UNIVERSITY OF TORONTO | Engineering

Machine Learning in SCM

Opportunities, Challenges and Possible Solutions

University of Toronto
Chi-Guhn Lee

1

Agenda

- OR to ML
- AI, ML and SCM
- Recent Projects
- Challenges and Attempts to Overcome
- Closing

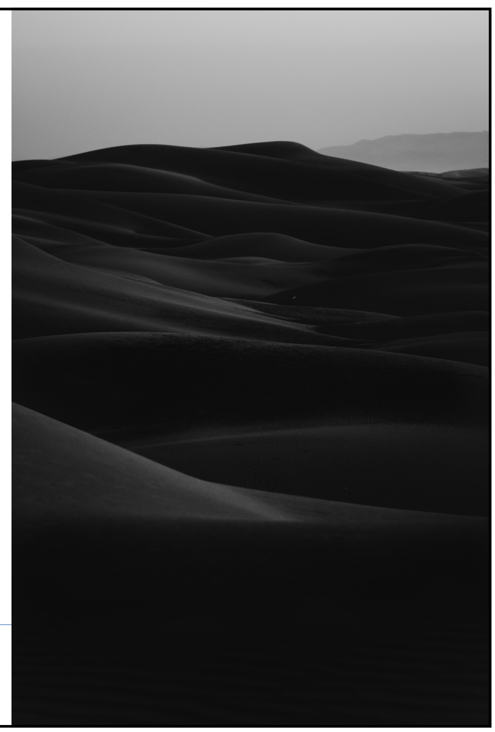

UNIVERSITY OF TORONTO | Engineering

2

UNIVERSITY OF TORONTO | Engineering

Where I turned from OR to ML

3

Where I turned around in 2016

TRIDEL®

HEC PARIS

- Line of vertically differentiable products
- Multi-attribute
  - Size
  - Floor
  - View
  - # of rooms
- Financial constraints
  - Sales targets

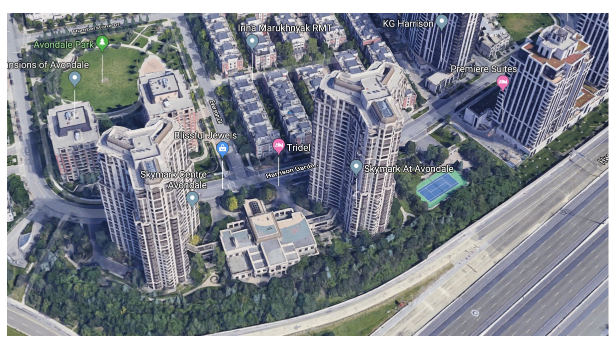

UNIVERSITY OF TORONTO | Engineering

4

Curse-of-dimensionality

- A few days of computation for a property of 10 unites

Time	States
Day 1-26	(1,1,1,1,1,1,1,1,1,1)
Day 27-32	(1,1,1,1,1,1,1,0,1,1)
Day 33	(1,1,1,0,1,1,1,0,1,1)
Day 34-52	(1,1,1,0,0,1,1,0,1,1)
Day 53-201	(0,1,1,0,0,1,1,0,1,1)
Day 202-274	(0,1,1,0,0,1,1,0,1,0)
Day 275-293	(0,0,1,0,0,1,1,0,1,0)
Day 294-360	(0,0,1,0,0,1,0,0,1,0)

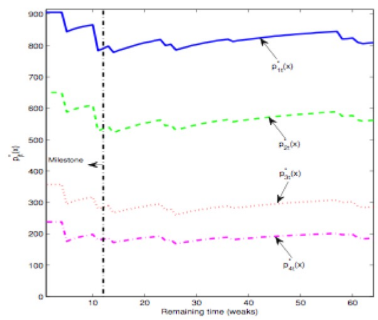

UNIVERSITY OF TORONTO | Engineering

6

7

New Era of AI (esp. among Koreans)

UNIVERSITY OF TORONTO | Engineering

8

History of AI

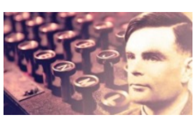

The Dawn of AI

The Boltzmann Machine

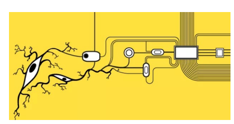

Deep Learning

AlphaGo

1950

1983

2006

2016

1956

1997

2011

The Dartmouth Conference

IBM's Deep Blue

IBM's Watson

UNIVERSITY OF TORONTO | Engineering

10

AI, ML and DL

- **AI** is to have a computer to mimic human behavior
- **ML** is to enable computers to learn from data
- **DL** is a type of ML algorithms inspired by human brain

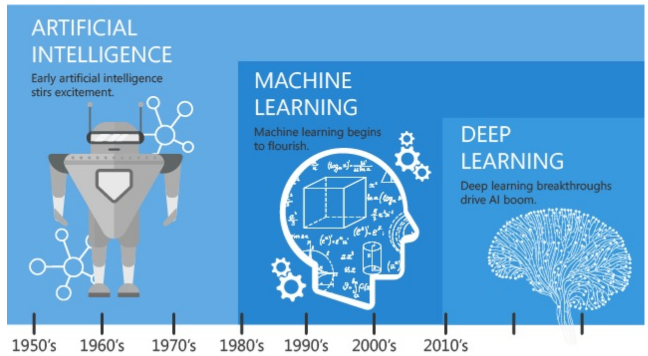

Super-human Performance

- ImageNet: Large Scale Visual Recognition Challenge
  - 1,000 classes and 1,431,167 images

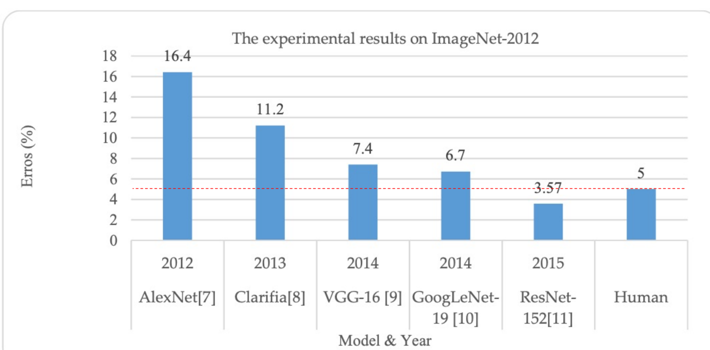

Super-human Performance

- Reinforcement Learning (Mnih, et al., 2014)

Game	Random Play	Best Linear Learner	Contingency (SARSA)	Human	DQN (± std)	Normalized DQN (% Human)
Alien	227.8	939.2	103.2	6875	3069 (±1093)	42.7%
Amidar	5.8	103.4	183.6	1676	739.5 (±3024)	43.9%
Assault	222.4	628	537	1496	3359(±775)	246.2%
Asterix	210	987.3	1332	8503	6012 (±1744)	70.0%
Asteroids	719.1	907.3	89	13157	1629 (±542)	7.3%
Atlantis	12850	62687	852.9	29028	85641(±17600)	449.9%
Bank Heist	14.2	190.8	67.4	734.4	429.7 (±650)	57.7%
Battle Zone	2360	15820	16.2	37800	26300 (±7725)	67.6%
Beam Rider	363.9	929.4	1743	5775	6846 (±1619)	119.8%
Bowling	23.1	43.9	36.4	154.8	42.4 (±88)	14.7%
Boxing	0.1	44	9.8	4.3	71.8 (±8.4)	1707.9%
Breakout	1.7	5.2	6.1	31.8	401.2 (±26.9)	1327.2%
Centipede	2091	8803	4647	11963	8309(±5237)	63.0%
Chopper Command	811	1582	16.9	9882	6687 (±2916)	64.8%
Crazy Climber	10781	23411	149.8	35411	114103 (±22797)	419.5%
Demon Attack	152.1	520.5	0	3401	9711 (±2406)	294.2%
Double Dunk	-18.6	-13.1	-16	-15.5	-18.1 (±2.6)	17.1%
Enduro	0	129.1	159.4	309.6	301.8 (±24.6)	97.5%
Fishing Derby	-91.7	-89.5	-85.1	5.5	-0.8 (±19.0)	93.5%
Freeway	0	19.1	19.7	29.6	30.3 (±0.7)	102.4%
Frostbite	65.2	216.9	180.9	4335	328.3 (±250.5)	6.2%
Gopher	257.6	1288	2368	2321	8520 (±3279)	400.4%
Gravitar	173	387.7	429	2672	306.7 (±223.9)	5.3%
H.E.R.O.	1027	6459	7295	25763	19950 (±158)	76.5%
Ice Hockey	-11.2	-9.5	-3.2	0.9	-1.6 (±2.5)	79.3%
James Bond	29	202.8	354.1	406.7	576.7 (±175.5)	145.0%
Kangaroo	52	1622	8.8	3035	6740 (±2959)	224.2%
Krull	1598	3372	3341	2395	3805 (±1033)	277.0%
Kung-Fu Master	258.5	19544	29151	22736	23270 (±5955)	102.4%
Montezuma's Revenge	0	10.7	259	4367	0 (±0)	0.0%
Ms. Pacman	307.3	1692	1227	15693	2311(±525)	13.0%
Name This Game	2292	2500	2247	4076	7257 (±547)	278.3%
Pong	-20.7	-19	-17.4	9.3	18.9 (±1.3)	132.0%
Private Eye	24.9	684.3	86	69571	1788 (±5473)	2.5%
Q*Bert	163.9	613.5	960.3	13455	10596 (±3294)	78.5%
River Raid	1339	1904	2650	13513	8316 (±1049)	57.3%
Road Runner	11.5	67.7	89.1	7845	18257 (±4268)	232.9%
Robotank	2.2	28.7	12.4	11.9	51.6 (±4.7)	509.0%
Seaquest	68.4	664.8	675.5	20182	5286(±1310)	25.9%
Space Invaders	148	250.1	267.9	1652	1976 (±893)	121.5%
Star Gunner	664	1070	9.4	10250	57997 (±3152)	598.1%
Tennis	-23.8	-0.1	0	-8.9	-2.5 (±1.9)	143.2%
Time Pilot	3568	3741	24.9	5925	5947 (±1600)	100.9%
Tutankham	11.4	114.3	98.2	167.6	186.7 (±41.9)	112.2%
Up and Down	533.4	3533	2449	9082	8456 (±3162)	92.7%
Venture	0	66	0.6	1188	380.0 (±238.6)	32.0%
Video Pinball	16257	16871	19761	17298	42684 (±16287)	2539.4%
Wizard of Wor	563.5	1981	36.9	4757	3393 (±2019)	67.5%
Zaxxon	32.5	3365	21.4	9173	4977 (±1235)	54.1%

UNIVERSITY OF TORONTO | Engineering

14

SCM and ML

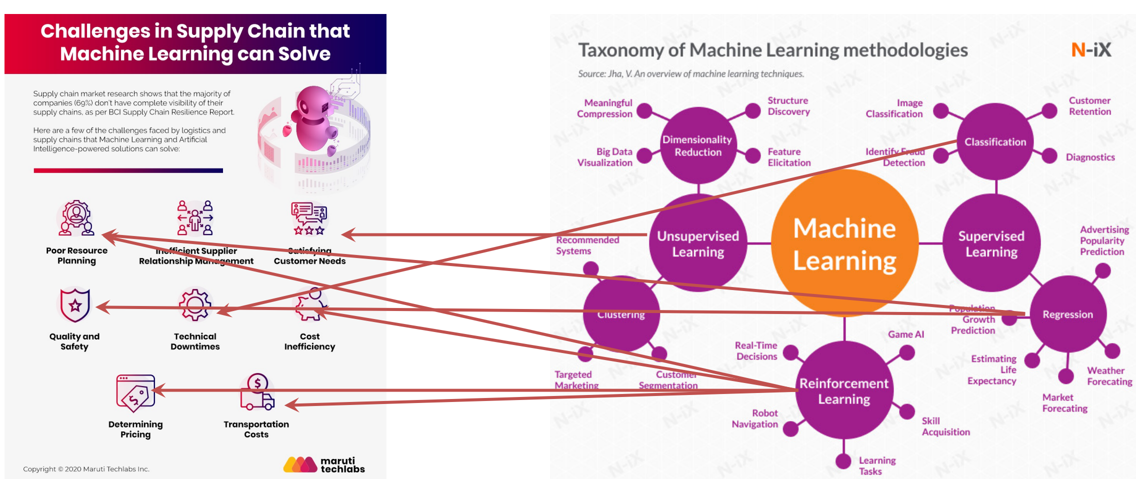

UNIVERSITY OF TORONTO | Engineering

15

UNIVERSITY OF TORONTO | Engineering

Recent Cases (in our group)

16

Anomaly Detection on Track

- Subway power-rail may overheat and cause defects
- IR camera is used to identify anomalies, but the video must be viewed manually to detect the anomalies
- Using TensorFlow Object Detection API & Microsoft Visual Object Tagging Tool to detect anomalies automatically

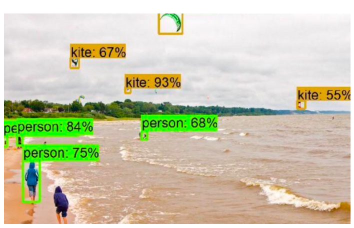

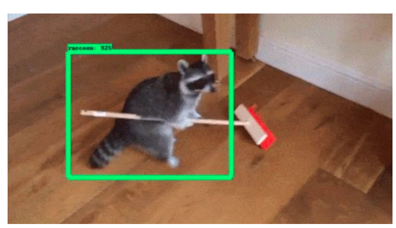

UNIVERSITY OF TORONTO | Engineering

17

Anomaly Detection on Track

- ANN using Tensorflow Object Detection API

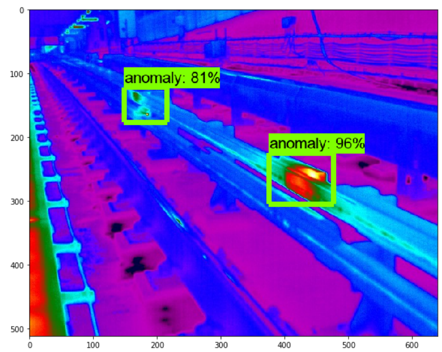

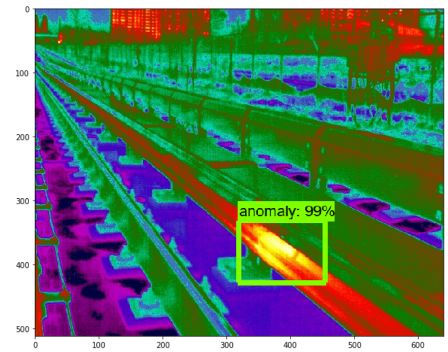

True Positives

UNIVERSITY OF TORONTO | Engineering

18

Underpayment Claims

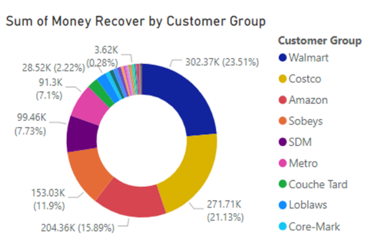

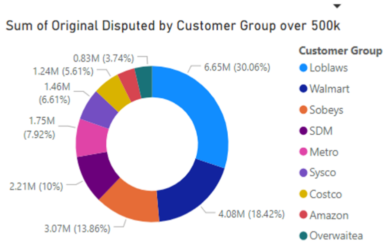

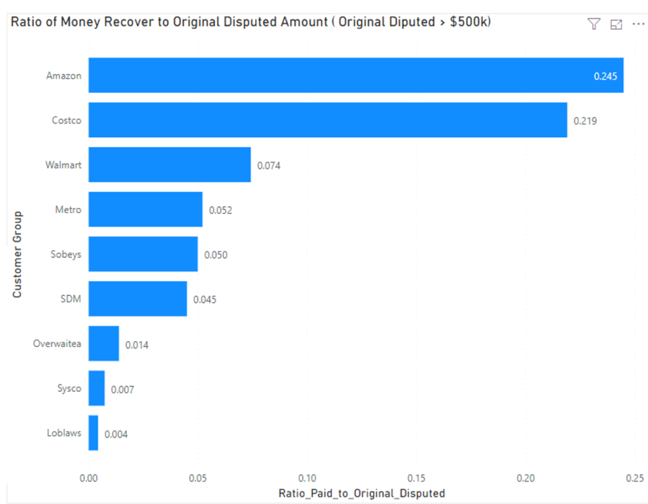

UNIVERSITY OF TORONTO | Engineering

19

Claims and Recovery

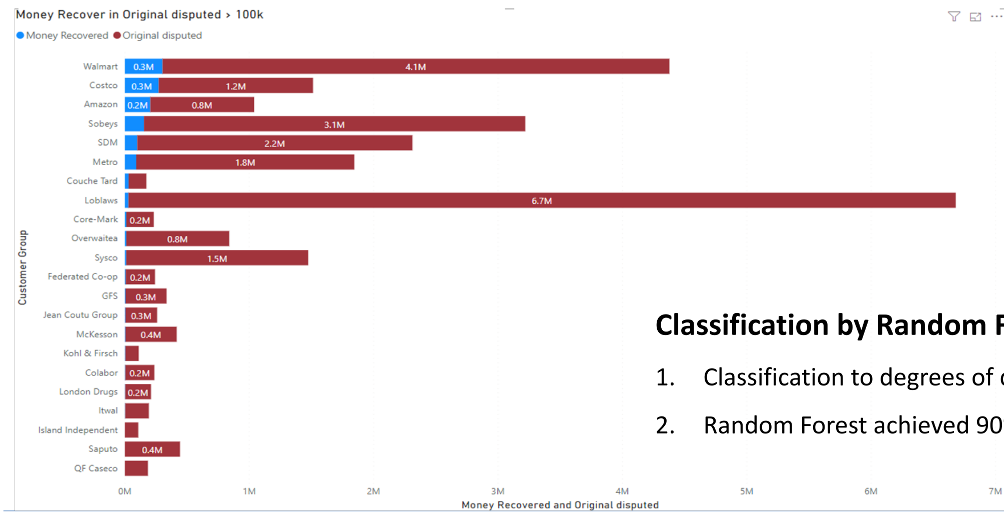

Classification by Random Forest

1. Classification to degrees of difficulty
2. Random Forest achieved 90% accuracy

UNIVERSITY OF TORONTO | Engineering

20

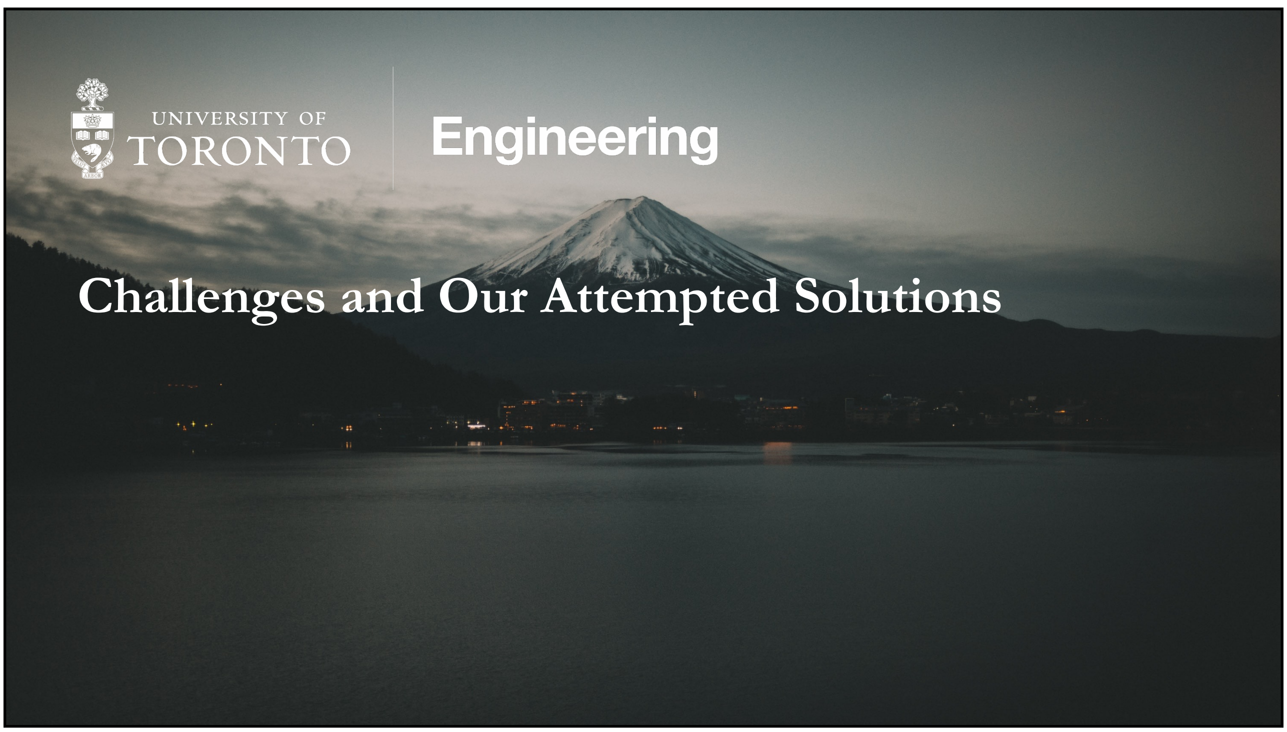

UNIVERSITY OF TORONTO | Engineering

Challenges and Our Attempted Solutions

23

Challenges to ML

- **Cost of Training**
  - Limited samples
  - Off-line training
- **Explainability**
- **Safety**
- Generalization
- **Delayed Reward** (RL)

UNIVERSITY OF TORONTO | Engineering

24

Transfer Learning

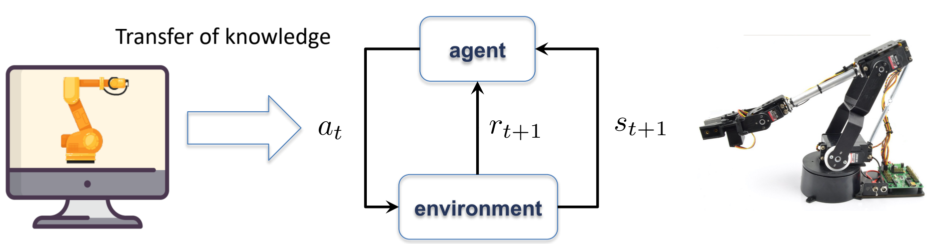

UNIVERSITY OF TORONTO | Engineering

25

Transfer RL with Multiple Experts

- Agent is to move from S to G
- Agent is to collect flags in the order of 1,2, 3 & 4
- Cost of moving is 1
- Invalid move will have an extra cost 1
  - For example, up move from the top edge
  - Flags collected in the wrong order

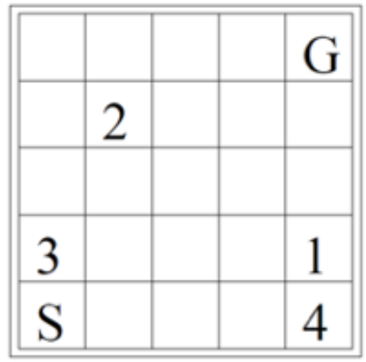

Gimelfarb, Sanner & Lee, "Reinforcement Learning with Multiple Expert: A Bayesian Model Combination Approach," NIPS 2018.

UNIVERSITY OF TORONTO | Engineering

26

Ex 1: Grid World

- Given 5 experts
  - Expert with optimal value function
  - Expert with a good heuristic: -22(5 – c – 0.5)/5
  - Expert with null advise: 0
  - Expert with uniformly distributed advise: U~[-20,20]
  - Expert with (-1)* optimal value function

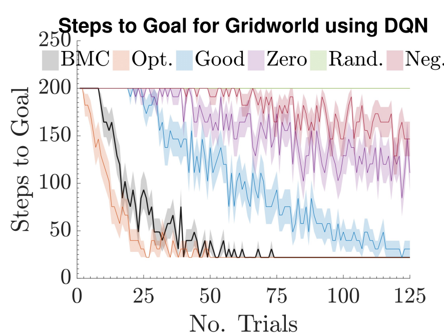

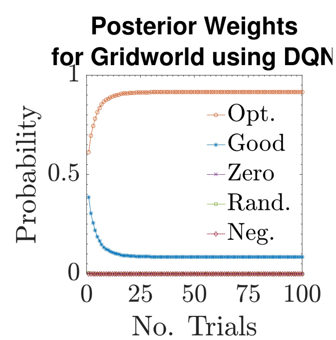

UNIVERSITY OF TORONTO | Engineering

27

Domain Adaptation

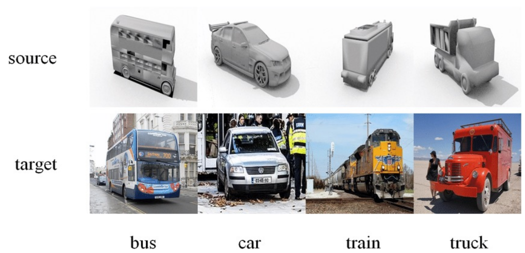

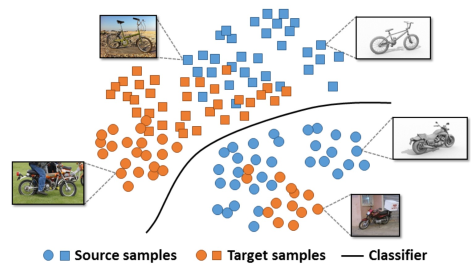

UNIVERSITY OF TORONTO | Engineering

28

Transfer Learning for Deep Diagnosis

- From Laboratory to the Field

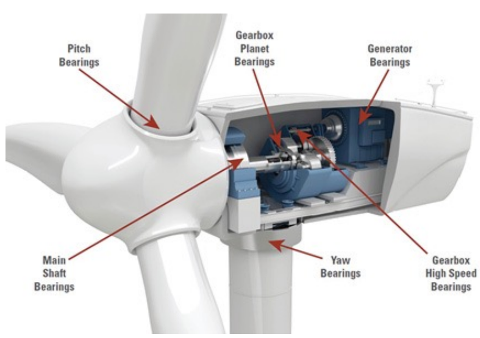

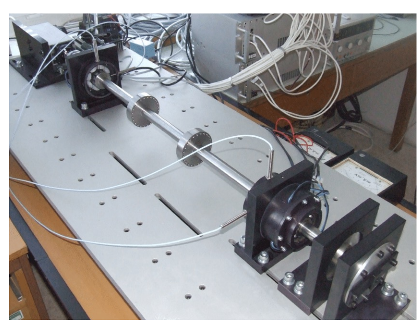

Yang, Lee, Lei, Li & Na, "Deep Partial Transfer Learning Network: A Method to Selectively Transfer Diagnostic Knowledge Across Related Machines," Mech Sys & Signal Processing, 2021

UNIVERSITY OF TORONTO | Engineering

29

Machine Learning without Data

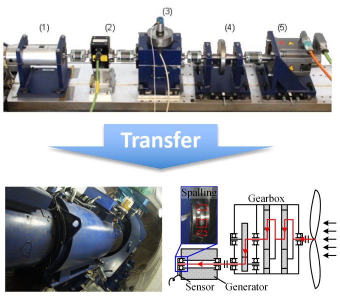

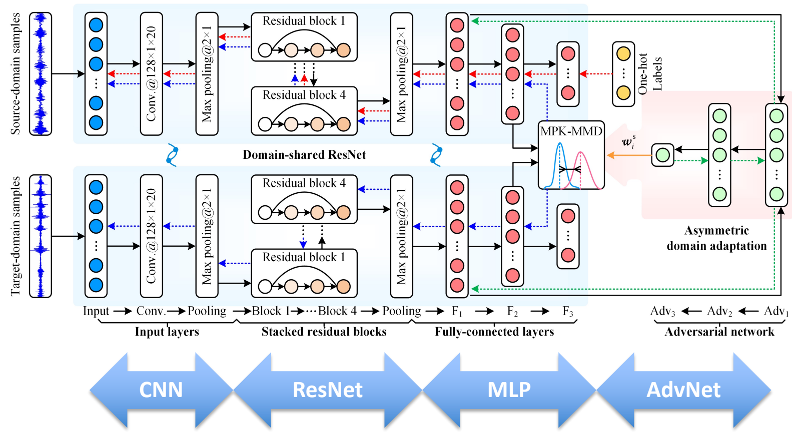

UNIVERSITY OF TORONTO | Engineering

30

Seq2Seq Learning

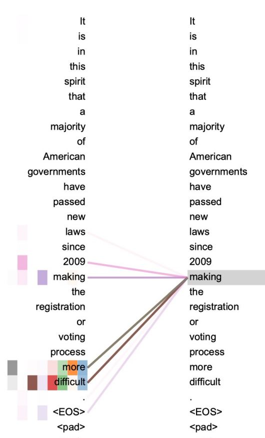

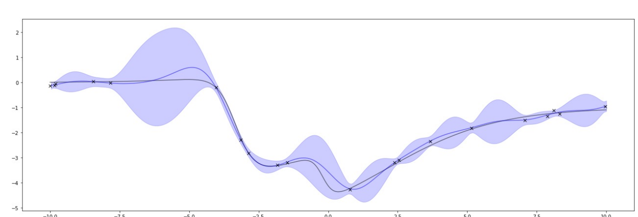

UNIVERSITY OF TORONTO | Engineering

31

Data-driven Digital Twin

- Digital twin can be a platform for RL training
- Engineering knowledge is limited, but data is sufficient

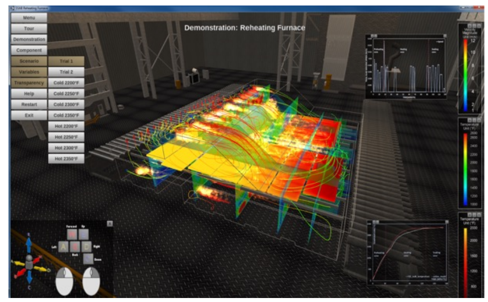

UNIVERSITY OF TORONTO | Engineering

32

Attentive-GP

- Transformer

$$\{\mathbf{x}_1, \dots, \mathbf{x}_N\} \ \& \ \{\mathbf{y}_1, \dots \mathbf{y}_{i-1}\}$$

$$\overline{\mathbf{x}}_i = \phi(\{\mathbf{x}_1, \dots, \mathbf{x}_N\}, \{\mathbf{y}_1, \dots \mathbf{y}_{i-1}\})$$

- GP regression layer

$$\mathbf{y}_i = f(\overline{\mathbf{x}}_i) + \epsilon_i$$

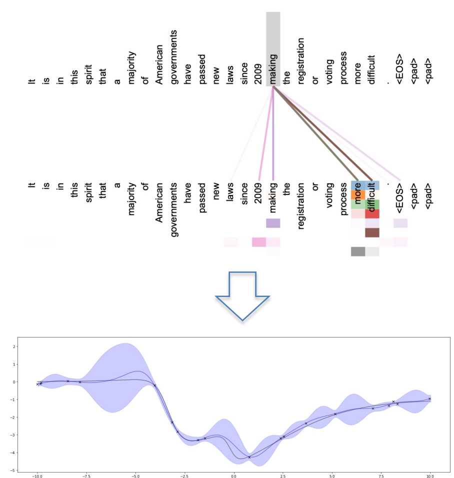

UNIVERSITY OF TORONTO | Engineering

33

Seq2Seq Learning

- The smart grid data from Kaggle
  - Global Energy Forecasting Competition
  - Input = hourly temperature at 11 areas,
  - Output = the total electrical load on the grid
  - Data collected over 2004 ~ 2008

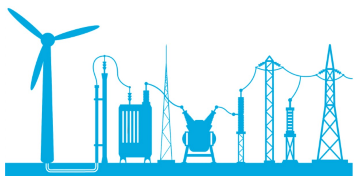

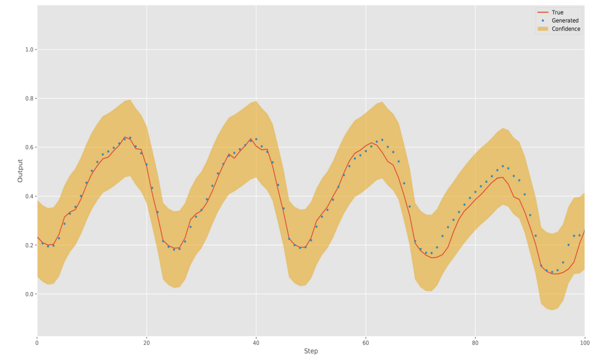

Generated Output Sequence of Smart Grid

Source: https://www.kaggle.com/

UNIVERSITY OF TORONTO | Engineering

34

Learning to Learn

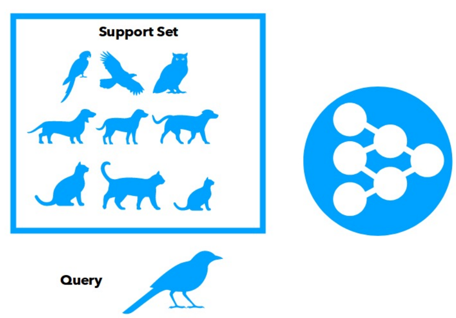

UNIVERSITY OF TORONTO | Engineering

35

Few-shot Learning

- Learn to classify using only a few samples per class
  - N-way k-shot: novel training set has N classes with k samples per class
- Our Algorithm
  - Meta-training for transferable knowledge
  - Parameter adaptation for a novel task
    - Bi-level optimization
    - One (or a few) gradient optimization step
    - Bayesian regression done by SWA

Chen & Lee, "Incremental Few-shot Learning via Vector Quantization in Deep Embedded Space," ICLR 2011.

UNIVERSITY OF TORONTO | Engineering

36

Incremental Few-shot Learning

- Performance Comparison
  - Caltech-UCSD Birds 200

Method	sessions										
	1	2	3	4	5	6	7	8	9	10	11
Fine-tune	77.30	46.23	34.71	25.35	23.16	20.65	16.21	13.32	11.98	11.17	10.76
Joint train	77.30	73.28	68.80	65.34	63.75	62.00	60.81	59.71	**59.06**	**58.69**	**58.23**
iCaRL	77.30	57.18	54.67	48.11	40.76	36.85	33.12	30.42	28.22	26.84	25.23
ProtoNet	77.30	69.76	66.01	62.29	59.58	57.10	55.13	54.09	52.40	51.65	50.36
SDC	77.34	74.45	69.45	65.27	61.81	58.26	56.14	55.71	53.31	52.79	51.52
Imprint	77.02	73.39	69.50	65.61	62.81	60.74	59.39	58.61	56.85	55.93	54.82
IDLVQ-C	**77.37**	**74.72**	**70.28**	**67.13**	**65.34**	**63.52**	**62.10**	**61.54**	**59.04**	**58.68**	**57.81**

UNIVERSITY OF TORONTO | Engineering

37

Incremental Few-shot Learning

- Regression task

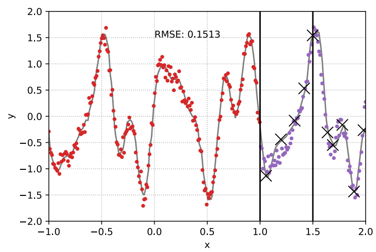

(a) Off-line NN

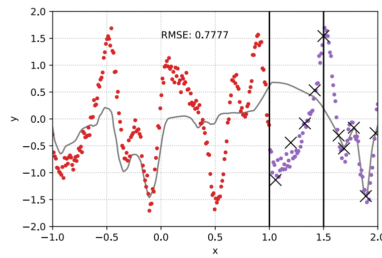

(b) Naive NN

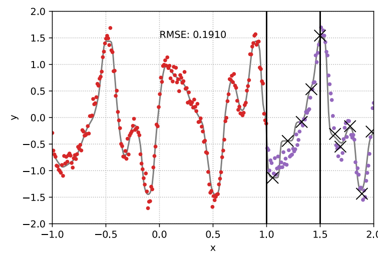

(c) Our algorithm

- **Extension**
  - Few-shot learning (not incremental)
  - Time-series regression (COVID-19, financial crisis, etc.)

UNIVERSITY OF TORONTO | Engineering

38

Closing

- **Opportunities in ML+SCM+OR**
  - Solution to complexity in modeling and solution
  - End-to-end solution due to DL
  - ML + OR is exciting
  - SCM has no shortage of interesting problems
- **Optimism with Caution**
  - ML is not a magic wand
  - Still immature technology
  - Investment and patience are required

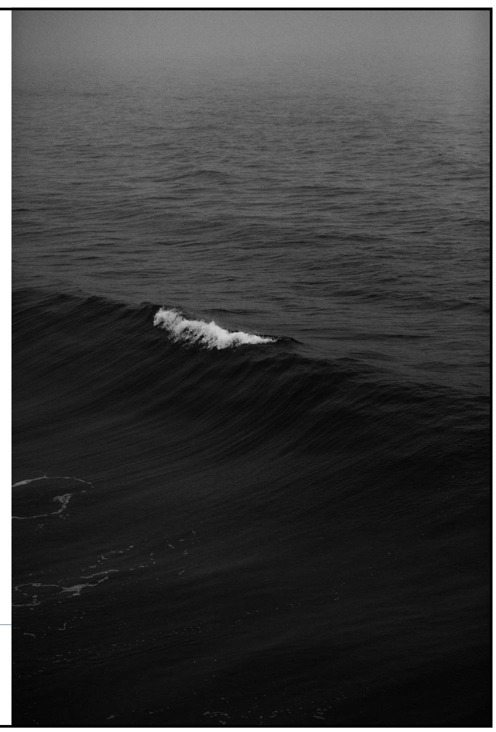

UNIVERSITY OF TORONTO | Engineering

39