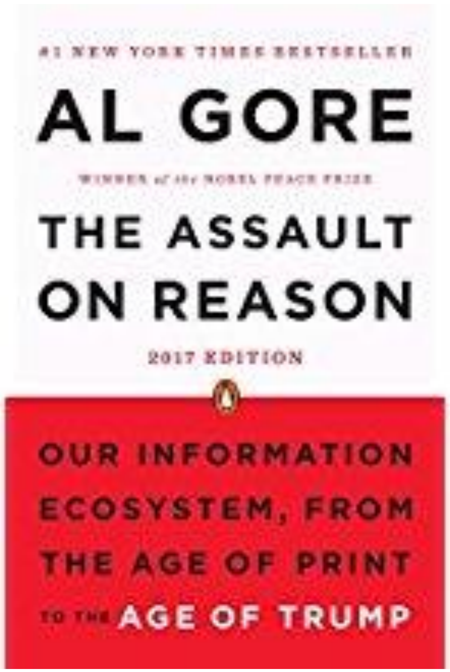

ISBN-13: 978-0143113621

ISBN-13: 978-1443422536

Thank You