MODERN AI/ML

REVIEW • OUTLOOK • EMPIRICAL RESEARCH • DEAKIN AI RESEARCH

Truyen Tran
HCMC 05/2019

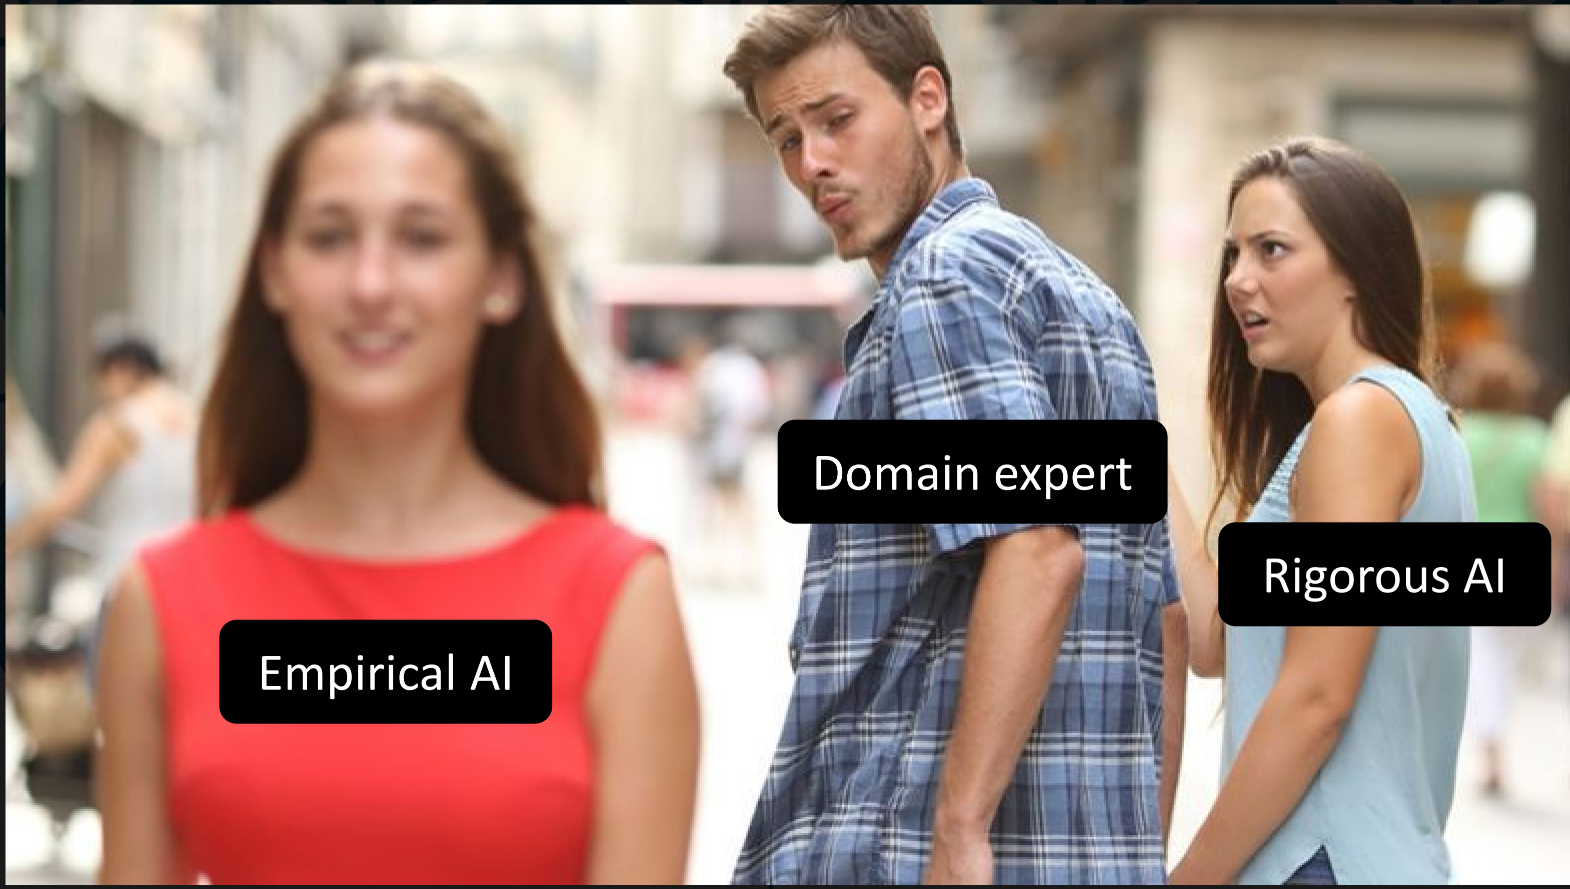

Agenda

- Refresher
- 2018-2019 Review
- Outlook
- Empirical AI Research
- DAIR: Deakin AI Research

AI/ML RESEARCH: DREAM

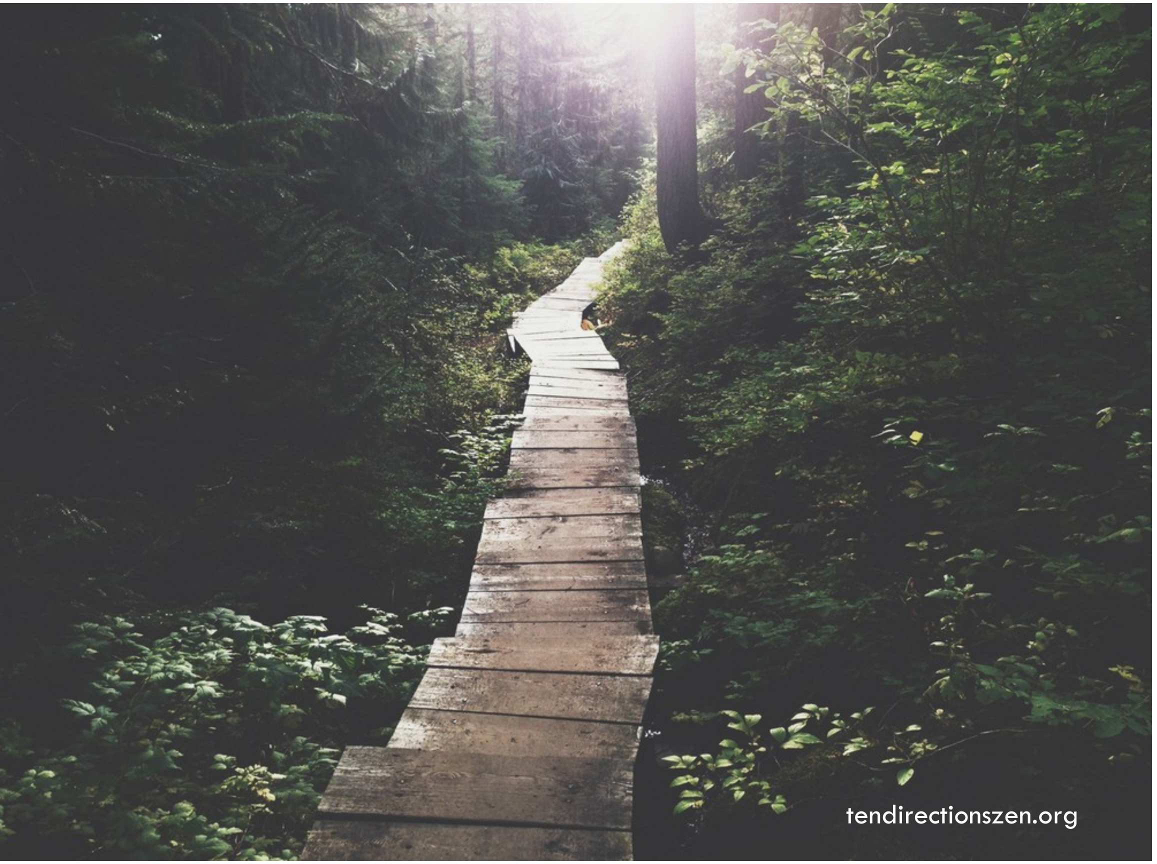

Among the most challenging scientific questions of our time are the corresponding analytic and synthetic problems:

- How does the brain function?
- Can we design a machine which will simulate a brain?

-- *Automata Studies*, 1956.

WHAT MAKES AI?

Perceiving

Learning

Reasoning

Planning

Acting

Robotics

Communicating

Consciousness

Automated discovery

Modern AI is mostly data-driven, as opposed to classic AI, which is mostly expert-driven.

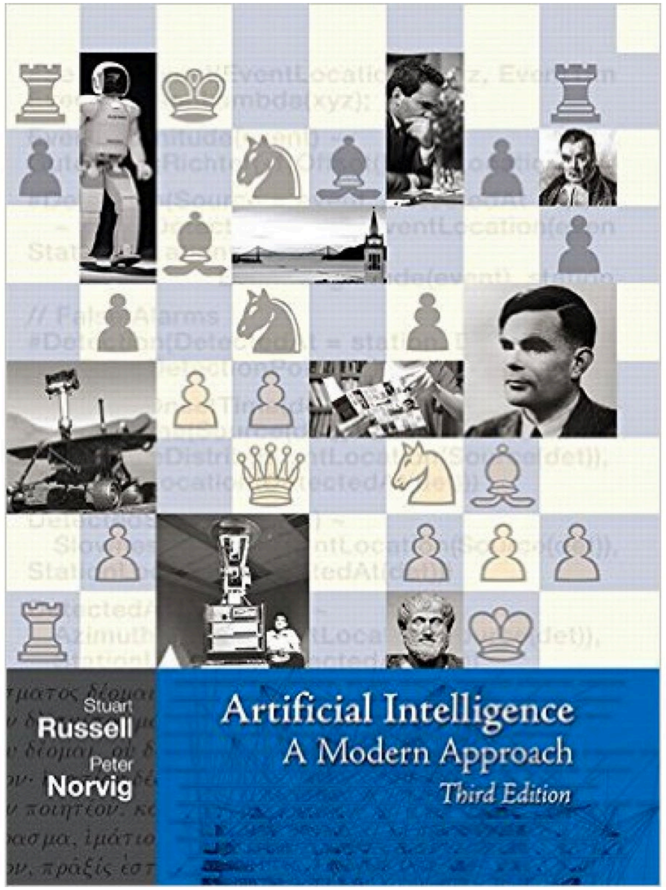

REASONING

Reasoning is to deduce knowledge from previously acquired knowledge in response to a query (or a cue)

Early theories of intelligence:

- focuses solely on reasoning,
- learning can be added separately and later! (Khardon & Roth, 1997).

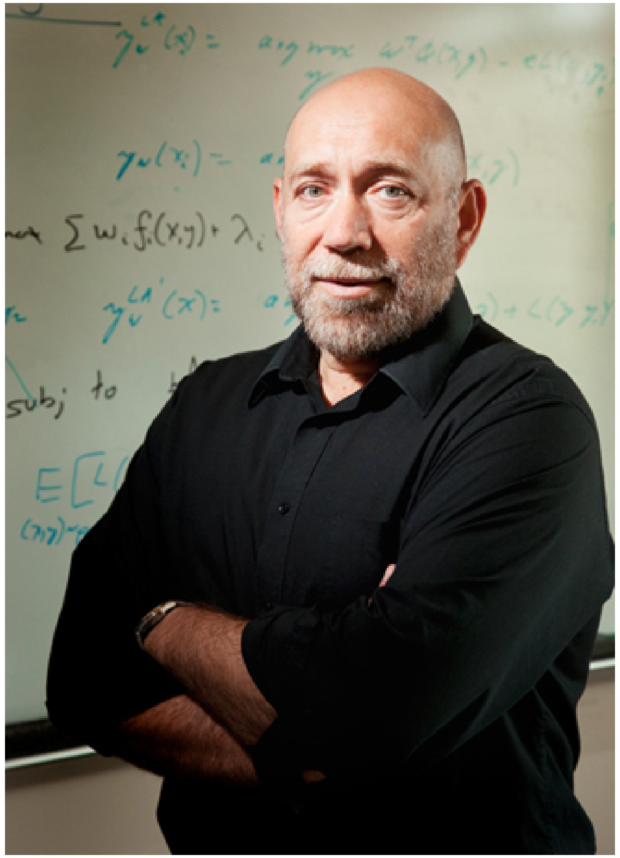

(Dan Roth; *ACM Fellow; IJCAI John McCarthy Award*)

Khardon, Roni, and Dan Roth. "Learning to reason." *Journal of the ACM (JACM)* 44.5 (1997): 697-725.

WHAT YOU CAN'T DESIGN, LEARN!
(AKA VARIATIONAL METHOD)

Filling the slot

- In-domain (intrapolation), e.g., an alloy with a given set of characteristics
- Out-domain (extrapolation), e.g., weather/stock forecasting
- Classsification, recognition, identification
- Action, e.g., driving
- Mapping space, e.g., translation
- Replacing expensive simulations

Estimating semantics, e.g., concept/relation embedding

Assisting experiment designs

Finding unknown, causal relation, e.g., disease-gene

Predicting experiment results, e.g., alloys → phase diagrams → material characteristics

MACHINE LEARNING SETTINGS

Supervised learning

(mostly machin

A → B

Will be quickly solved for "easy" problems (Andrew Ng)

Anywhere in between: semi-supervised learning, reinforcement learning, lifelong learning, meta-learning, few-shot learning, knowledge-based ML

Unsupervised learning

man)

$$\mathbf{v} \sim P_{model}(\mathbf{v})$$

$$(\mathbf{v}) \approx P_{data}(\mathbf{v})$$

Agenda

- Refresher
- 2018-2019 Review
- Outlook
- Empirical AI Research
- DAIR: Deakin AI Research

AI/ML RESEARCH: MODERN REALITY

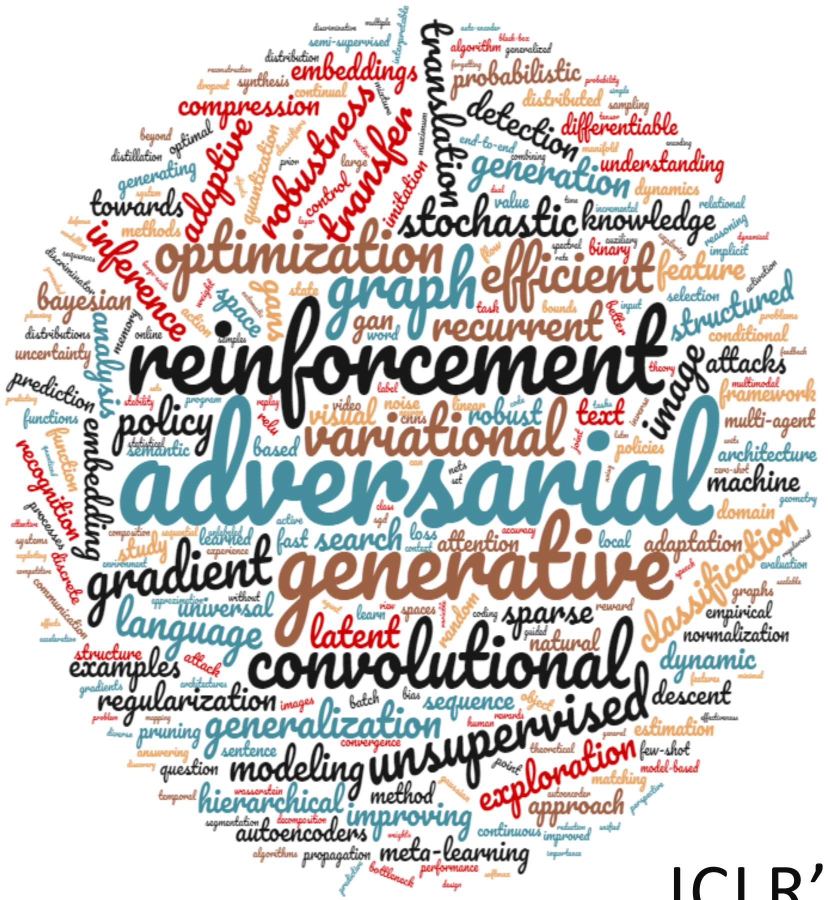

ICLR’19

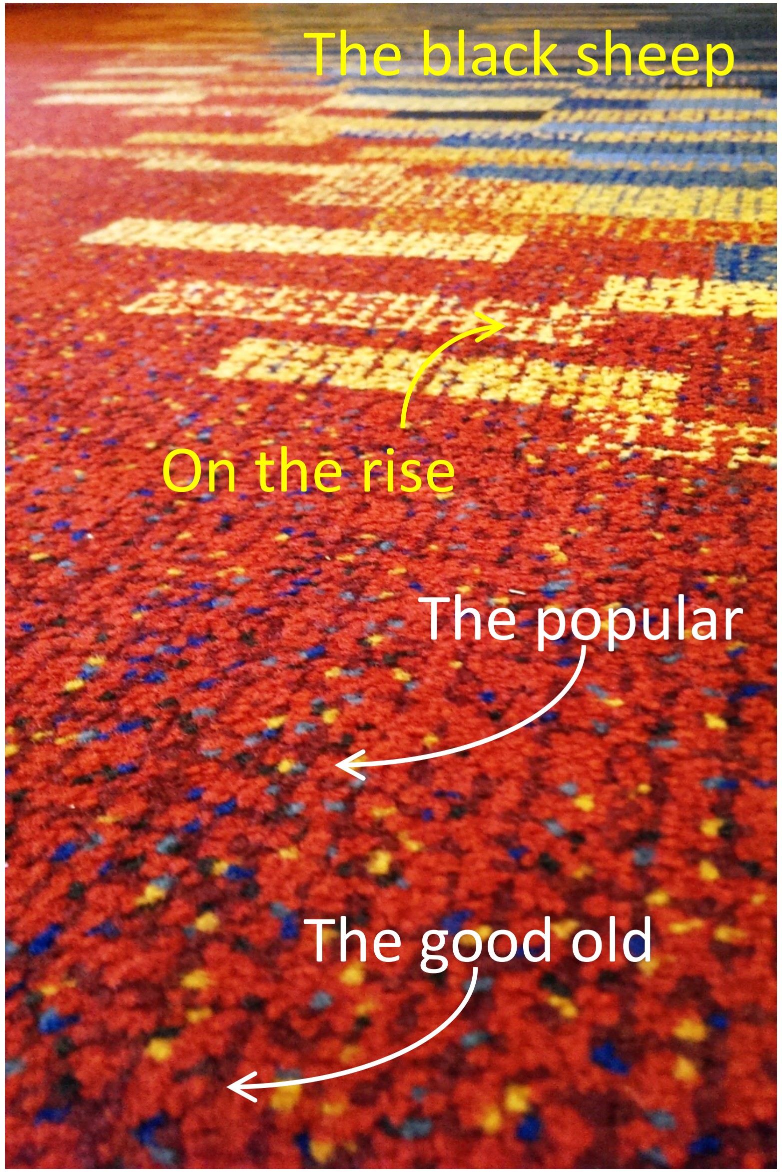

- Group theory
- Quantum ML/AI
- Theories of consciousness

- Reinforcement learning, imagination & planning
- Deep generative models + Bayesian methods
- **Memory & reasoning**
- Lifelong/meta/continual/few-shot/zero-shot learning
- Universal transformer

- Attention
- Batch-norm
- ReLU & skip-connections
- Highway nets, LSTM/GRU & CNN

- Representation learning (RBM, DBN, DBM, DDAE)
- Ensemble
- Back-propagation
- Adaptive stochastic gradient

"deep learning" "artificial intelligence" "machine learning" Worldwide, 10/7/12 - 11/7/18

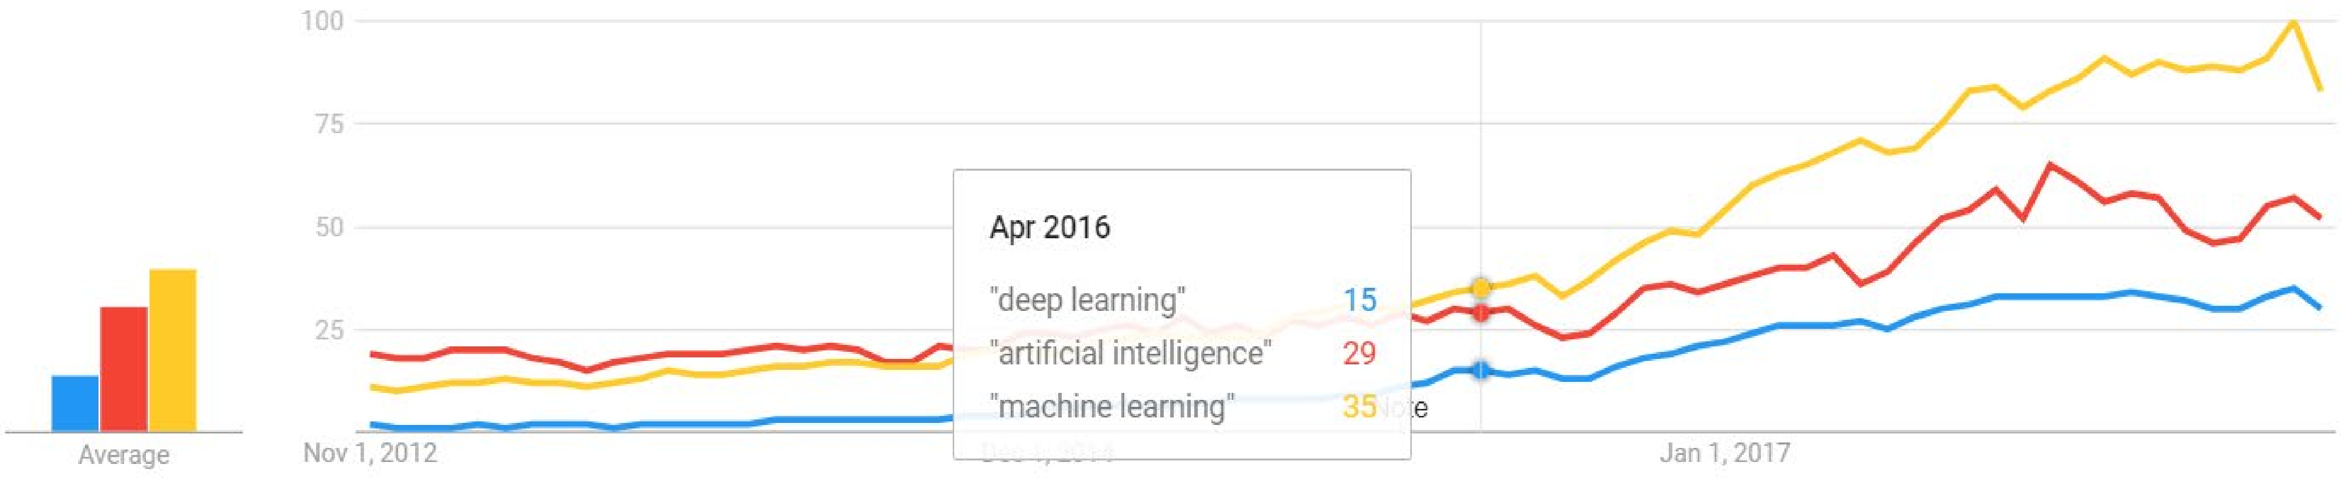

Compared breakdown by region

Region

"deep learning" "artificial intelligence" "machine learning"

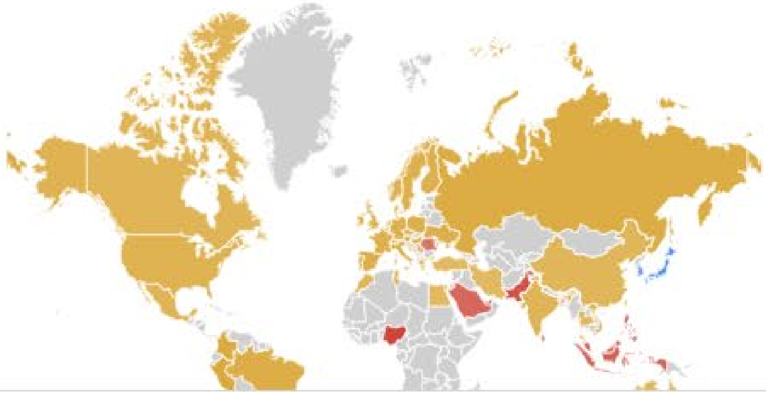

Sort: Interest for "deep learning"

1 Japan

2 South Korea

3 China

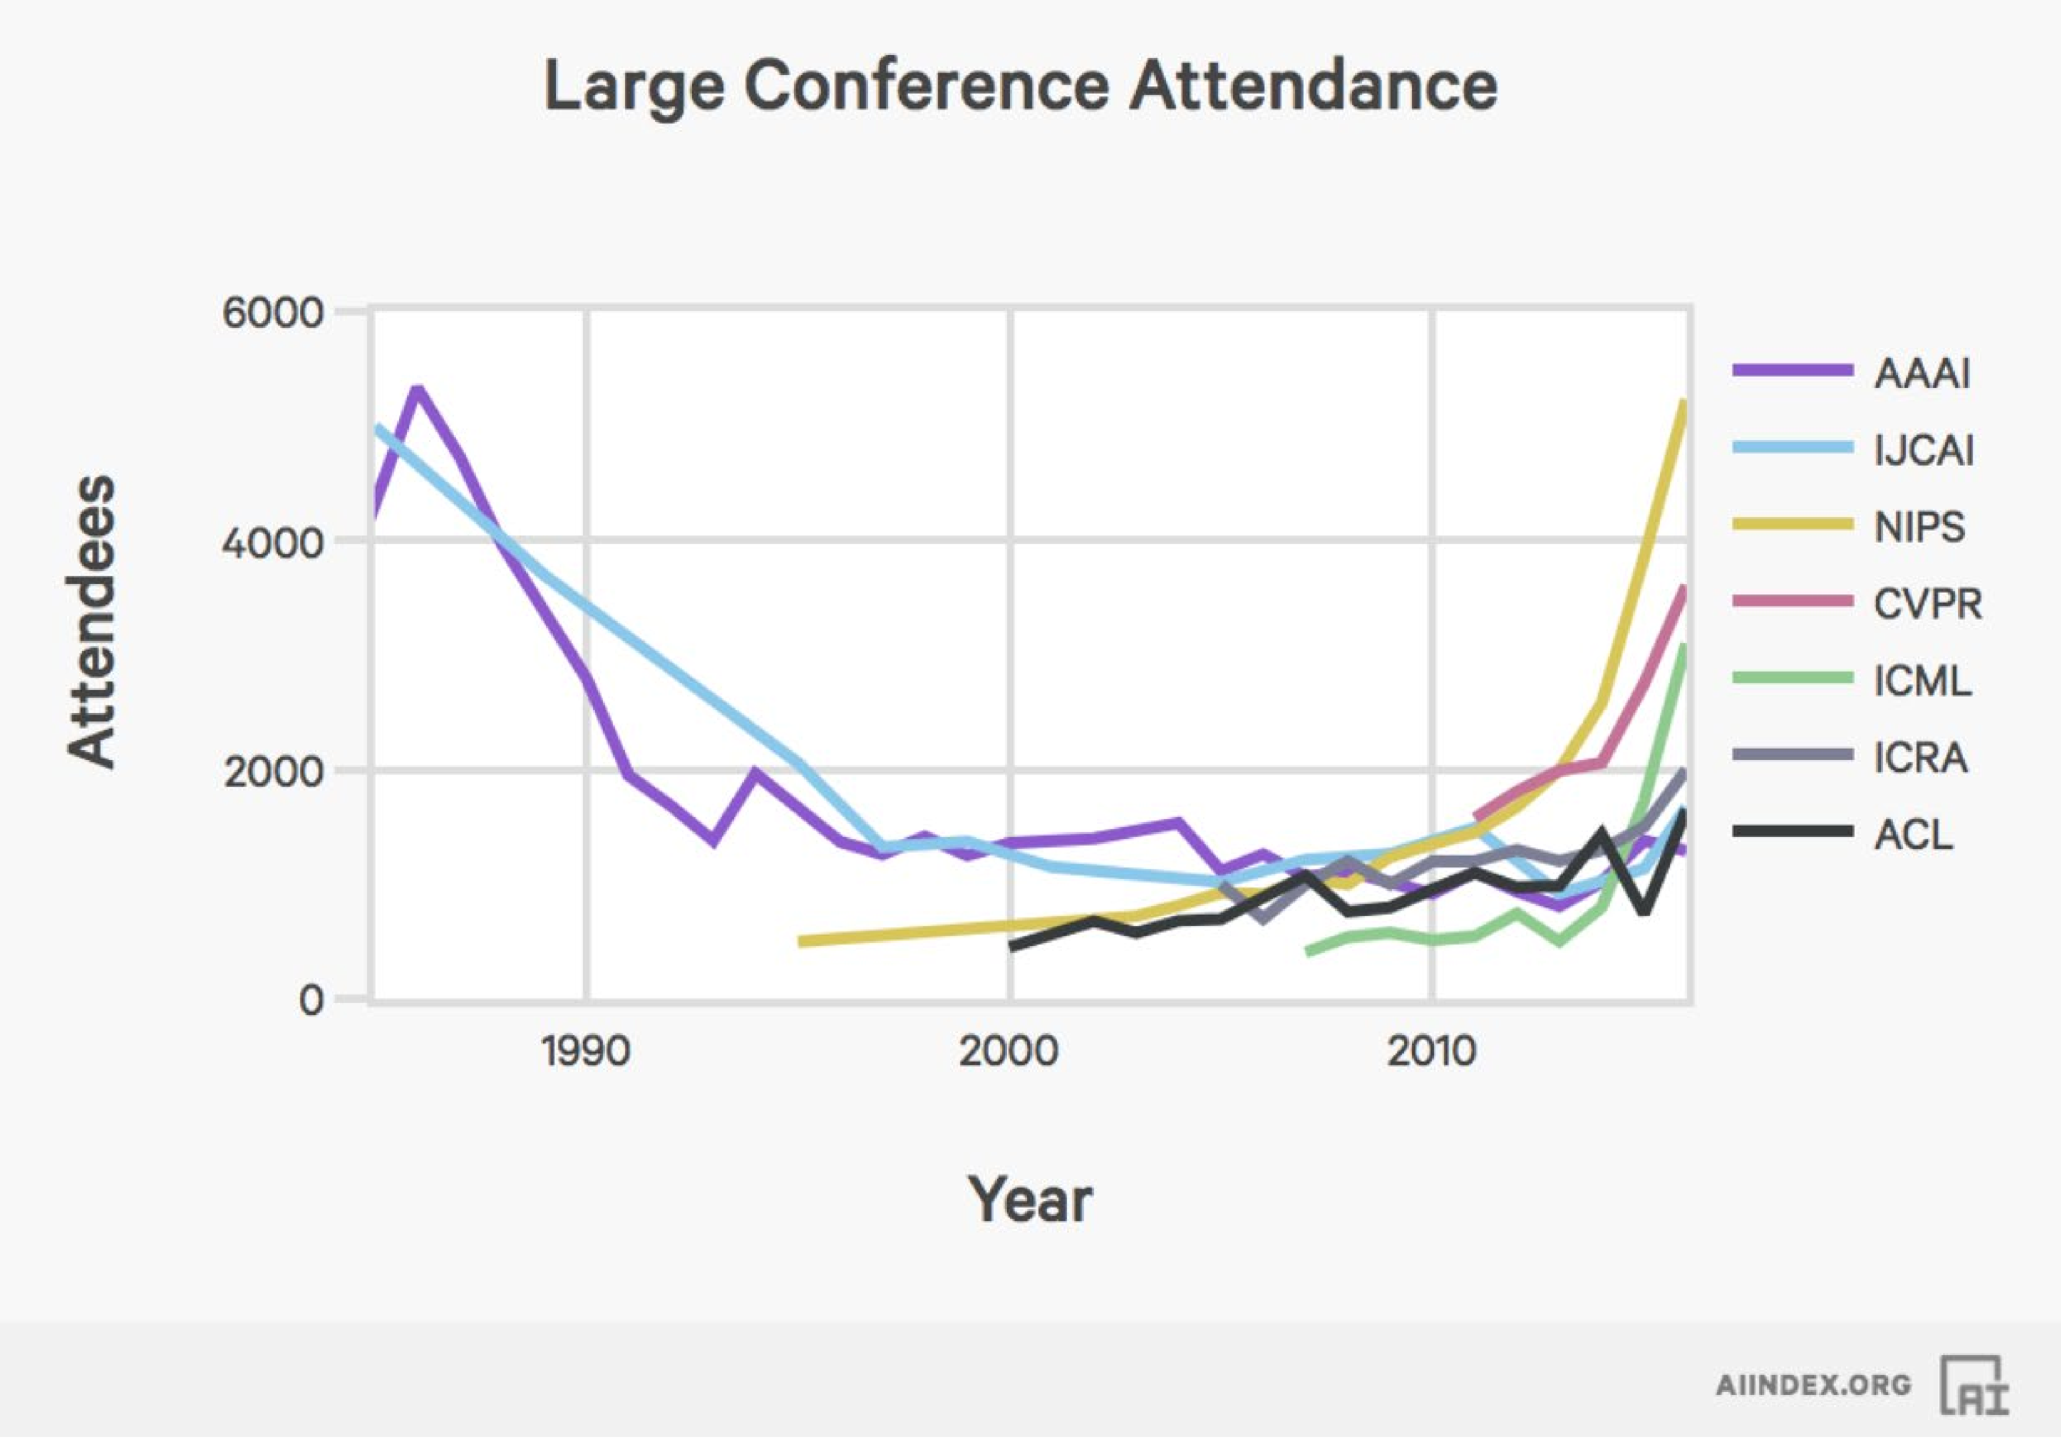
Large Conference Attendance
Attendees
6000
4000
2000
0
1990
2000
2010
Year
AAAI
IJCAI
NIPS
CVPR
ICML
ICRA
ACL
AIINDEX.ORG

MIT TECH REVIEW: 10 BREAKTHROUGHS OF 2018

3-D Metal Printing (Markforged, Desktop Metal, GE)

Artificial Embryos, aka stem cells (University of Cambridge; University of Michigan; Rockefeller University)

Sensing City, aka smart cities (Sidewalk Labs and Waterfront Toronto, 1-2 years)

AI for Everybody, aka Cloud AI (Amazon; Google; Microsoft)

Dueling Neural Networks, aka GAN (Google Brain, DeepMind, Nvidia)

Babel-Fish Earbuds, aka near real time translation (Google and Baidu)

Zero-Carbon Natural Gas (8 Rivers Capital; Exelon Generation; CB&I; 3-5 years to come)

Perfect Online Privacy (Zcash; JPMorgan Chase; ING)

Genetic Fortune-Telling, aka DNA-based predictions (Helix; 23andMe; Myriad Genetics; UK Biobank; Broad Institute)

Materials' Quantum Leap, aka Quantum computing for molecules (IBM; Google; Harvard's Alán Aspuru-Guzik; 5-10 years to come)

#Ref: https://www.technologyreview.com/lists/technologies/2018/

International Conference on Machine Learning (ICML 2018)

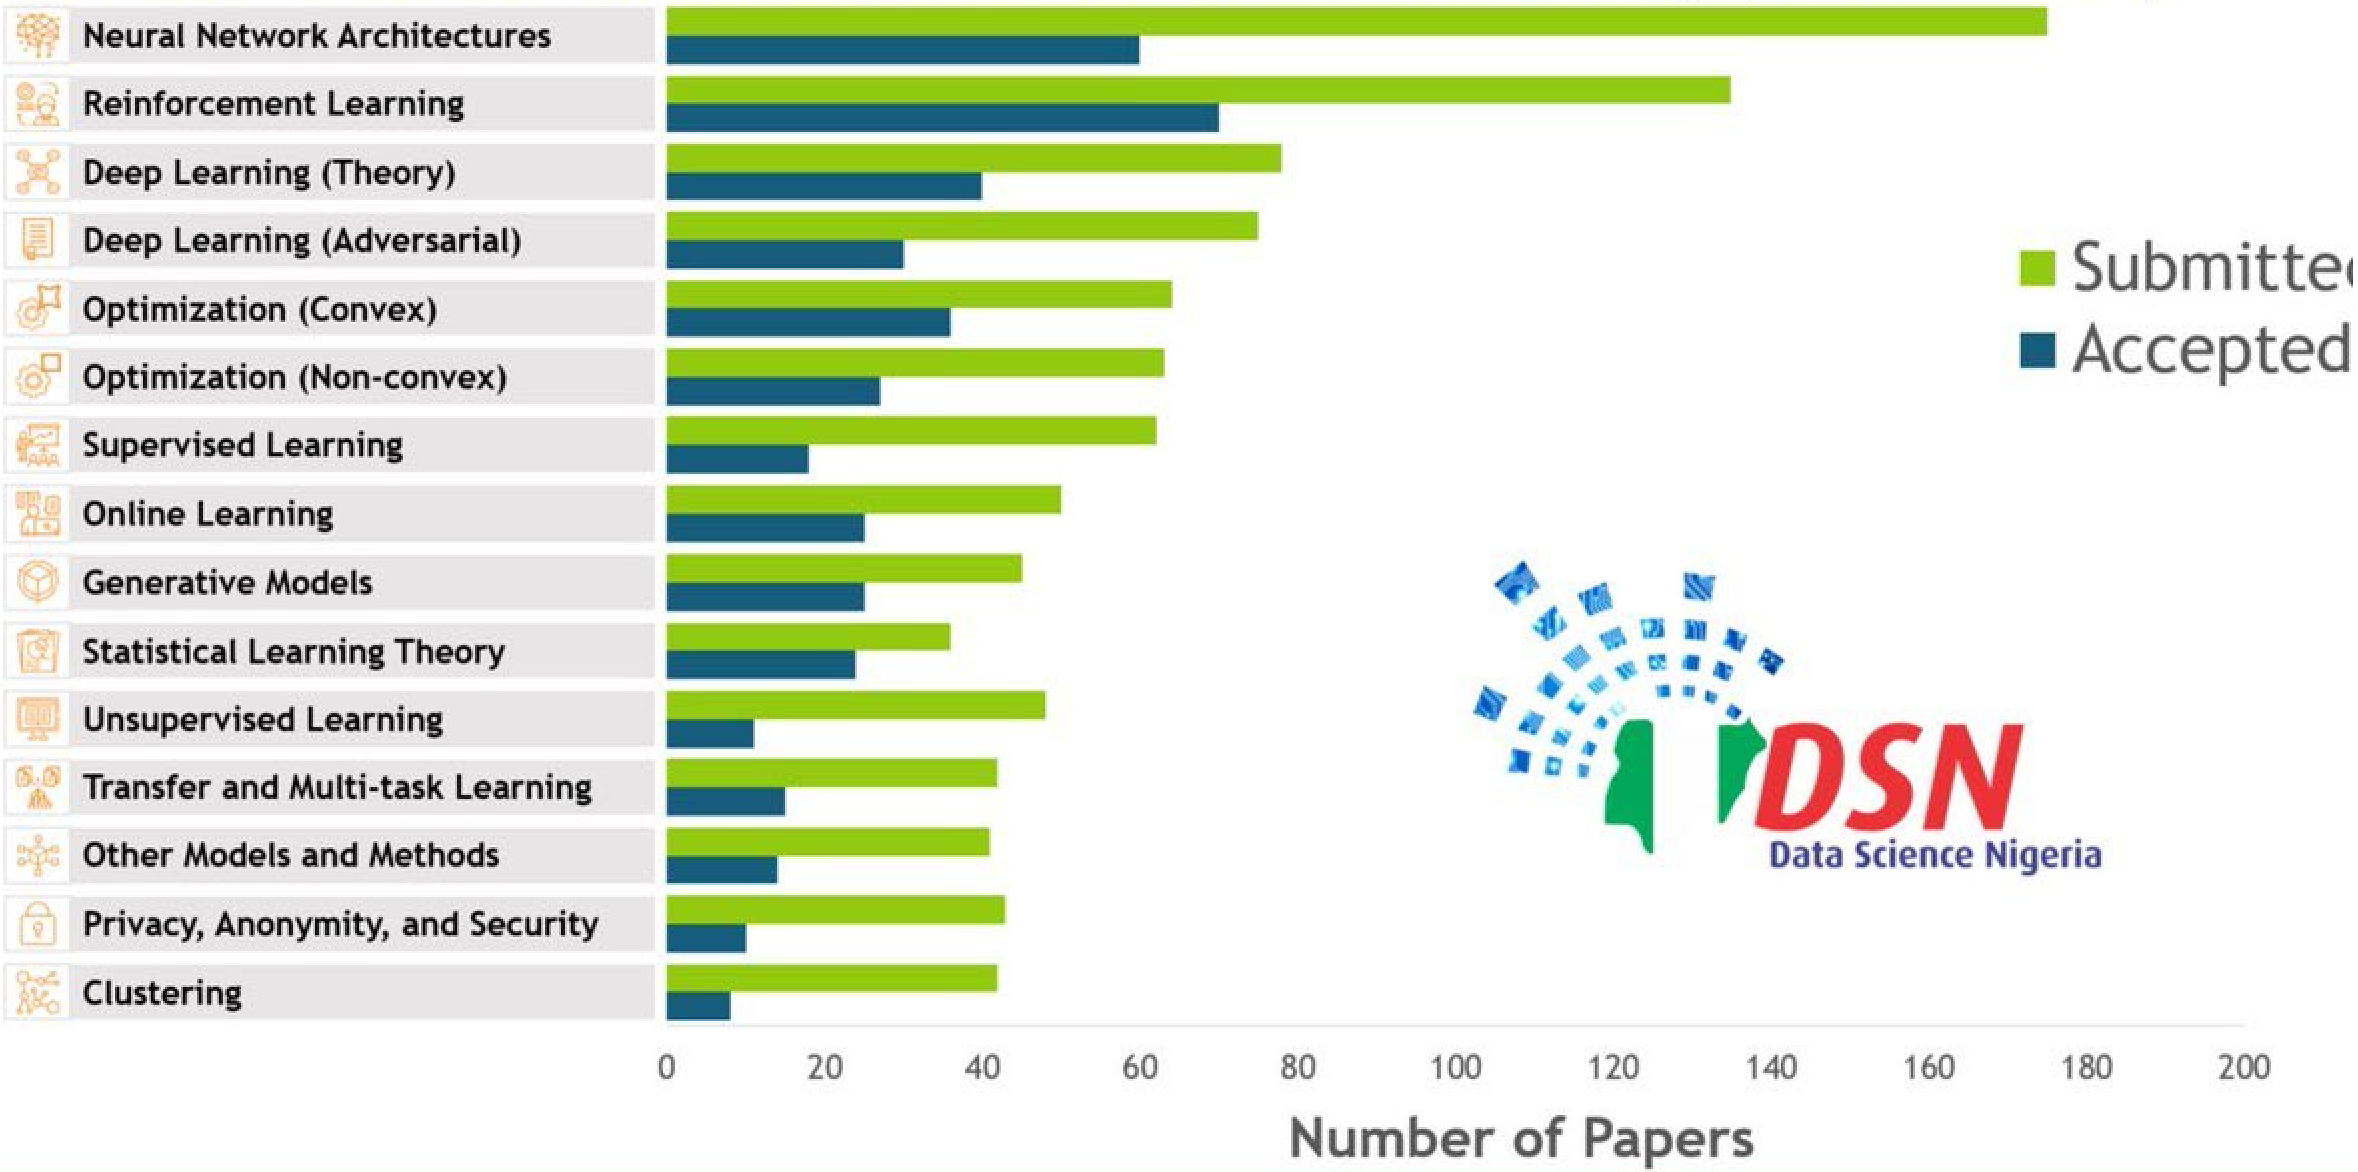

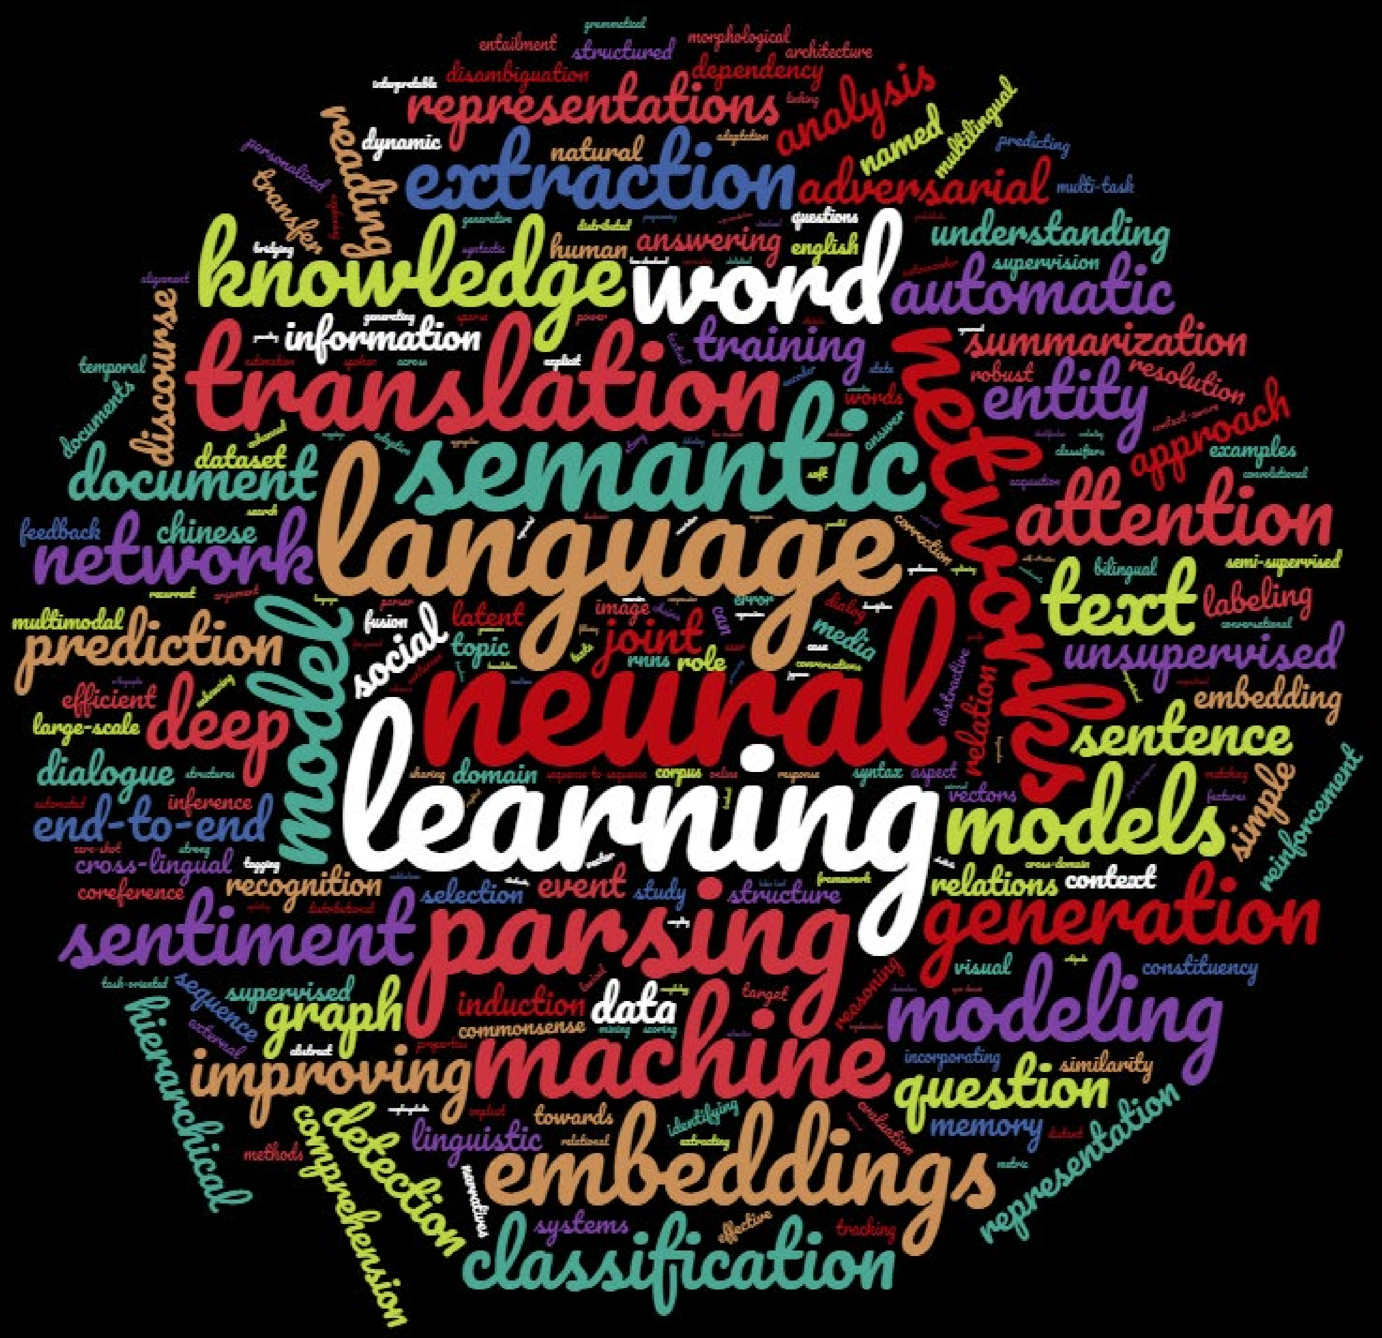
neural
learning
language
semantic
translation
networks
word
knowledge
extraction
representations
parsing
machine
embeddings
classification
model
models
generation
modeling
attention
text
network
document
sentiment
prediction
deep
reading
analysis
adversarial
automatic
summarization
entity
training
information
understanding
unsupervised
sentence
question
improving
graph
detection
comprehension
hierarchical
representation
dialogue
end-to-end
recognition
approach
embedding
reinforcement
simple
relations
labeling
answering
discourse
documents
temporal
feedback
chinese
multimodal
efficient
large-scale
inference
cross-lingual
coreference
supervised
sequence
induction
commonsense
linguistic
systems
towards
memory
similarity
constituency
visual
data
event
selection
structure
study
context
domain
topic
joint
social
latent
dataset
named
multilingual
natural
dynamic
transfer
human
english
supervision
resolution
robust
examples
semi-supervised
bilingual
media
relation
role
corpus
syntax
aspect
vectors
tracking
methods
disambiguation
dependency
structured
morphological
entailment
architecture
personalized
multi-task
predicting
questions
convolutional
conversational
fusion
image
abstractive
reasoning
identifying
incorporating
effective
evaluation
external
words
answer
classifiers
error
recurrent
parser
generating
task-oriented
zero-shot
tagging
framework
sequence-to-sequence
response
matching
features
distant
metric
relational
extracting
target

ACL’18

ACL'18

Ignore the obvious


learning
embedding
networks
prediction
deep
detection
search
recommendation
graph
neural
online
discovery
dynamic
clustering
efficient
adversarial
information
optimization
framework
reinforcement
scalable
approach
representation
large-scale
system
models
modeling
analysis
machine
model
time
fast
towards
text
knowledge
spatio-temporal
heterogeneous
interpretable
supervised
metric
active
based
visual
ranking
adaptive
transfer
planning
quality
improving
multi-task
real-time
mining
inference
recommender
recurrent
robust
mobile
tree
local
user
memory
streams
structural
extraction
mechanism
semantic
event
big
regression
social
systems
latent
distributed
hashing
risk
behavior
new
estimation
attention
interactive
optimal
discovering
feature
algorithm
recommendations
detecting
product
empirical
application
e-commerce
classification

KDD’18

KDD’18

Ignore the obvious

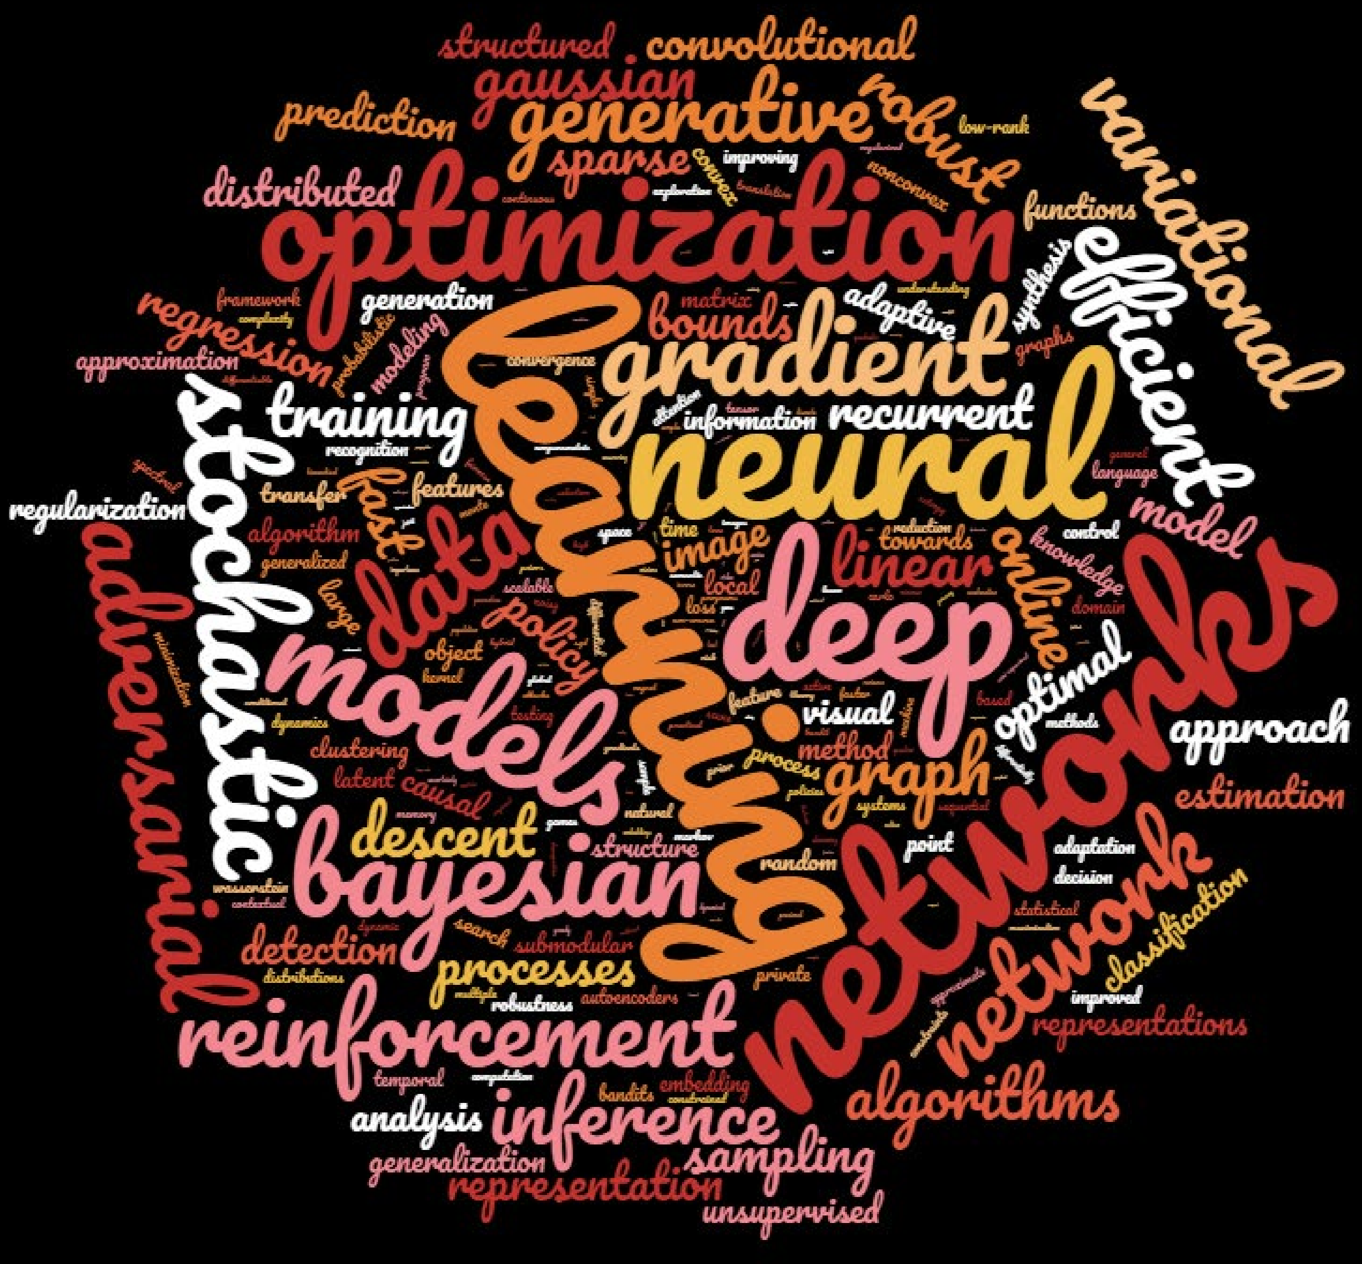
learning
networks
neural
optimization
deep
models
stochastic
bayesian
reinforcement
gradient
adversarial
generative
network
variational
efficient
inference
data
algorithms
training
graph
recurrent
descent
linear
gaussian
robust
convolutional
structured
prediction
sparse
distributed
regression
bounds
fast
image
online
model
optimal
approach
processes
detection
representations
representation
sampling
analysis
generalization
unsupervised
classification
estimation
policy
regularization
approximation
framework
generation
adaptive
matrix
functions
information
recognition
features
transfer
algorithm
time
towards
knowledge
control
visual
method
process
clustering
latent
causal
structure
random
point
adaptation
decision
statistical
search
submodular
autoencoders
improved
robustness
embedding
bandits
temporal
private
low-rank
convex
nonconvex
improving
synthesis
modeling
probabilistic
convergence
local
large
object
kernel
loss
domain
scalable
generalized
language
methods
dynamics
systems
natural
wasserstein
distributions

NIPS'18

NIPS’18

Ignore the obvious


learning
neural
networks
network
deep
model
embedding
data
classification
optimization
adversarial
knowledge
efficient
image
detection
reinforcement
approach
recognition
convolutional
unsupervised
representation
prediction
attention
graph
models
generation
planning
recommendation
hierarchical
temporal
information
semantic
machine
text
online
feature
framework
analysis
bayesian
video
language
object
dynamic
games
selection
matching
transfer
translation
robust
visual
inference
optimal
clustering
recurrent
question
answering
multi-agent
sampling
partial
estimation
modeling
adaptation
reasoning
systems
variational
retrieval
spatial
segmentation
stochastic
relation
sequential
probabilistic
programming
factorization
label
videos
action
towards
domain
chinese
search
decision
joint
distribution
training
natural
multi-task
constraints
representations
problems
improving
adaptive
distributed
understanding
heterogeneous
semi-supervised
multi-view
algorithm
random
gradient
graphs
word
answer
tasks
automatic
approximate
social
problem
person
sequence
regression
supervised
completion
features
captioning
large
extraction
time
human
method
attentive
better
memory
generative
complexity
re-identification
metric
sparse
face
multiple
structure
exploiting
machine
latent
sentence
alignment
community
robustness
academy

AAAI’19


adversarial
reinforcement
generative
convolutional
unsupervised
optimization
graph
efficient
variational
stochastic
translation
inference
robustness
transfer
adaptive
gradient
language
generalization
regularization
examples
modeling
hierarchical
improving
exploration
classification
latent
sparse
recurrent
knowledge
detection
generation
structured
policy
image
embedding
analysis
bayesian
towards
compression
embeddings
probabilistic
differentiable
understanding
feature
attacks
framework
multi-agent
architecture
machine
domain
adaptation
attention
search
dynamic
descent
normalization
empirical
meta-learning
autoencoders
approach
method
sequence
estimation
robust
text
natural
question
sentence
pruning
universal
structure
recognition
prediction
uncertainty
distributions
methods
generating
distillation
quantization
imitation
control
continual
synthesis
distributed
dynamics
selection
conditional
policies
semantic
visual
noise
linear
local
fast
loss
study
learned
few-shot
matching
propagation
continuous
improved
algorithm
semi-supervised
distribution

ICLR’19

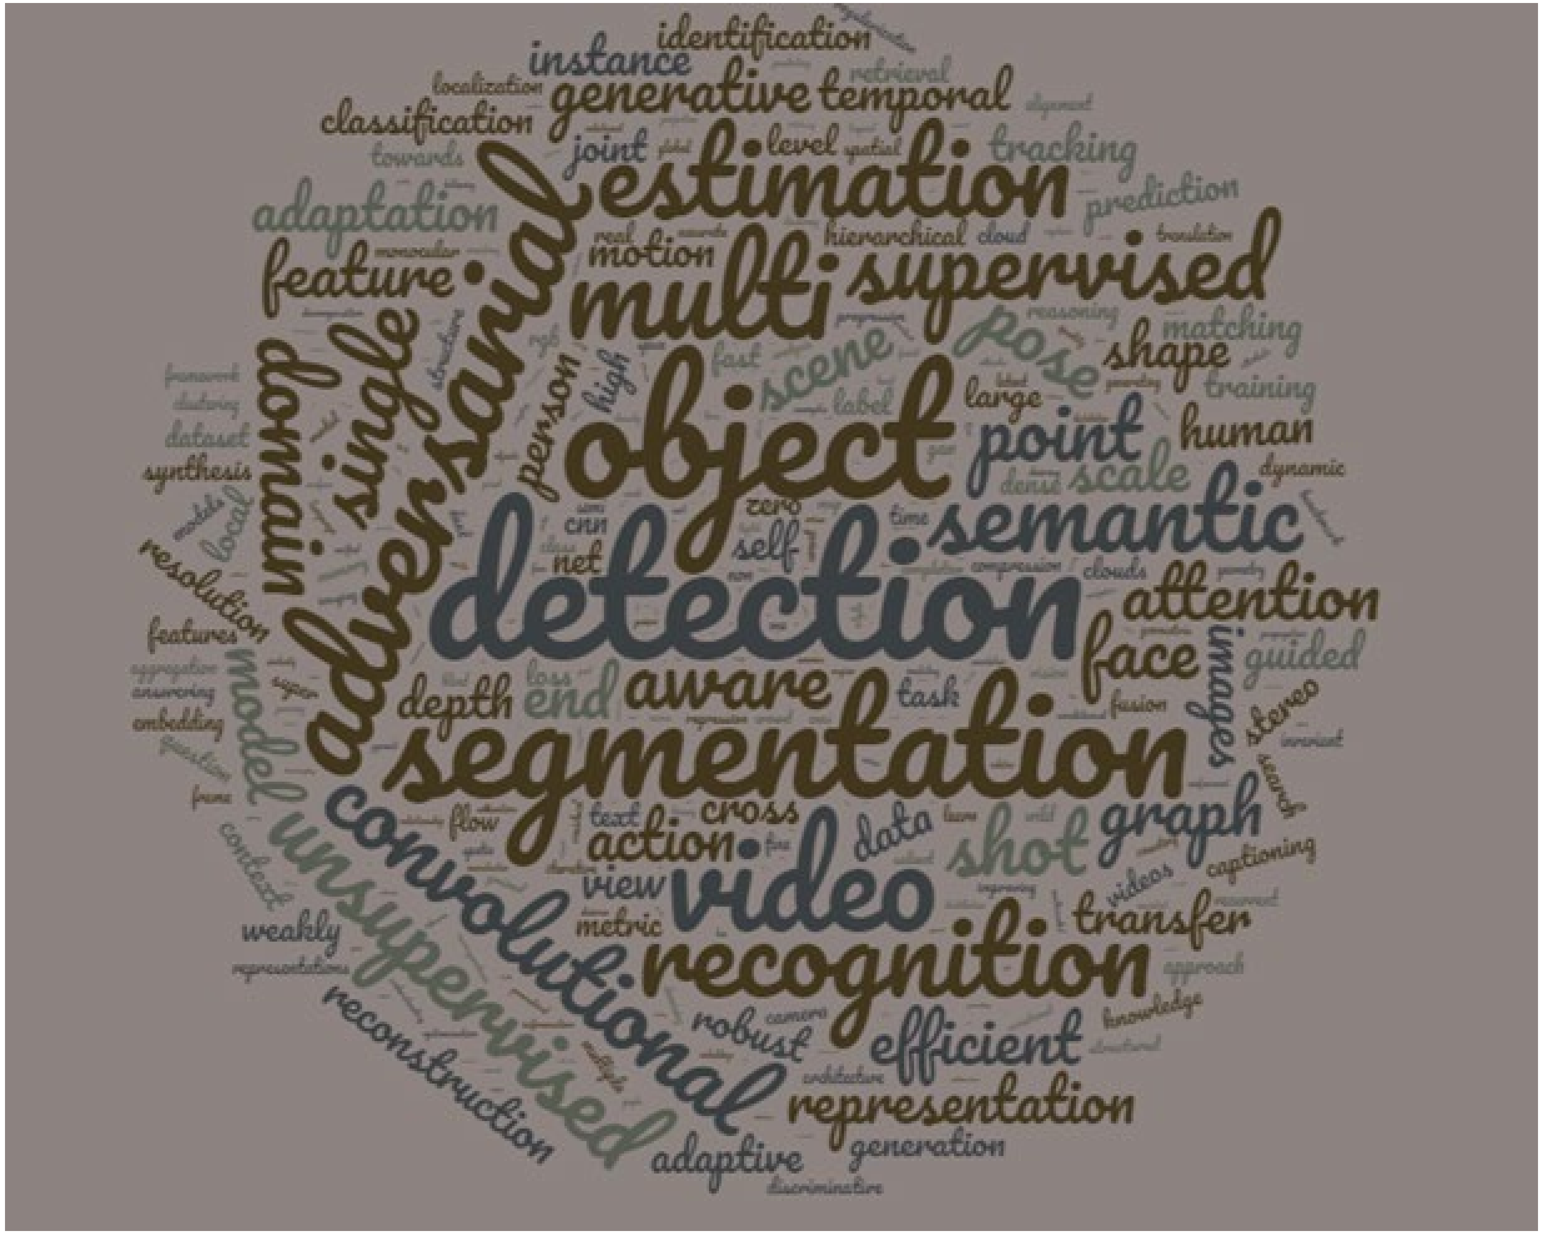
detection
object
segmentation
adversarial
multi
estimation
supervised
video
recognition
unsupervised
convolutional
semantic
domain
single
point
attention
face
aware
graph
shot
efficient
representation
reconstruction
generative
temporal
feature
adaptation
pose
scene
shape
human
scale
images
action
transfer
model
depth
end
cross
data
view
metric
robust
adaptive
generation
identification
instance
classification
tracking
prediction
motion
hierarchical
joint
level
matching
training
large
person
high
fast
label
self
net
zero
cnn
local
resolution
features
synthesis
dataset
weakly
context
flow
text
stereo
guided
captioning
videos
knowledge
approach
camera
architecture
discriminative
task
loss
dynamic
clouds
compression
fusion
time
retrieval
localization
towards
real
cloud
reasoning
embedding
answering
aggregation
framework
clustering
representations
invariant
translation
monocular

CVPR’19

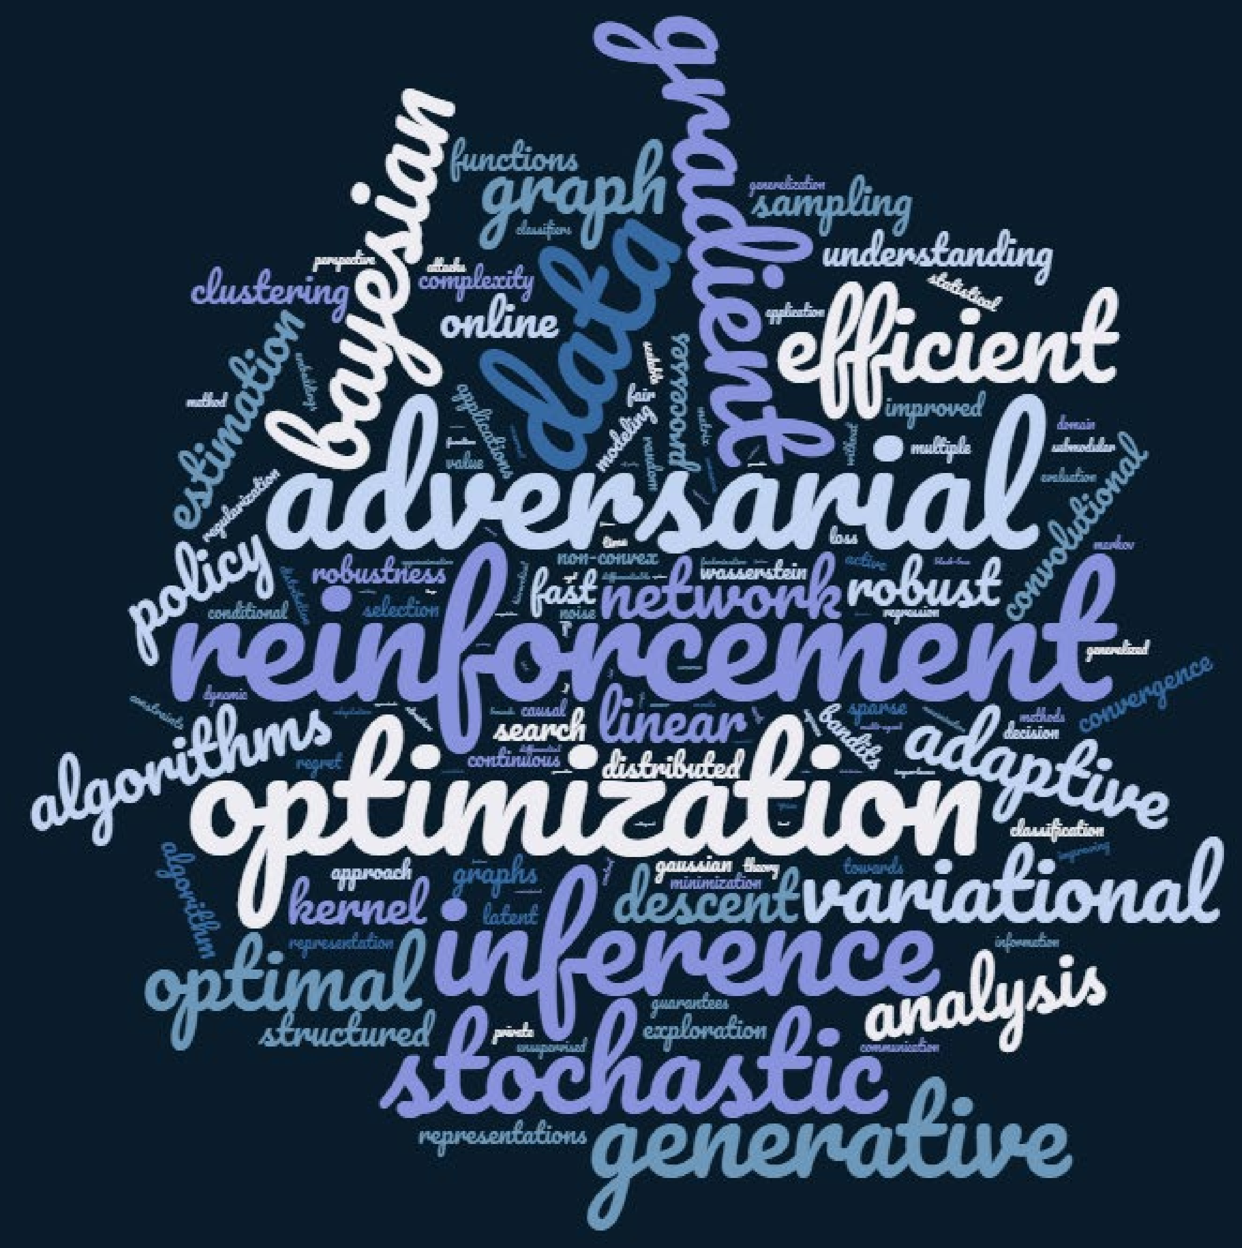

ICML'19

Ignore the obvious

A LOOK INTO NIPS'18 WORKSHOPS (1)

Uncertainty

Bayesian Deep Learning

All of Bayesian Nonparametrics (Especially the Useful Bits)

Data efficiency

Continual Learning

NIPS 2018 Workshop on Meta-Learning

Reinforcement learning

Deep Reinforcement Learning

Imitation Learning and its Challenges in Robotics

Reinforcement Learning under Partial Observability

Infer to Control: Probabilistic Reinforcement Learning and Structured Control

A LOOK INTO NIPS'18 WORKSHOPS (2)

Communication

Learning by Instruction

Emergent Communication Workshop

The second Conversational AI workshop – today's practice and tomorrow's potential

Visually grounded interaction and language

Theories & modelling

Integration of Deep Learning Theories

Relational Representation Learning

Modeling the Physical World: Learning, Perception, and Control

Impact of AI

Workshop on Ethical, Social and Governance Issues in AI

AI for social good

RICH SUTTON'S BITTER LESSON

"The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. "

"The two methods that seem to scale arbitrarily in this way are *search* and *learning*."

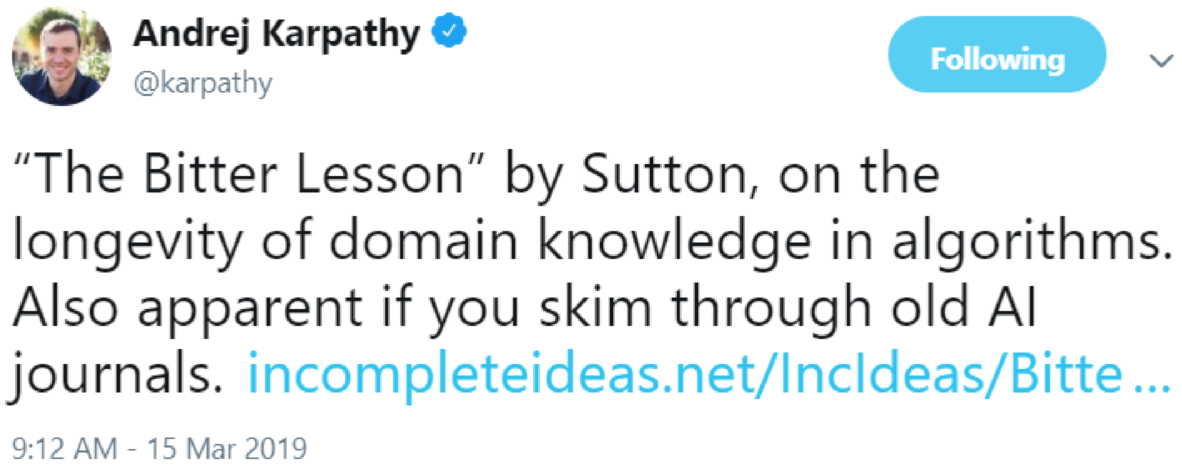

http://www.incompleteideas.net/IncIdeas/BitterLesson.html

Agenda

- Refresher
- 2018-2019 Review
- Outlook
- Empirical AI Research
- DAIR: Deakin AI Research

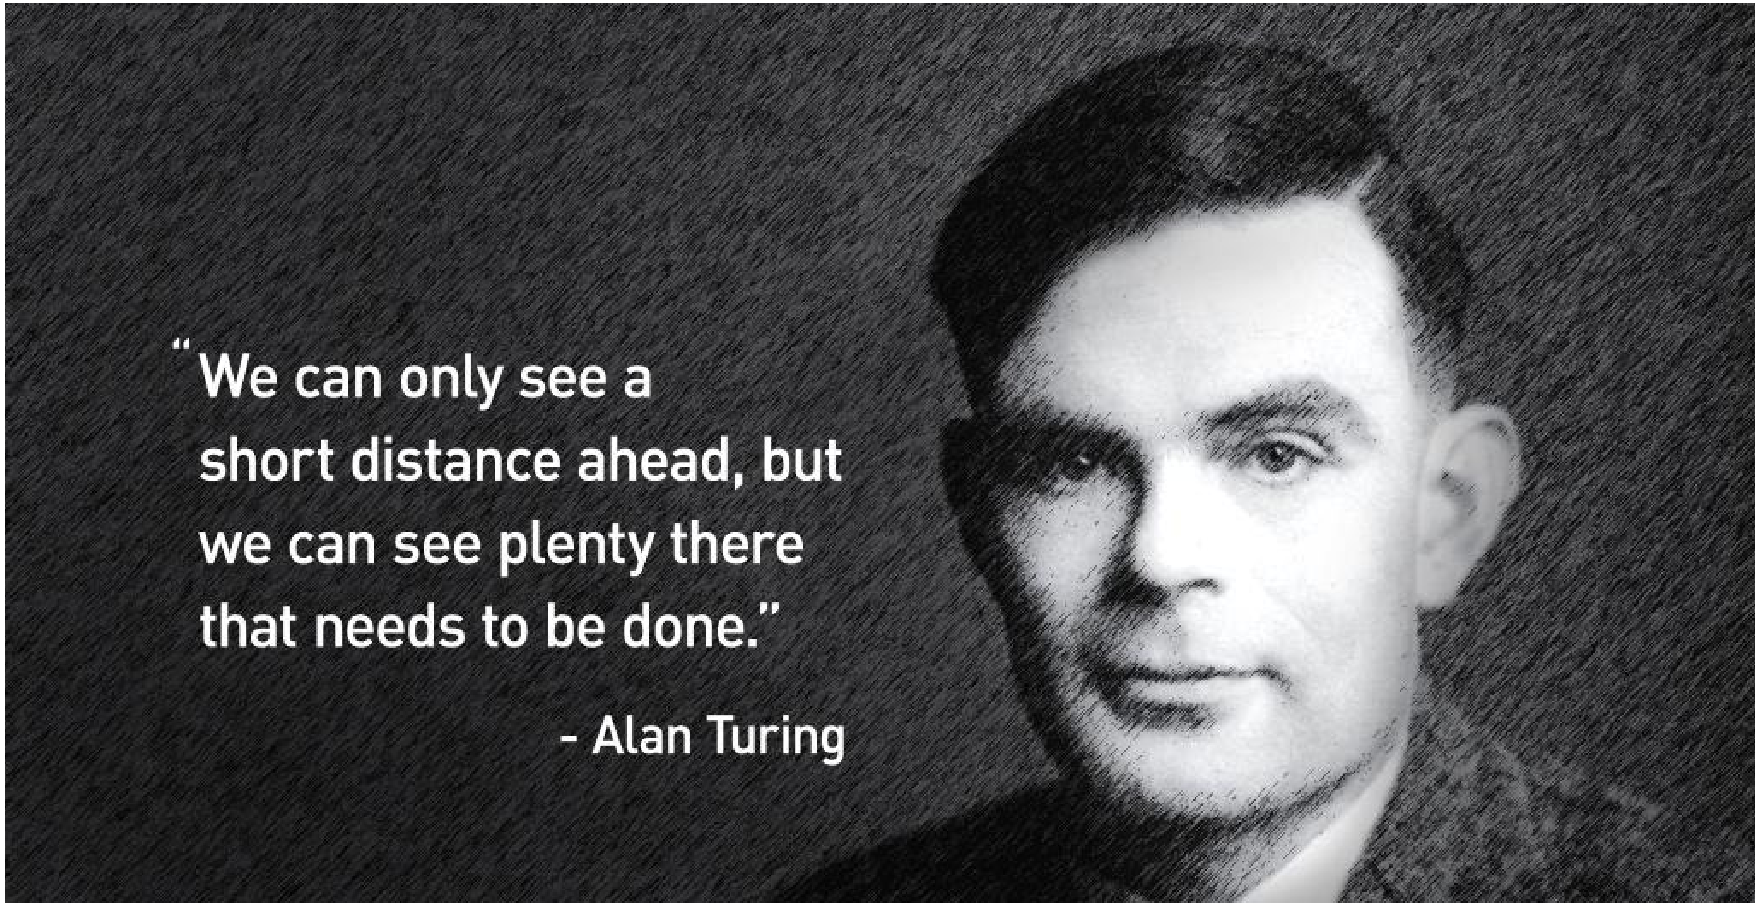

https://twitter.com/nvidia/status/1010545517405835264

G. MARCUS' CRITICISM OF DL (JAN 2018)

1. Data hungry
2. No "deep" abstraction, sensitive to noise, hard to transfer
3. No natural way to deal with hierarchy
4. Struggles with open-ended inference (e.g., answer not in the data)
5. Black-box
6. Uses little prior knowledge
7. Correlation, not causation
8. Assumes stable world
9. Sensitive to adversarial examples
10. Difficult to engineer

WHEN IN DOUBT, ASK WHAT DOES BRAIN DO?

Perceive the world

Conceptualize

Reason about things

Imagine the future

Plan what to do

Act

Learn

Think about self

Have emotion

Be conscious

Love

Be happy and pursue happiness

Be ethical

Socialize

Then ask what computer can do.

Which ones are Turing computational?

RECURSION | BOOTSTRAPPING

The entire history of human's invention: Tool that produces tools

For past 50 years: Software that writes software

Basis for the prediction of Singularity in 2045 by Ray Kurzweil

- The Law of Accelerating Returns

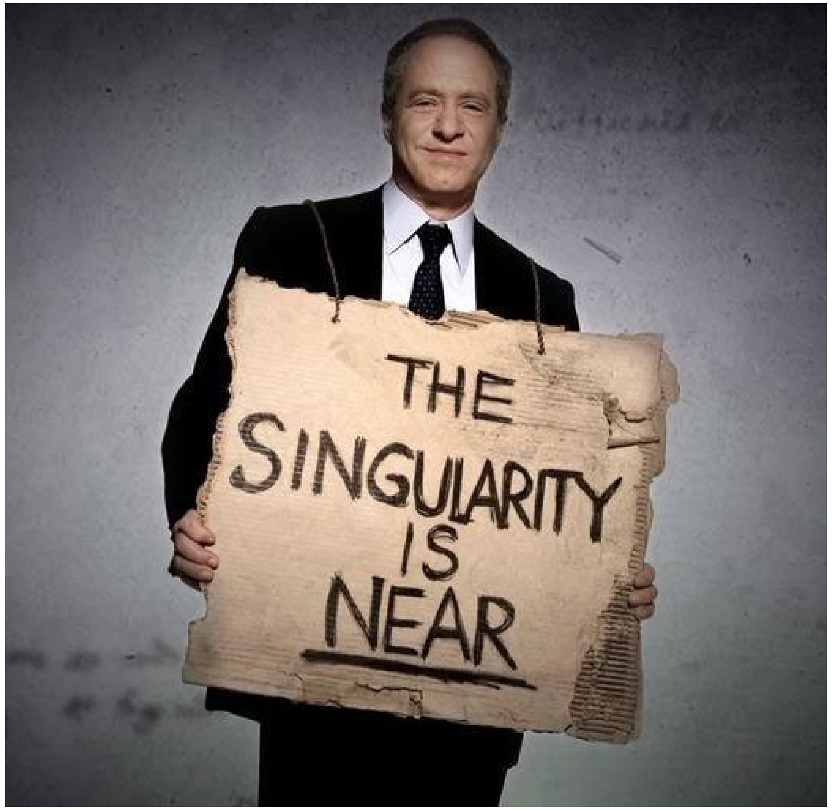

DARPA: 3 WAVES OF AI

First wave (1960s-1990s): hand-crafting rules, domain-specific, logic-based

- Can’t scale.
- Fail on unseen cases.

Second wave (1990s-2010s): machine learning, general purpose, statistics-based

- Needs lots of data
- Less adaptive
- Little explanation

Third wave (2010s-2030s): learning + reasoning, general purpose, human-like

- Requires less data
- Adapt to change
- Has contextual and common-sense reasoning
- Explainable

RODNEY BROOKS' PREDICTION

2020: The popular press starts … **the era of Deep Learning is over.**

2022: Driverless "taxi" service in a major US city, restricted settings.

2023: **A conversational agent** for **long term context**, and without recognizable and repeated patterns.

2030: An AI system with an **ongoing existence** at the level of a mouse.

2048: A robot that seems as intelligent, as attentive, and as faithful, as a dog.

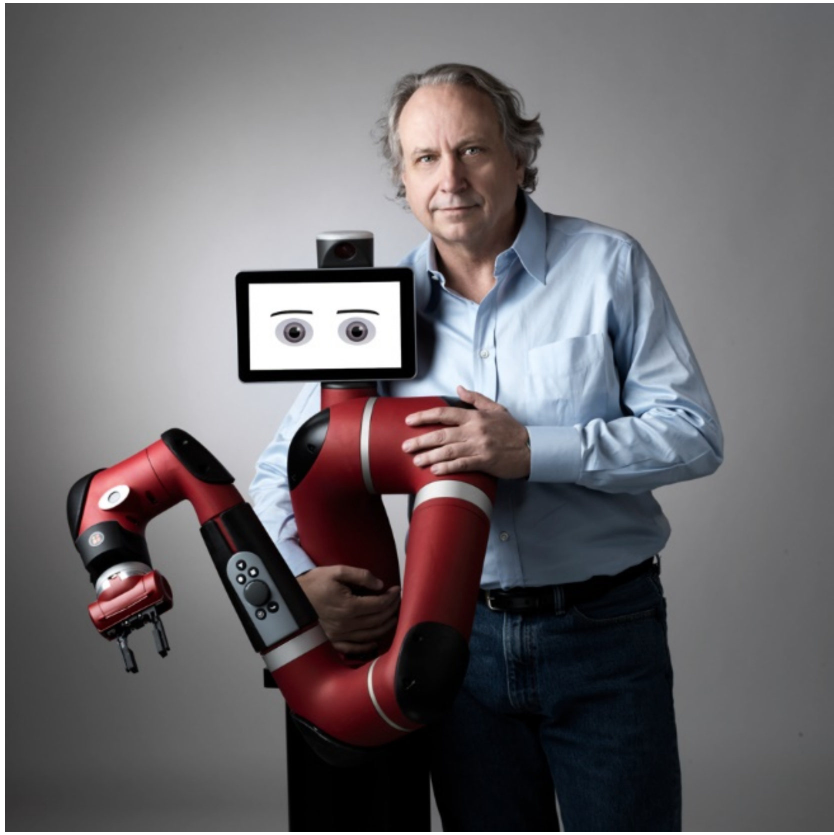

http://rodneybrooks.com/my-dated-predictions/

BY FRANCOIS CHOLLET (JULY 2017)

Models closer to general-purpose computer programs, built on top of rich primitives—this is how we will get to reasoning and abstraction.

Non-backprop: Move away from just differentiable transforms.

AutoML: Models that require less involvement from human engineers.

Meta-learning systems based on reusable and modular program subroutines.

https://blog.keras.io/the-future-of-deep-learning.html

SOFTWARE 2.0 (ANDREJ KARPATHY), NOV 2017

"Software 2.0 is written in neural network weights"

"a large portion of real-world problems […] it is significantly easier to collect the data than to explicitly write the program."

"Google is […] re-writing large chunks of itself into Software 2.0 code."

"neural networks as a software stack and not just a pretty good classifier"

"future of Software 2.0 is bright because […] when we develop AGI, it will certainly be written in Software 2.0."

https://medium.com/@karpathy/software-2-0-a64152b37c35

The Next AI Revolution (LeCun 2018)

With thanks
To
Alyosha Efros

UNSUPERVISED LEARNING

DARPA's program: Learning with less labels

Current trick: pre-training with word2vec, glove, BERT, filling the slots

In RL: intrinsic motivation

What are the principles?

- Energy-based, i.e., pull down energy of observed data, pull up every else
- Filling the missing slots (aka predictive learning, self-supervised learning)
- Multiple objectives, or no objective at all?
- Minimizing "actual" energy given a Markov blanket?
- Compress the history?
- Emergence from many simple interacting elements?

SELF-SUPERVISED LEARNING (LECUN 2018)

Predict any part of the input from any other part.

Predict the future from the recent past.

Predict the past from the present.

Predict the top from the bottom.

Predict the occluded from the visible

Pretend there is a part of the input you don't know and predict that.

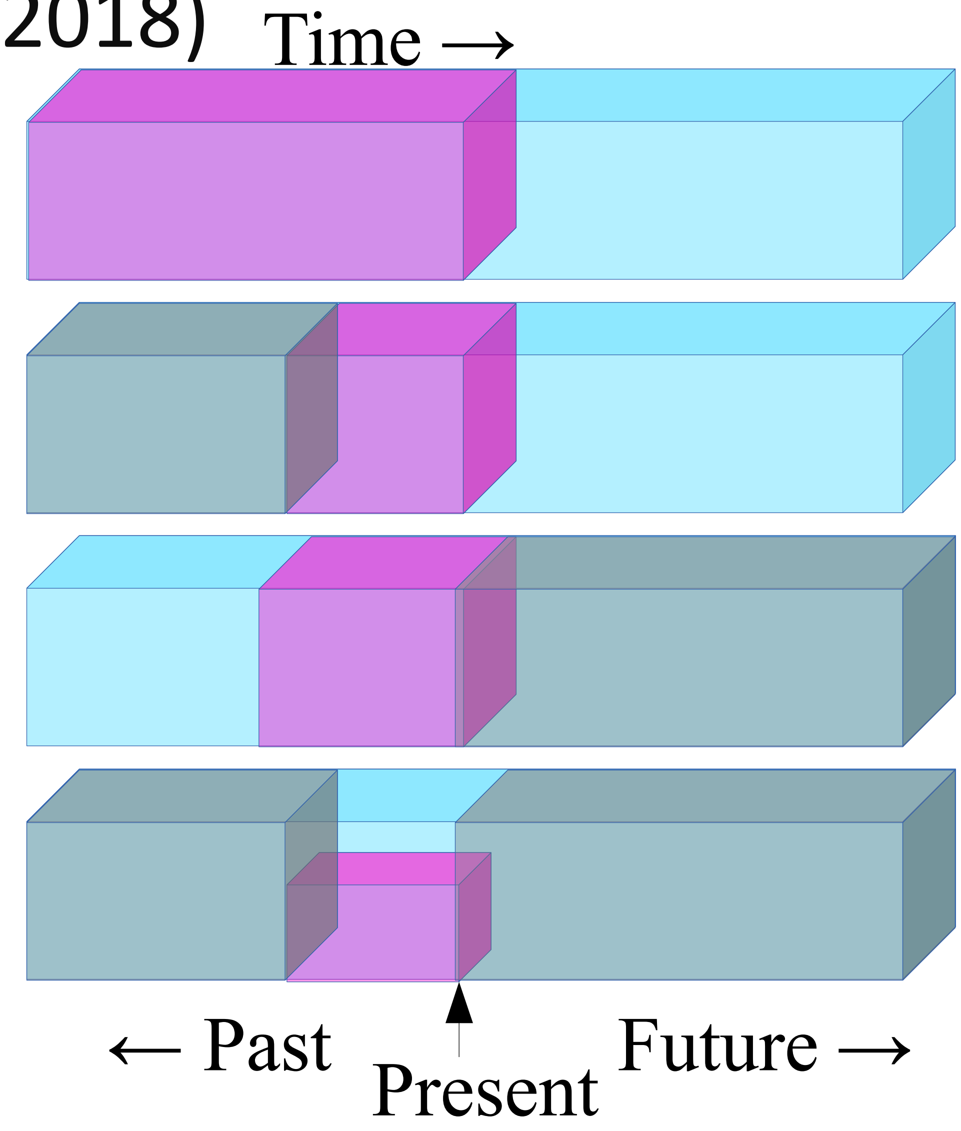

COULD SELF-SUPERVISED LEARNING LEAD TO COMMON SENSE? (LECUN 2018)

Successfully learning to predict everything from everything else would result in the accumulation of lots of background knowledge about how the world works.

Perhaps this ability to learn and use the regularities of the world is what we call common sense.

If I say "john picks up his briefcase and leaves the conference room"

You can infer a lot of facts about the scene.

John is a man, probably at work, he is extending his arm, closing his hand around the handle of his briefcase, standing up, walking towards the door. He is not flying to the door, not going through the wall....

DATA EFFICIENCY FROM DARPA

Program on Learning with Less Labels

- "reducing the amount of labeled data required to build a model by **six or more orders of magnitude**, and by reducing the amount of data needed to adapt models to new environments to **tens to hundreds of labeled examples**."

Technical Area 1: Learn and Adapt Efficiently

- "given a dataset, algorithms must be able to automatically determine which exemplars they would like to have labeled, select from existing corpora or existing models for potential transfer, and create models of a given task without human intervention. Algorithms can create data as part of this process, but they cannot manually create labels. "

Technical Area 2: Limits of Machine Learning and Adaptation

- "seeks theories that prove tight bounds on learning in the presence of transfer and meta-transfer learning. The scope of this TA includes extensions to PAC (Probably Approximately Correct) learning theory (and variants) or alternative formalisms to prove tight class-specific problem bounds (e.g., extensions to VC theory), and statistical theory needed to characterize data complexity (e.g., extensions to Johnson-Lindenstrauss) and domain mismatch."

Source: https://www.darpa.mil/program/learning-with-less-labels

“DEEP” REASONING

Fusion of perception (CV), communication (NLP), action (RL), and reasoning (classical AI)

Aka neural-symbolic integration

- Perception done by neural nets
- Symbolic integration needs reinforcement learning

http://www.neural-symbolic.org/CoCoSym2018/index.html

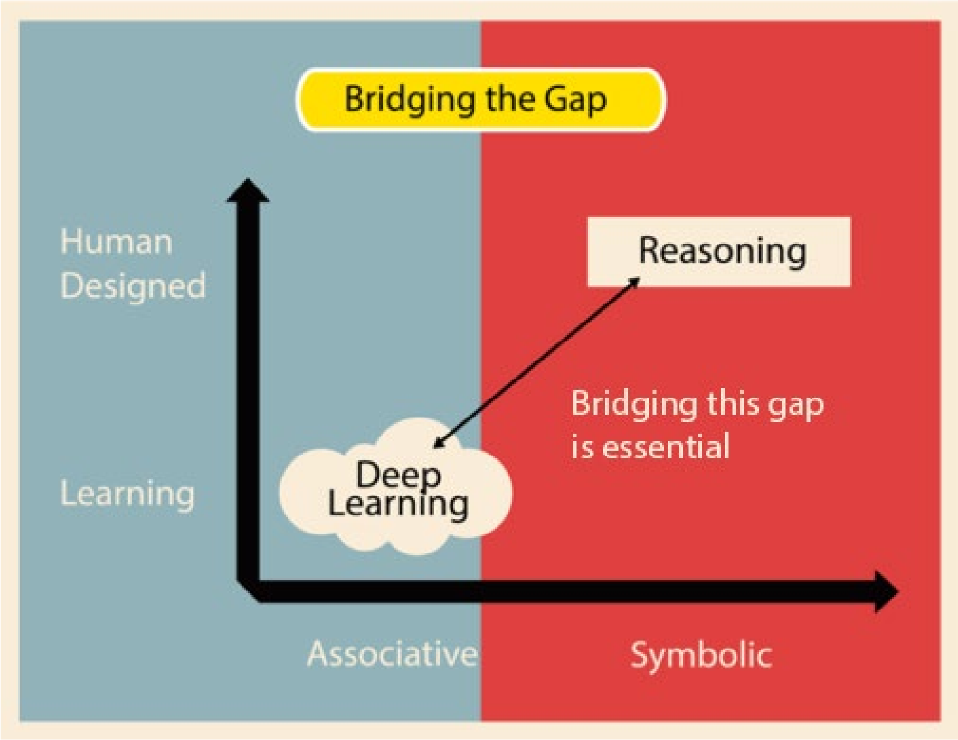

(DEEP) REINFORCEMENT LEARNING

Rewards engineering → Rewards learning

Engineering more learning signals

Model-based learning

Better priors

Model reuse

Uncertainty modelling

Hierarchical RL

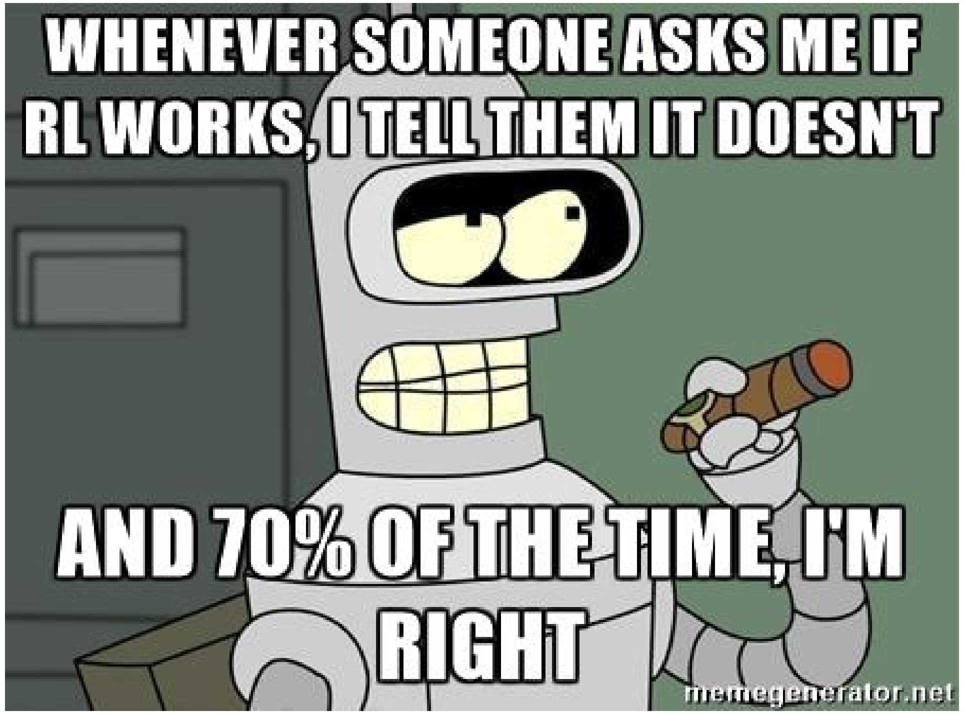

https://www.alexirpan.com/2018/02/14/rl-hard.html

THEORETICAL DL

Deep versus shallow, (over)-representation, generalization issues

- Role of compositionality on AI and learning.

Loss surface of deep nets, and whether SGD can find good local minima or global minima

Invariance and equivariance properties of CNNs

Limits of deep nets, especially RNNs, Graph NNs & neural Turing machines

Compressibility of RNNs & NTMs

The rise of Optimal Transport Theory (e.g., with GAN)

Complex-valued neural nets

https://stats385.github.io/lecture_slides
https://sites.google.com/site/deeplearningtheory/schedule
http://dalimeeting.org/dali2018/workshopTheoryDL.html
http://nips2018dltheory.rice.edu

PROBABILISTIC PROGRAMMING

Whether the brain is probabilistic/Bayesian is debatable, but probabilistic thinking is useful.

- This was the reason behind the conference of UAI and later AISTATS

Related to Probabilistic Graphical Models (Bayesian nets, Markov random fields and Factor graphs), but more flexible.

- Naturally support probabilistic reasoning (aka inference, either marginalization or posterior estimation)

http://probabilistic-programming.org/wiki/Home

ANY HOPE FOR LAWS OF INTELLIGENCE?

Be predictive (aka Popper's falsifiability)

Beauty (simple and revealing)

Elements of surprise

Universality

Invariance

Symmetry

Conservation of quantities

A discovery of an anti-thermal dynamics law

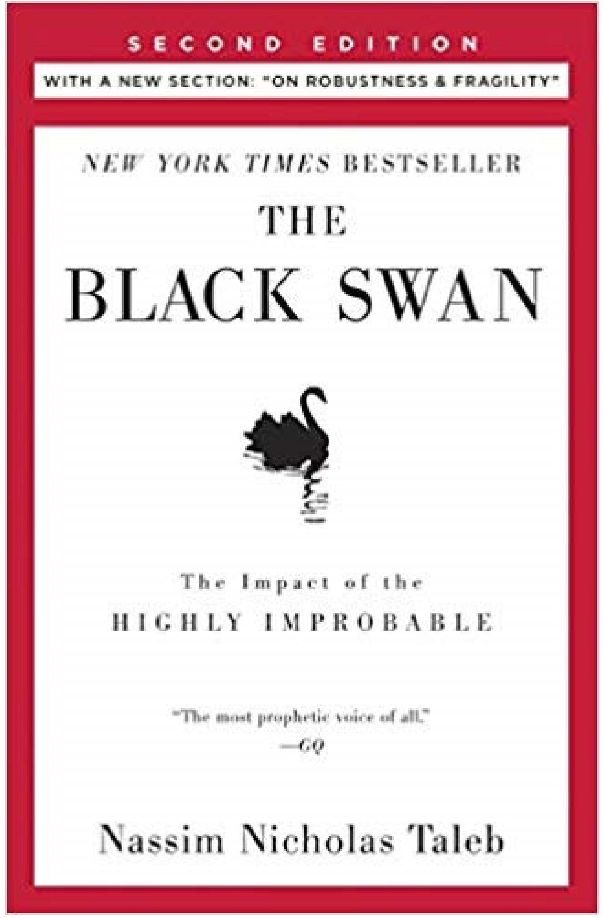

THEORIES OF INTELLIGENCE

Penrose's consciousness theory

Friston's free-energy principle

Russell's bounded rationality

Domingos' Master algorithms

Newell's unified theories of cognition

Cognitive architecture. The latest: Sigma (symbolic + factor graph)

Priors of intelligence (innate)

Intelligence as emerging phenomenon

Minsky's society of minds

Quantum theories of cognition

Agenda

- Refresher
- 2018-2019 Review
- Outlook
- **Empirical AI Research**
- DAIR: Deakin AI Research

Empirical observations of empirical AI research

Concepts • Positioning • Planning • Empirical methods • Theorization

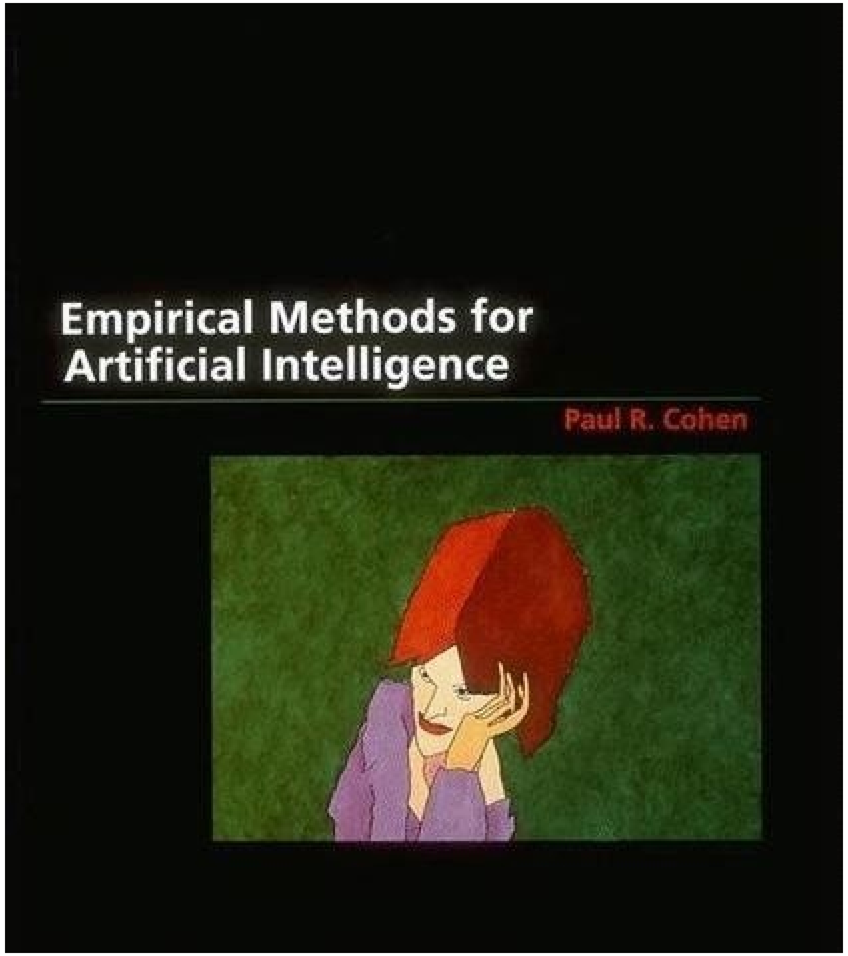

em·pir·i·cal

/əmˈpirik(ə)l/

Origin: Greek *empeirikos* ("experienced")

adjective

based on, concerned with, or verifiable by observation or experience rather than theory or pure logic.
"they provided considerable empirical evidence to support their argument"

Empirical != Ad-hoc

Empirical !=
Applied research

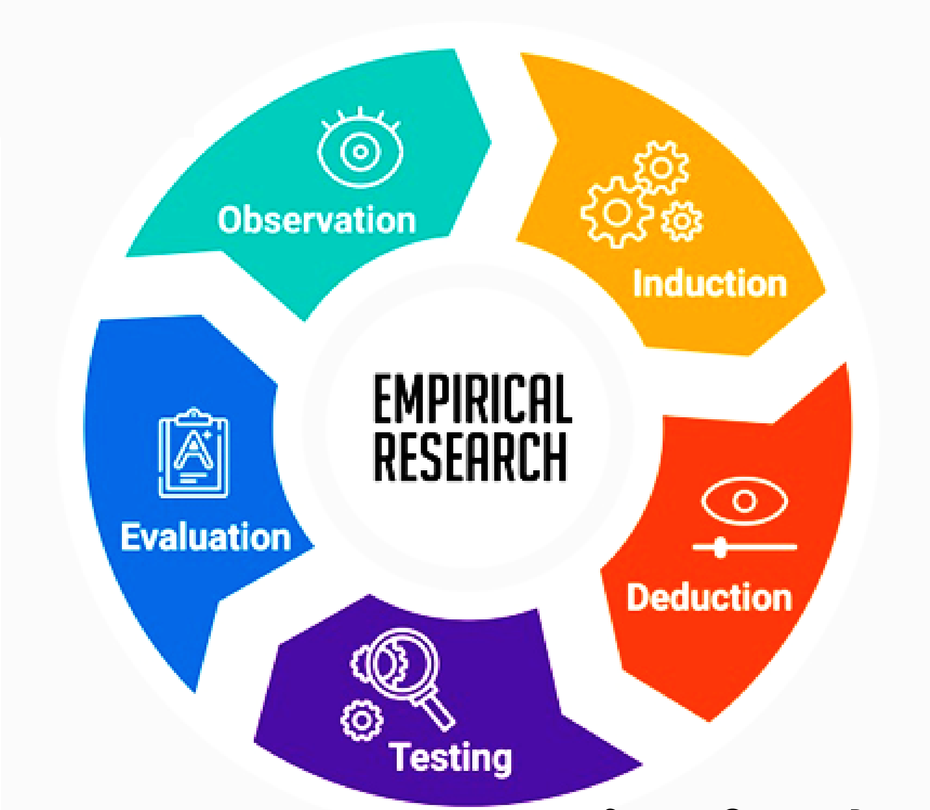

Source: QuestionPro

HEAVIER-THAN-AIR FLIGHT: OBSERVATION

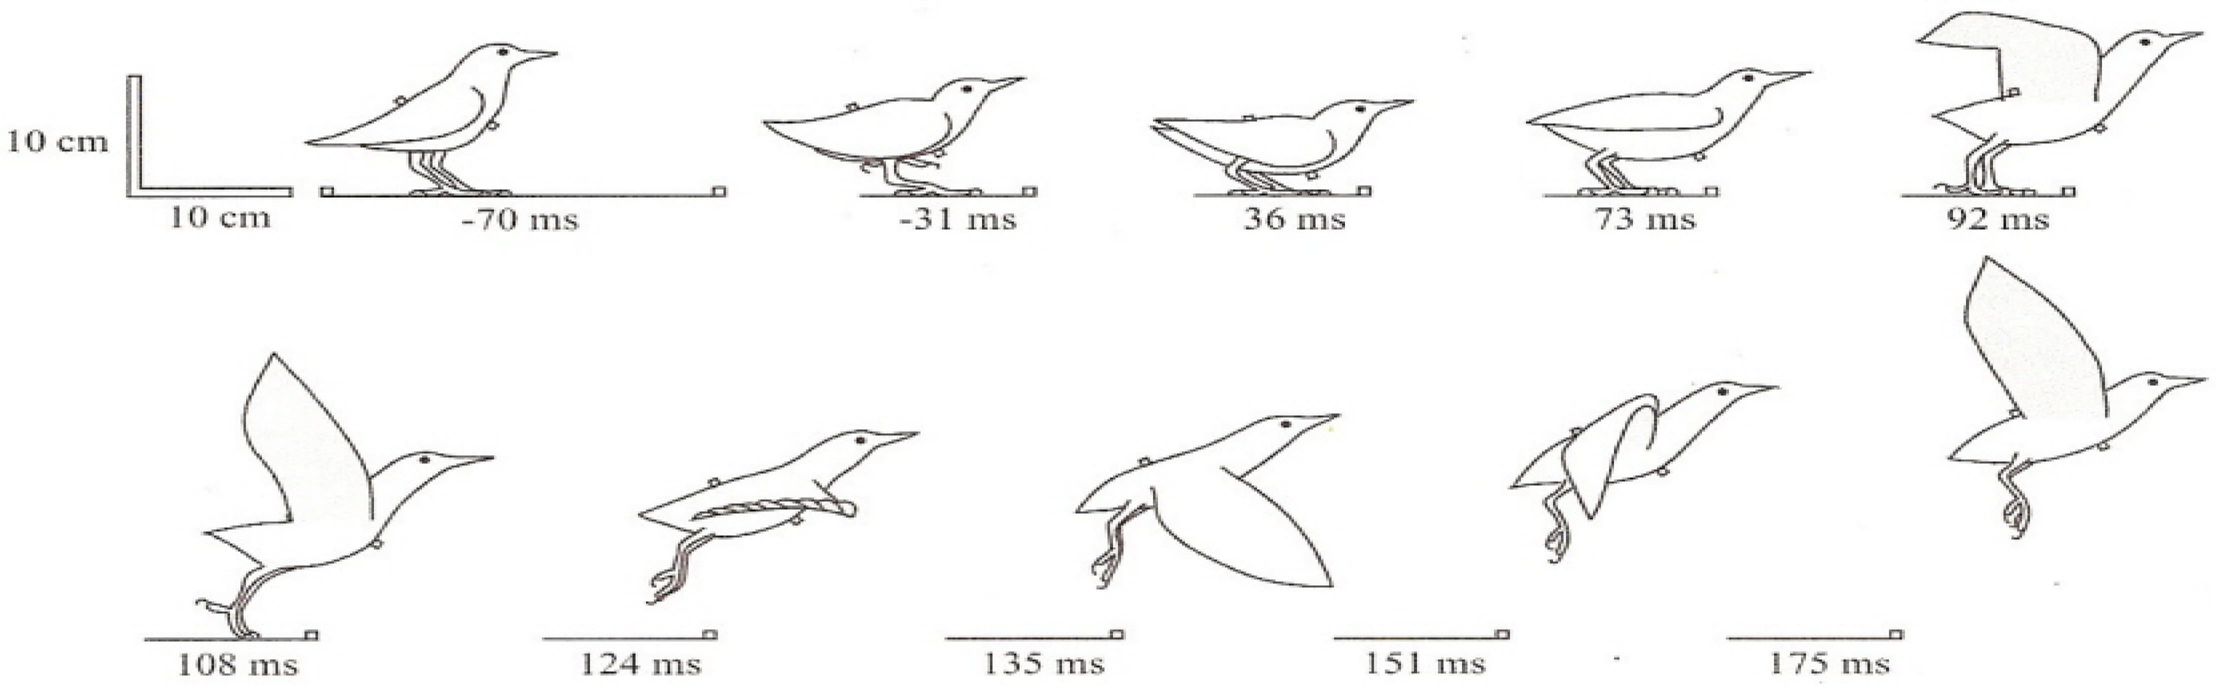

http://people.eku.edu/ritchisong/554notes2.html

HEAVIER-THAN-AIR FLIGHT: THEORY

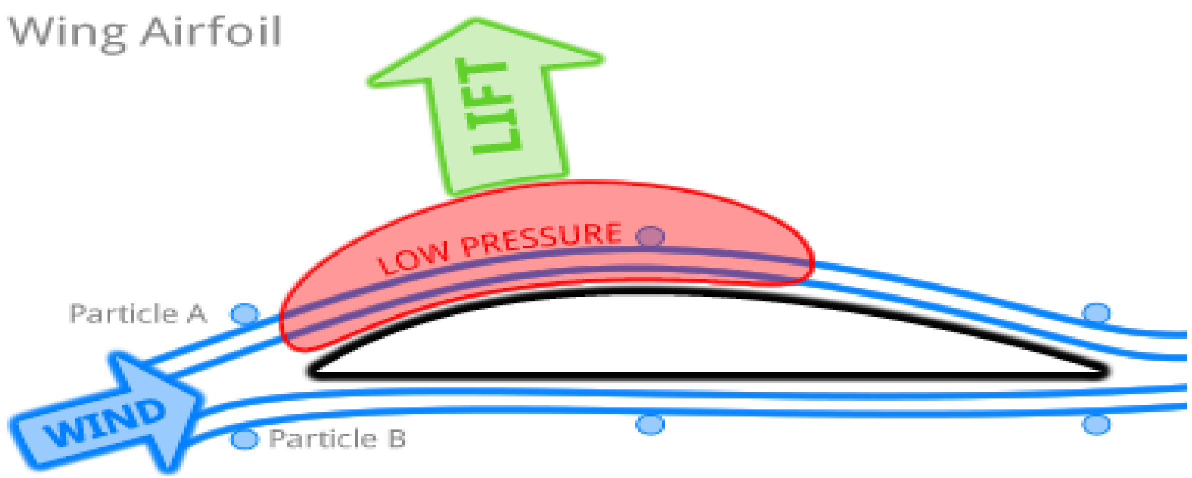

Enabling factors

- Aerodynamics
- Powerful engines
- Light materials
- Advances in control
- Established safety practices

Sources:
http://aero.konelek.com/aerodynamics/aerodynamic-analysis-and-design
http://www.foolishsailor.com/Sail-Trim-For-Cruisers-work-in-progress/Sail-Aerodynamics.html

WHY EMPIRICAL?

- Unless theoretically proven, we need realistic validation.
  - Physics theories are great, but they are based on assumptions → need experimental validation
- Toy theories don’t scale → have zero impact → Winter AI.
  - Hence, the theories have gone down a bit, e.g., down the Chomsky hierarchy in NLP, and down to surface pattern recognition in CV.
- Real difficult problems lead to new theoretical set ups.
- Empirical observations offer insights, suggests problem formulation, induction and abduction
- Push the field forward. AI now strongly driven by industry.

EMPIRICAL RESEARCH != IGNORANCE OF THEORY

Theory guides the search. It quantifies the goals. Making sure what we obtained is correct.

Current AI theories:

- Optimization. Mostly non-convex. Loss landscape of deep nets.
- Learnability and generalization. Learning dynamics, convergence and stability.
- Probabilistic inference. Quantification of uncertainty. Algorithmic assurance.
- First-order logic.
- Constraint satisfaction.

IS EMPIRICAL ML/AI DIFFICULT?

Yes! This is the real world. It is messy and must be difficult.

- Robotics
- Healthcare
- Drug design

No! The "effective" solution space is small and constrained.

- Machine translation
- Self-driving cars
- Drug design

TWO RESEARCH STYLES

Methodological: One method, many application scenarios. The method will fail at some point.

Applied: One app, many competing methods. Some methods are intrinsically better than others. E.g., CNN for image processing.

HOW TO POSITION YOURSELF

"[…] the dynamics of the game will evolve. In the long run, the right way of playing football is to position yourself intelligently and to wait for the ball to come to you. You'll need to run up and down a bit, either to respond to how the play is evolving or to get out of the way of the scrum when it looks like it might flatten you." (*Neil Lawrence, 7/2015, now with Amazon*)

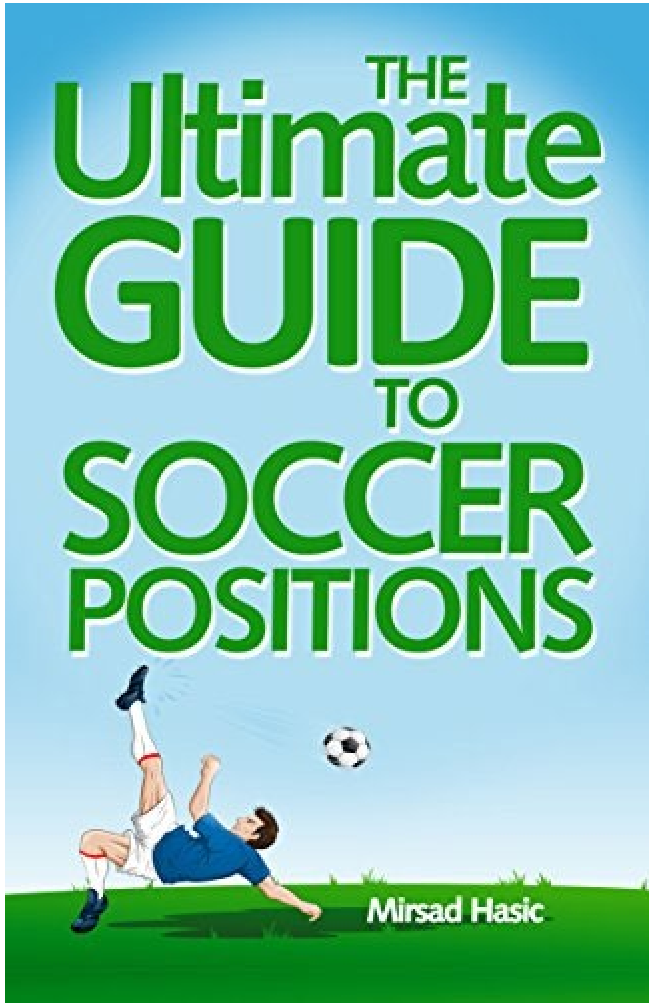

http://inverseprobability.com/2015/07/12/Thoughts-on-ICML-2015/

"Then ask yourself, what is your unique position? What are your strengths and advantages that people do not have? Can you move faster than others? It may be by having access to data, access to expertise in the neighborhood, or borrowing angles outside the field. Sometimes digging up old ideas is highly beneficial, too.

Alternatively, just calm down, and do boring-but-important stuffs. Important problems are like the goal areas in ball games. The ball will surely come."

https://letdataspeak.blogspot.com/2016/12/making-dent-in-machine-learning-or-how.html

WISDOM OF THE ELDERS

Max Planck

- Science advanced by one death at a time

Geoff Hinton:

- When you have a cool idea but everyone says it is rubbish, then it is a really good idea.
- The next breakthrough will be made by some grad students

Richard Hamming:

- After 7 years you will run out of ideas. Better change to a totally new field and restart.
- If you solve one problem, another will come. Better to solve all problems for next year, even if we don’t know what they are.

http://paulgraham.com/hamming.html

STRATEGIES TO SELECT RESEARCH AREAS

Ask what intelligence means, and what human brain would do

Set an intelligence level (bacteria, insect, fish, mouse, dog, monkey, human)

Check solvable unsolved problems. There are LOTS of them!!!

Look for impact

Look for being difference. Fill the gap. Be unique. Be recognizable.

Riding the early wave.

Ask what a physicist would do:

- Simplest laws/tricks that would do 99% of the job
- Minimizing some free-energy
- Looking for invariance, symmetry

INSIGHTS FROM NEUROSCIENCE

Brain is the only known intelligence system

Can provide inspirations for AI

Can validate AI claims

Indeed, many AI concepts have had originated from neuroscience

- However, the influence is only at the initial step
- AI doesn’t have the constraints that neuroscience does
- AI can do things that are impossible within neuroscience

Hassabis, Demis, et al. "Neuroscience-inspired artificial intelligence." *Neuron*95.2 (2017): 245-258.
Schmidhuber, Jürgen. "Deep learning in neural networks: An overview." *Neural networks* 61 (2015): 85-117.

NEUROSCIENCE CONCEPTS

Memory

Reinforcement learning

Control

Attention

Planning

Imagination

Symbolic manipulation

Emotion

Consciousness

Transfer learning

Continual learning

Spatial memory

Social interaction

Developmental

Catastrophic forgetting

Nature versus nurture

Newell's Time Scale of Human Action

Scale (sec)	Time Units	System	Band
10^7	months		
10^6	weeks		Social
10^5	days		
10^4	hours	Task	
10^3	10 min	Task	Rational
10^2	minutes	Task	
10^1	10 sec	Unit task	
10^0	1 sec	Operations	Cognitive
10^{-1}	100 ms	Deliberate act	
10^{-2}	10 ms	Neural Circuit	
10^{-3}	1 ms	Neuron	Biological
10^{-4}	100 µs	Organelle	

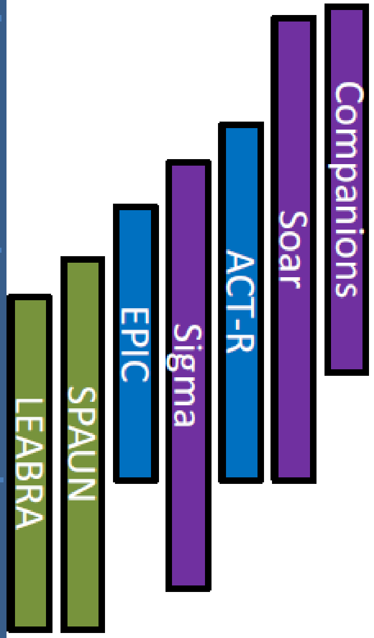

#REF: A Cognitive Architecture Approach to Interactive Task Learning, John E. Laird, University of Michigan

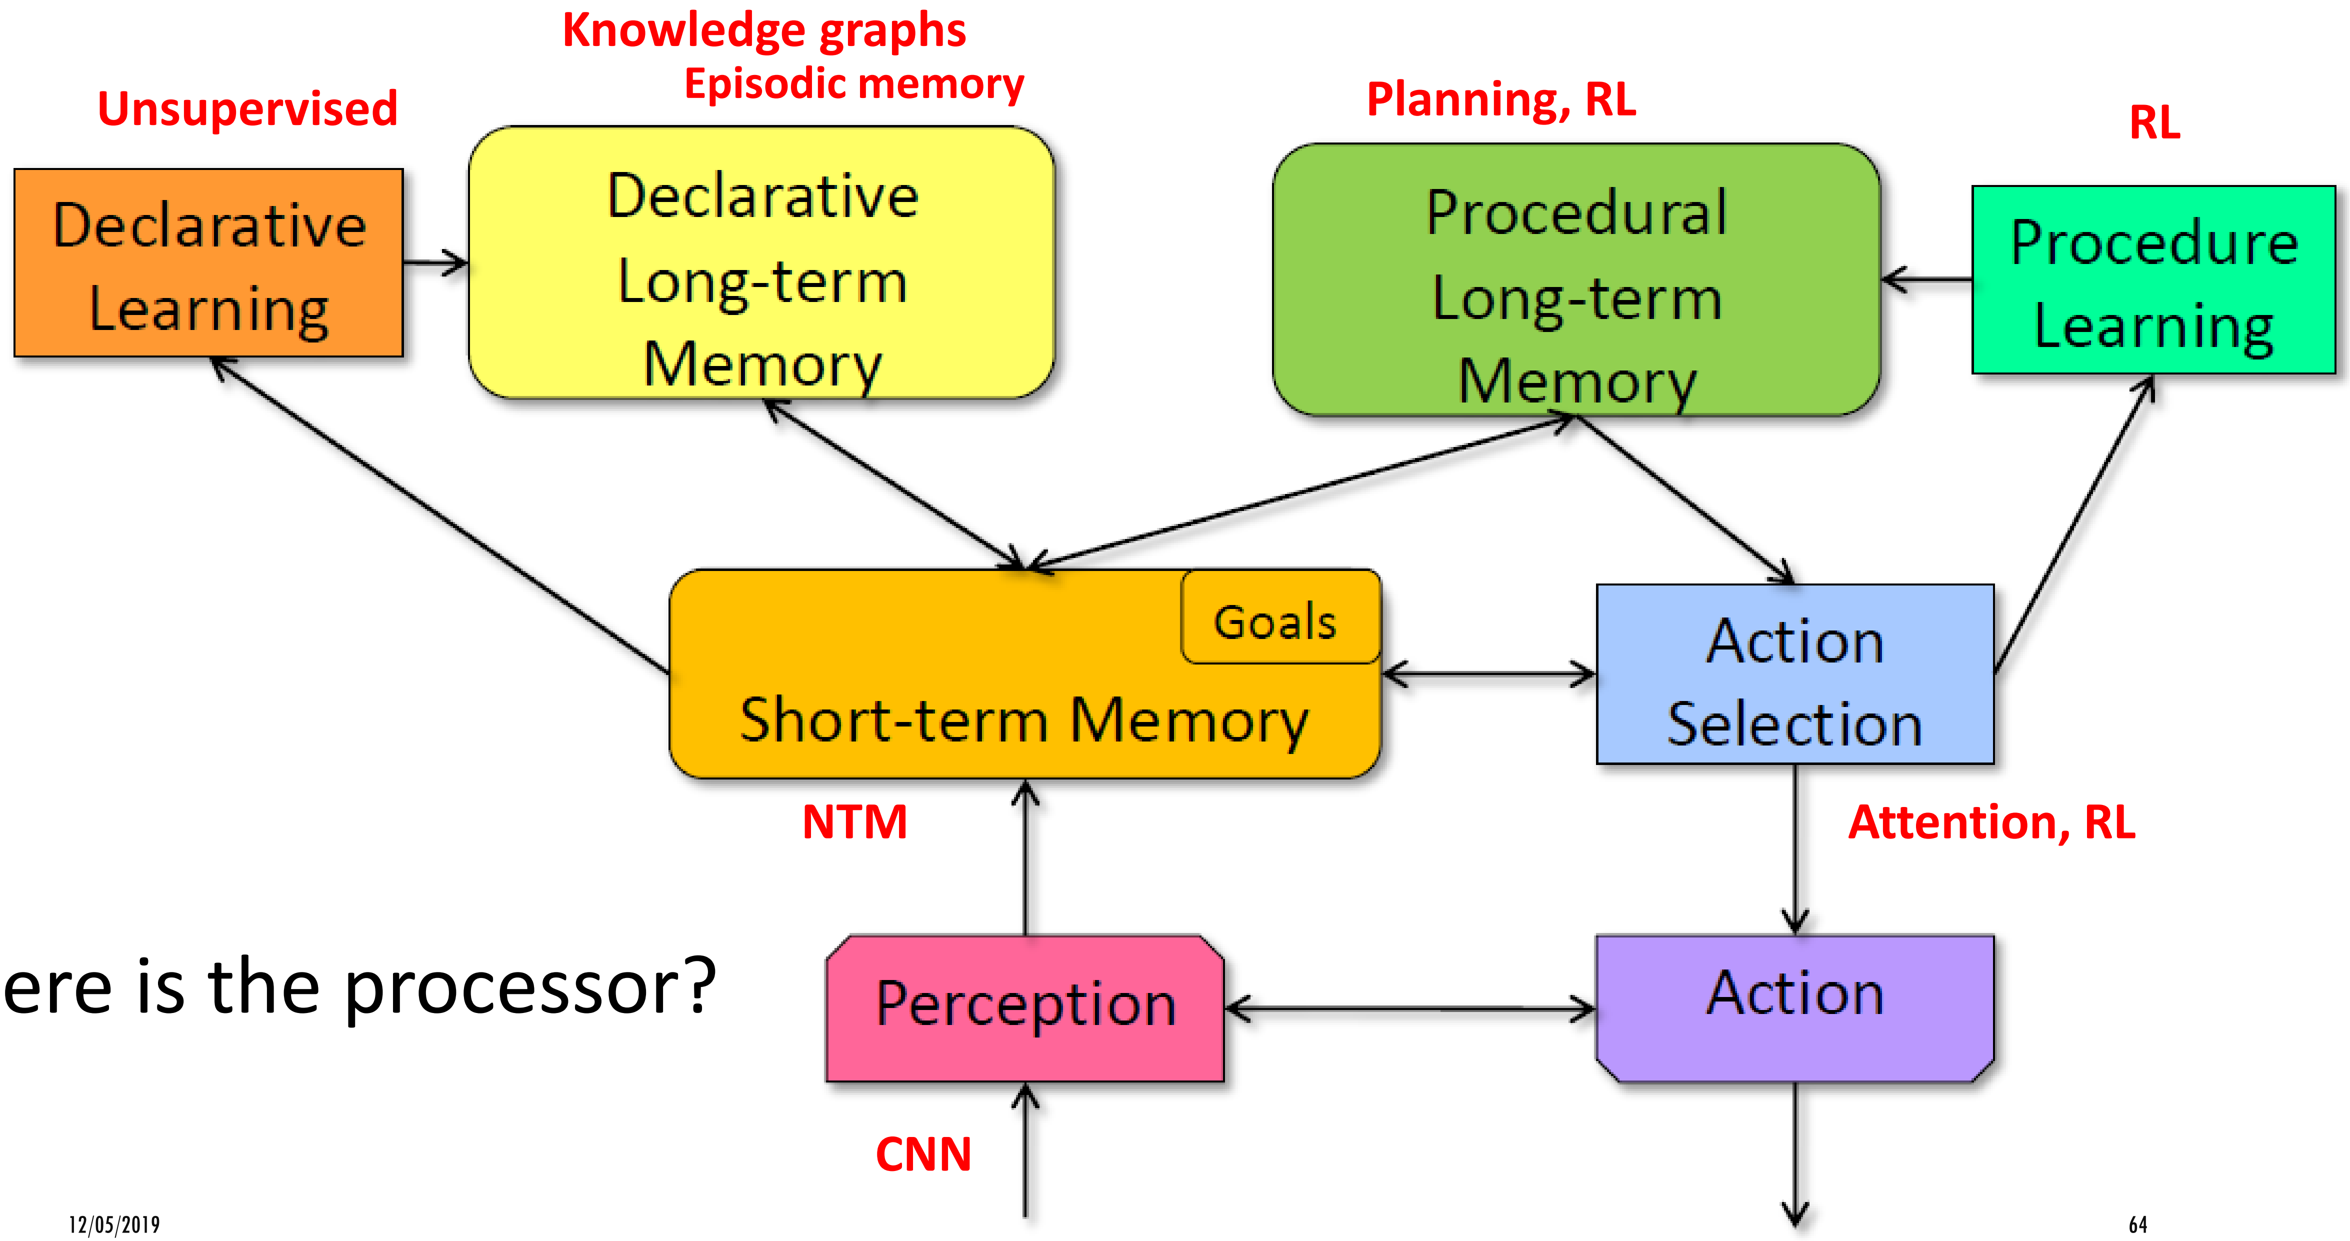

Where is the processor?

THINKING & LEARNING, FAST AND SLOW

Thinking

- Fast: reactive, primitive, associative, animal like
- Slow: deliberative, reasoning, logical.

Learning

- Fast: often memory-based, one-shot learning style, fast-weights, e.g., k-NN
- Slow: often model- based, iterative, slow-weights, e.g., backgrop

UNSOLVED PROBLEMS (1)

Learn to organize and remember ultra-long sequences

Learn to generate arbitrary objects, with zero supports

Reasoning about object, relation, causality, self and other agents

Imagine scenarios, act on the world and learn from the feedbacks

Continual learning, never-ending, across tasks, domains, representations

Learn by socializing

Learn just by observing and self-prediction

Organizing and reasoning about (common-sense) knowledge

Grounding symbols to sensor signals

Automated discovery of physical laws

Solve genetics, neuroscience and healthcare

Automate physical sciences

Automate software engineering

UNSOLVED PROBLEMS (2)

[Click here]

LECUN'S LIST (2018)

Deep Learning on new domains (beyond multi-dimensional arrays)

Graphs, structured data...

Marrying deep learning and (logical) reasoning

Replacing symbols by vectors and logic by algebra **Self-supervised learning of world models**

Dealing with uncertainty in high-dimensional continuous spaces

Learning hierarchical representations of control space

Instantiating complex/abstract action plans into simpler ones

Theory!

Compilers for differentiable programming.

A TYPICAL EMPIRICAL PROCESS

Hypothesis generation

- Mostly by our brain
- But AI can help the process too!

Hypothesis testing by experiments

Two types:

- **Discovery**: it is there, awaiting to be found. Mostly theories of nature.
- **Invention**: not there, but can be made. Mostly engineering.

HOW TO PLAN AN EMPIRICAL RESEARCH

Get the right people. Smart. Self-motivated. Can read maths. Can code. Can do dirty data cleaning.

Choose a right topic. Impactful ones at the right time. Looks for current failures. Set the targets.

Explore theoretical options.

Breaking into solvable parts

- Before: 3 years of PhD, one problem.
- Now: every 3-6 months to meet a conference deadline (Feb-March, May-June, Aug-Sept, Nov-Dec)

Collect data. LOTS of them. Clean up.

Code it up. In DL, it means using your favourite programming frameworks (e.g., PyTorch, TensorFlow)

Evaluate the options on some small data, then big data, then a variety of data

Repeat

Question: When should we stop evaluating?

CHOOSING HIGH IMPACT PROBLEMS

Impact to the population (e.g., 100M people or more – Andrew Ng)

Start a new paradigm – a new trend

Ride the wave

No need to have a competitive mind – “Mine beats yours”

CHOOSING IMPORTANT BUT LESS SHINY PROBLEMS

Computer vision & NLP are currently shinny areas. Dominated by industry. You can't compete well without good data, computing resources, incentives and high concentration of talent.

Some of the best lie in the intersection of areas. Esp. those with some domain knowledge (which you can learn fast).

- Biomedicine, e.g., drug design, genomics, clinical data.
- Sustainability, e.g., Earth health, ocean, crop management, land use, climate changes, energy, waste management.
- Social good, law, developing countries, access to healthcare and education, clean water, transportation.
- Value-alignment, ethics, safety, assurance.
- Accelerating sciences, e.g., molecule design/exploration, inverse materials design, quantum worlds.

HOW TO BE PRODUCTIVE AND BE HAPPY

Maximize your quantity and quality of enlightenment.

Fail fast. Fail smart. Start small. Scale latter.

Use simplest possible method to confirm or reject a hypothesis

If the hypothesis is partially confirmed, refine the hypothesis, or improve the method, and repeat.

Time spent per experiment should be limited to an hour or similar range. Winning code seems to run for 2 weeks.

Productivity = knowledge unit gained / time unit.

PATTERNS REPEATED EVERYWHERE

Lots of problems can be formulated in similar ways

We often need an intermediate representation, e.g., graphs that can be seen in many fields.

Solving the generic problems are often faster and generalizable. See Sutton’s remarks.

Many domains can be solved at the same time. After all, we have just one brain that does everything.

Methods that survive need to be simple, generic and scalable. PCA, CNN, RNN, SGD are examples.

METHODS VERSUS CONCEPTS

Methods come and go, but concepts stay

- Decision trees, kernels → Template matching, data geometry
- Random forests, gradient boosting → Smoothing, variance reduction
- Neural nets → Differentiable programming, function reuse, credit assignment

“Machine learning is the new algorithms”

- https://nlpers.blogspot.com.au/2014/10/machine-learning-is-new-algorithms.html

RIDING THE EARLY WAVES

Neural network, cycle 1: 1985-1995
- Backprop

SVM:1995-2005
- Nice theory, strong results for moderate data

Ensemble methods: 1995-2005
- Nice theory, strong results for moderate data

Probabilistic graphical models: 1992-2002
- Nice theory, open up new thinking, difficult in practice

Topic models & Bayesian non-parametrics: 2002-2012
- Nice theory, indirect applications

Feature learning: 2005-2015
- Little theory, toy results, great excitements

Neural network, cycle 2: 2012-present
- Differentiable programming, backprop + adaptive SGD
- Little theory, strong results for large data

For PhD students: be an expert when the trend is at its peak!

MY OWN JOURNEY

In 2004 I bet my PhD thesis on "Conditional random fields". It died when I finished in 2008.

In 2007 I bet my theoretical research area on Deep Learning, back then it was Restricted Boltzmann Machines. RBM died around 2012. But the field moves on and now reaches its peak.

In 2012 I bet my applied research area on Biomedicine. This is super-hot now.

In 2019 I am betting on "learning to reason" → cognitive architectures.

THE 10 YEAR MINI-CYCLES IN ML

Neural network, cycle 1: 1985-1995

- Backprop

SVM:1995-2005

- Nice theory, strong results for moderate data

Ensemble methods: 1995-2005

- Nice theory, strong results for moderate data

Probabilistic graphical models: 1992-2002

- Nice theory, open up new thinking, difficult in practice

Topic models & Bayesian non-parametrics: 2002-2012

- Nice theory, indirect applications

Feature learning: 2005-2015

- Little theory, toy results, great excitements

Neural network, cycle 2: 2012-present

- Differentiable programming, backprop + adaptive SGD
- Little theory, strong results for large data

THE CURSE OF TEXT BOOKS

Truyen's law: "whenever a universally accepted textbook is published, the field is saturated."

- decision tree (1984),
- neural nets, cycle 1 (1995),
- information retrieval (1998),
- reinforcement learning (1998),
- multi-agent system (2001),
- ensemble methods (2001),
- kernels (2002),
- probabilistic graphical models (2006),
- structured output learning (2007),
- learning to rank (2009),
- Bayesian non-parametrics (2012), and
- deep learning (2016).

HOW MUCH COMPUTING POWER?

The larger the better. It gives you freedom.

- Sutton ("the bitter lesson"): generic methods + lots of compute win!
- See the history of speech recognition, machine translation.

However, for faster discovery and learning, better maximize the number of reasonably good GPUs → more people, more parallel experiments.

On the other hand, little compute forces us to be smart. There are lots of important problems with lean data. See "You and your research" by Richard Hamming.

Our brain uses very little energy (somewhere like 10 Watts). There must be clever ways to arrange computation.

BUT … BY ALL MEANS, ADVANCE SCIENCE

J. Pearl: Uncertainty in AI (1980s-1990s) & causality (2000s-present)

G. Hinton (Turing Award): parallel distributed processing (1980s-present)

J. Schmidthuber: Agent that learns to do, and to learn (1980s-present)

C. Sutton: Agent that learns by reinforcement (1980s-present)

M. Jordan: Probabilistic inference (1990s-present)

WARNING! PREDICTION V.S. UNDERSTANDING

We can predict well without understanding (e.g., planet/star motion Newton).

Guessing the God's many complex behaviours versus knowing his few universal laws.

→Without theorization, we can have intelligent behaviours, but not true intelligence.

→Might **abduction** (selecting the simplest explaining theory) be the goal.

A NECESSARY CAUTION

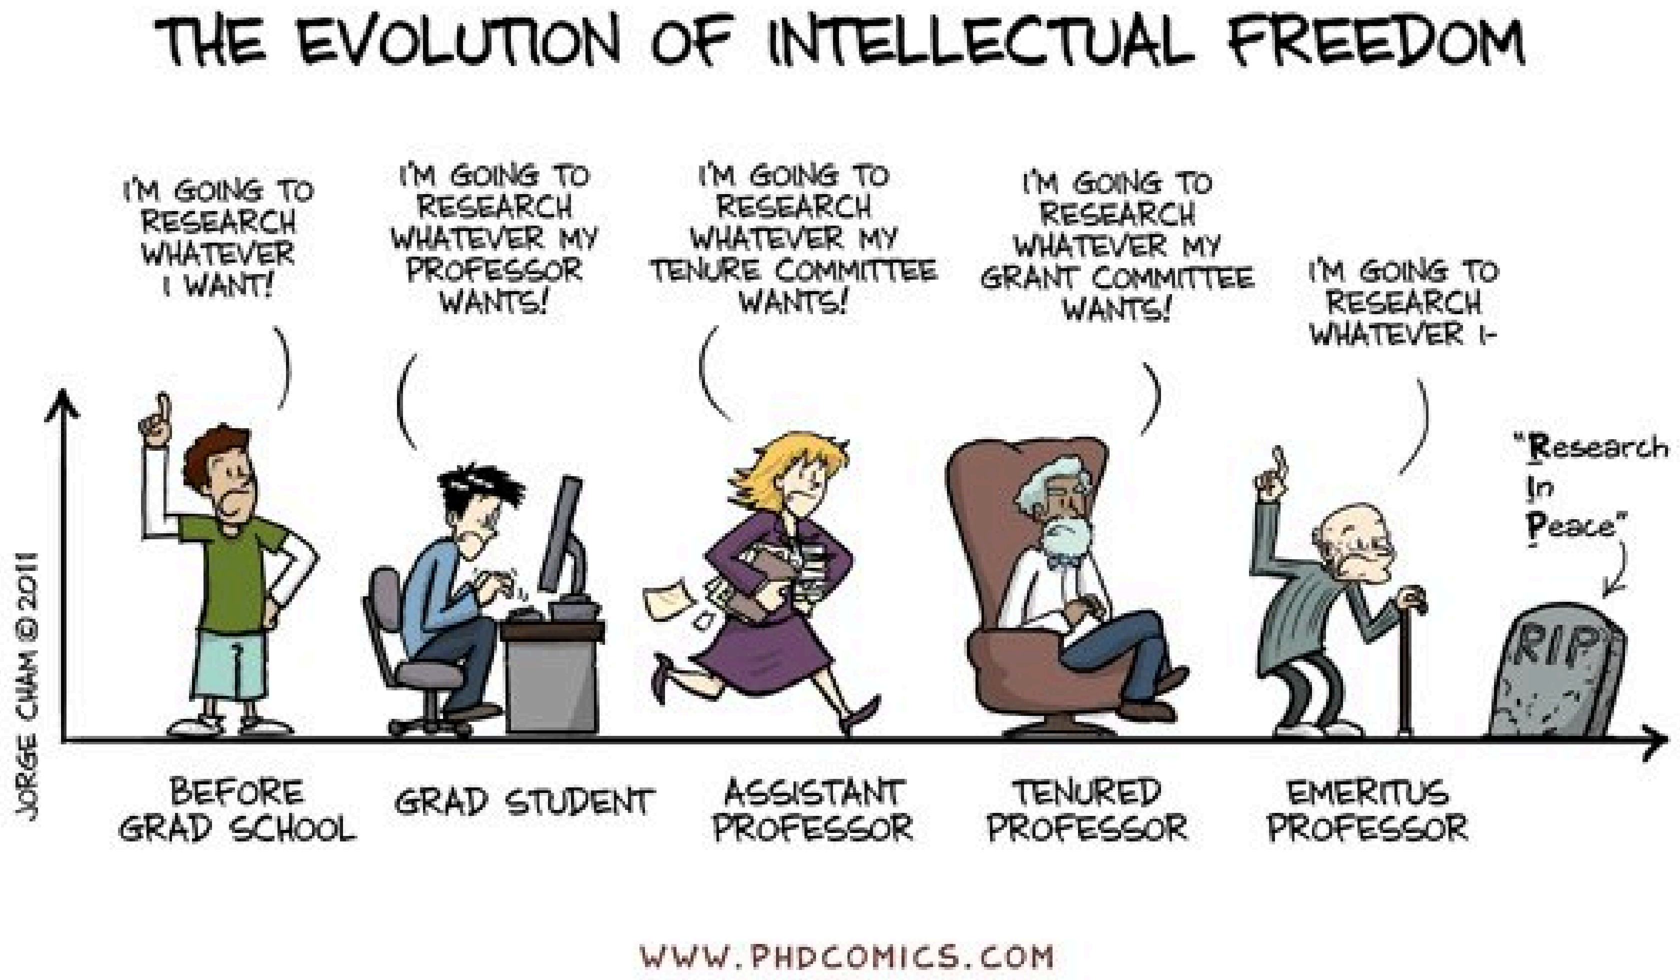

Agenda

- Refresher
- 2018-2019 Review
- Outlook
- Empirical AI Research
- **DAIR: Deakin AI Research**

DAIR: Deakin AI Research

Part of the newly found Applied AI Institute

THE TEAM @DAIR

A/Prof. Truyen Tran

Dr Vuong Le

Dr Phuoc Nguyen

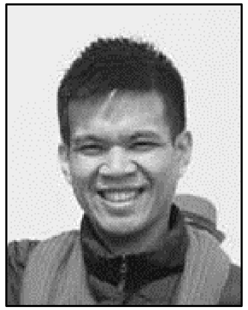

Dr Khanh Tran

Dat Tran @UTS

Romero Barata

Dung Nguyen

Tung Hoang

Tin Pham

Kien Do

Thao Le

Hung Le

Duc Nguyen

Fundamentals

- New inductive biases
- Memory and Turing machines
- Learning to reason
- Learning with less labels
- Deep reinforcement learning

Applied projects

- Vision: Visual reasoning
- NLP: QA and dialog systems
- Accelerating physical & bio sciences
- Transforming healthcare
- Automating software developments
- Building smarter smart homes
- Value-aligned ML

NEW INDUCTIVE BIASES

- Search for new priors
- In DL, we have a small set of primitive operators (signal filtering, convolution, recurrence, gating, memory and attention)
- Sets, tensors, sequences, trees, graphs.
- Inspirations from neocortex
  - Columns, thalamus routing, memory structures.

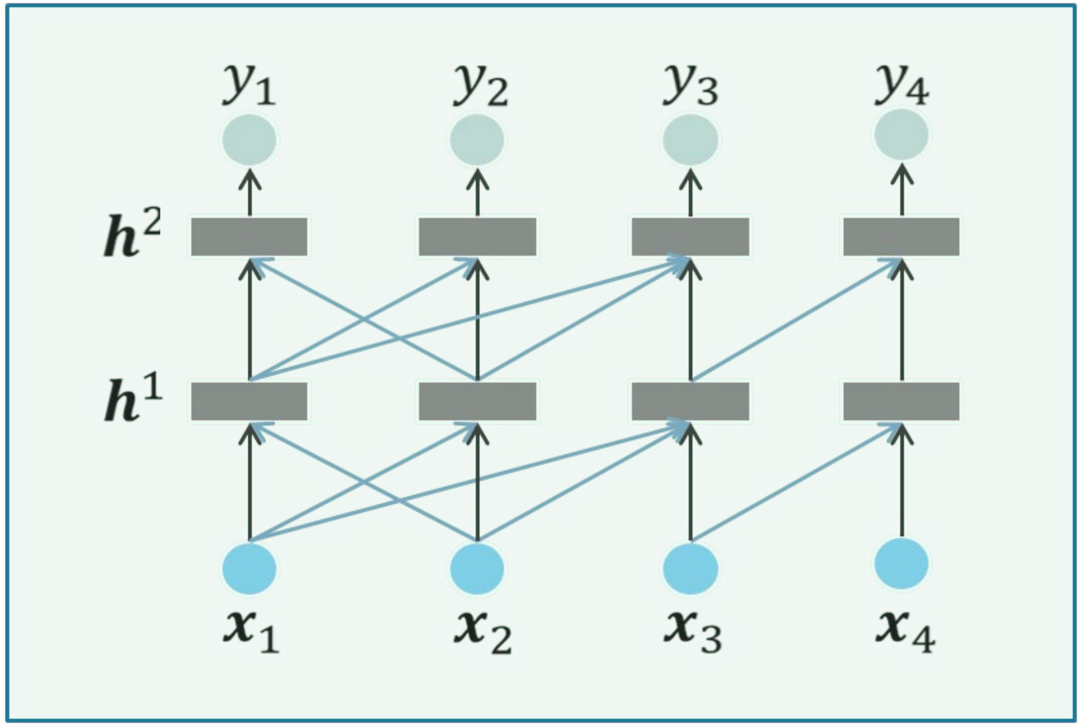

GRAPH: WHAT WE HAVE DONE

Column Network (CLN): Generic message passing architecture (AAAI'17)

Relation basis for knowledge graph completion (ICPR'18)

Graph with virtual nodes (IJCAI'17 workshop)

Graph Memory Network (ICLR'18) → Relational Dynamic Memory Network (In submission to MLJ)

- Support **graph-graph interaction** (e.g., chemical-chemical interaction)

GRAPH: WHAT WE HAVE DONE (2)

Graph transformation policy network (Submitted to ICLR’19)

Graph Attentional Multi-label Learning (GAML): Graph2X, where X = set (MLJ, Minor revision)

Column Bundle: Generic architecture for multi-X learning, where X = view, instance, label, task

Matrix representation of graphs & relation

MEMORY-AUGMENTED NEURAL NETWORKS

- Complex, long-range dependencies demand external memory
- Resemble how modern computer works (CPU & RAM)

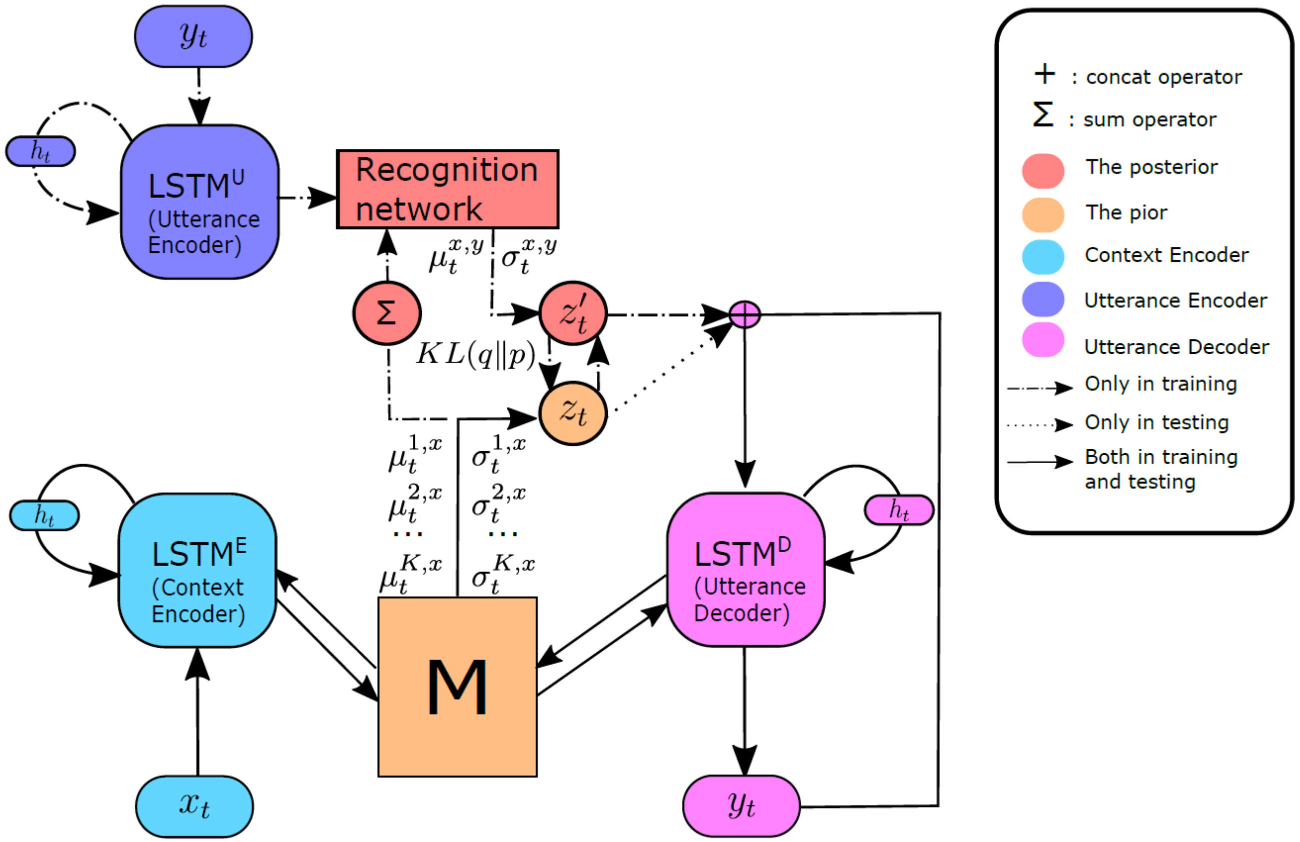

MEMORY AND TURING MACHINES

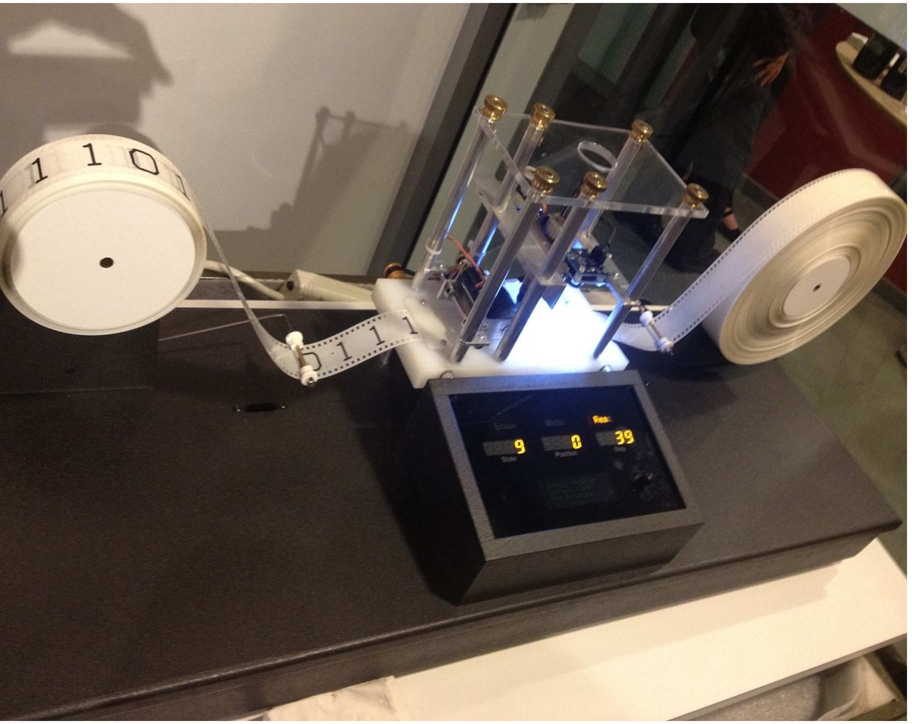

Can we learn from data a model that is as powerful as a Turing machine?

It is possible to invent a *single machine* which can be used to compute *any* computable sequence. If this machine **U** is supplied with the tape on the beginning of which is written the string of quintuples separated by semicolons of some computing machine **M**, then **U** will compute the same sequence as **M**.

Wikipedia

MANN: WHAT WE HAVE DONE

Dual-view in sequences (KDD'18)

Bringing variability in output sequences (NIPS'18)

Bringing relational structures into memory (ICPR'18)

Learning to skim read (ICLR'19)

Learning to program (on going)

ON GOING: NEURAL STORED-PROGRAM MEMORY

- Truly Turing machine: programs can be stored and called when needed.
- Can solve BIG problem with many sub-modules.
- Can do continual learning.
- Can help mitigate catastrophic forgetting.

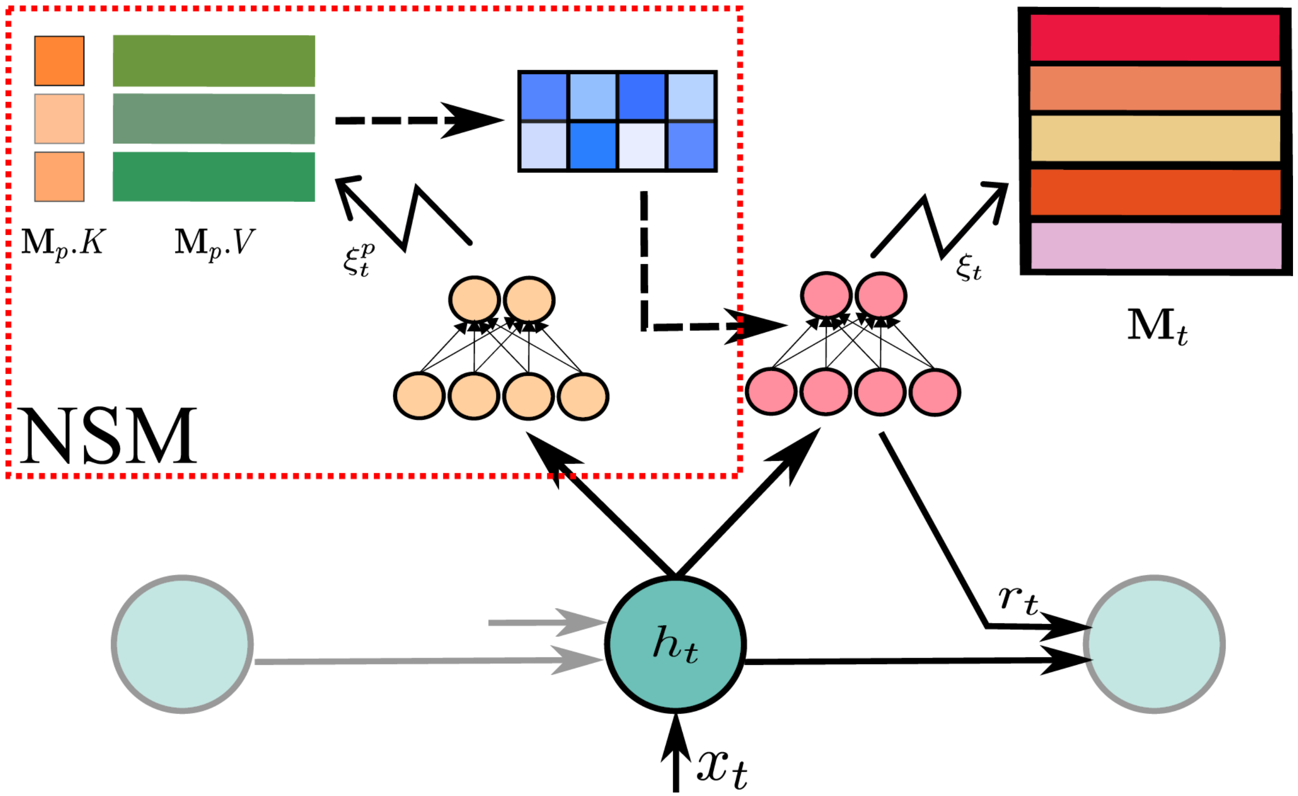

LEARNING TO REASON (L2R)

- L2R is learning to decide if a knowledge base entails a predicate.
- Can be cast as QA.
- Where neural meets symbolic
- Deep learning solves symbol grounding problem
- Integrating if fast and slow thinking
- Role of episodic memory to piece things together
- Role of working memory to support deliberative multi-step reasoning
- Knowledge-base L2R

CURRENT WORKS

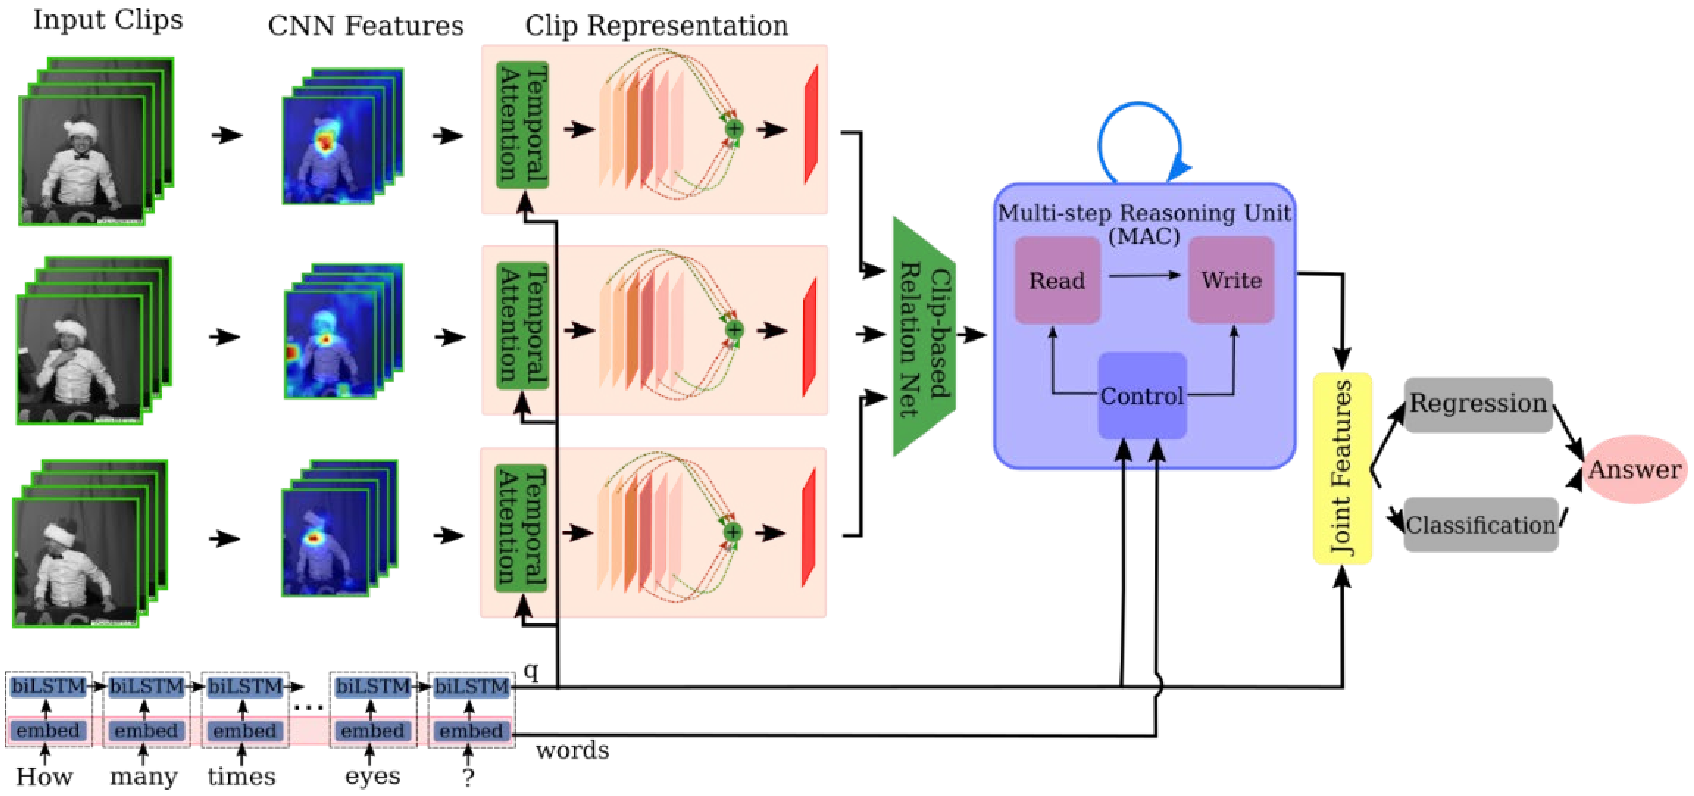

Graph reasoning

Video question answering

Protein-drug design: affinity prediction (drug as query) & drug generation (bioactivities as answers, what is the question?)

Conditional software execution

Knowledge graph querying and completion

QUERYING A GRAPH

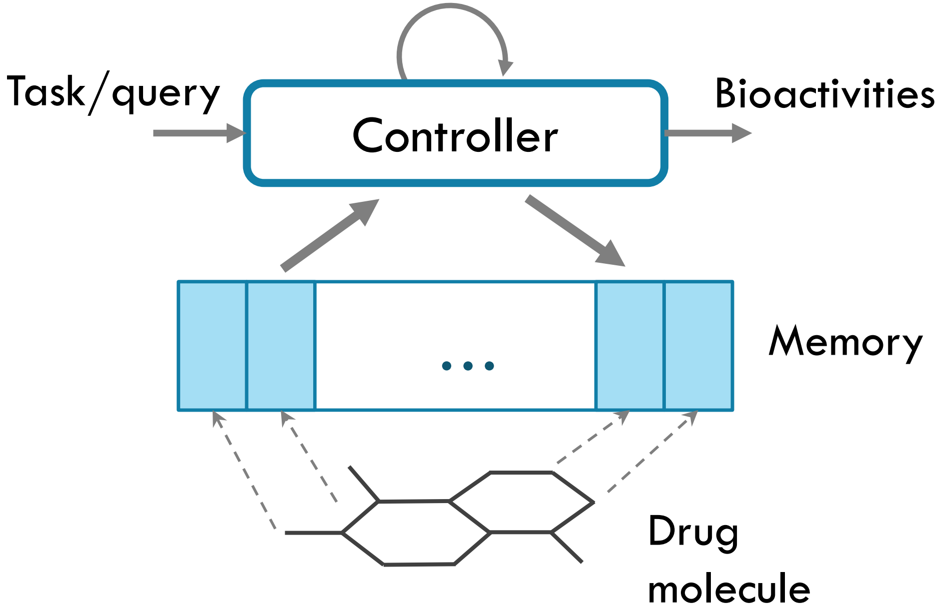

Message passing as refining atom representation

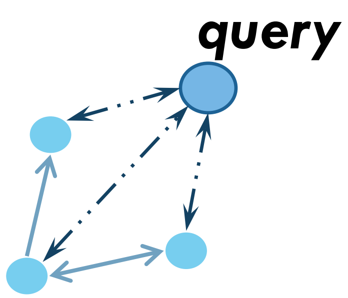

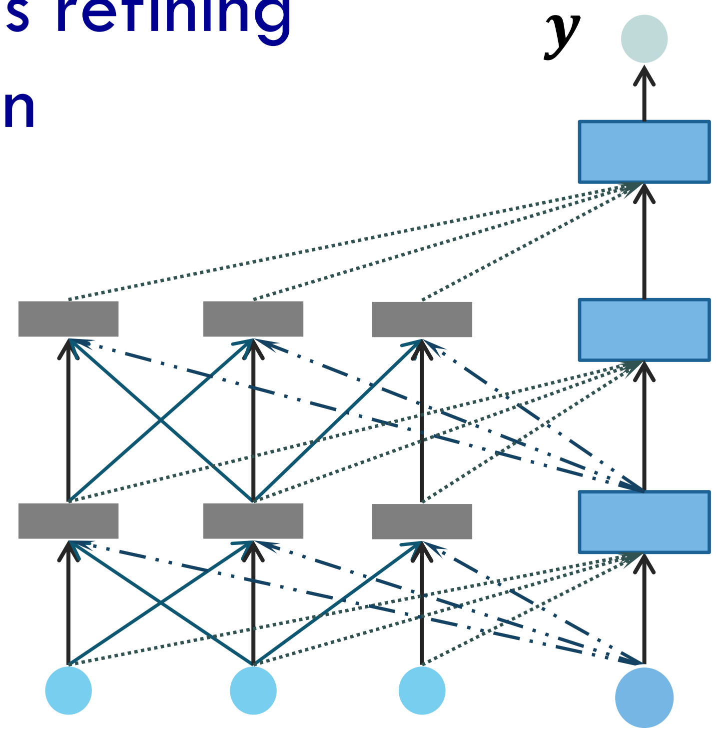

#Ref: Pham, Trang, Truyen Tran, and Svetha Venkatesh. "Graph Memory Networks for Molecular Activity Prediction." *ICPR'18*.

QUERYING MULTIPLE GRAPHS

A working memory view

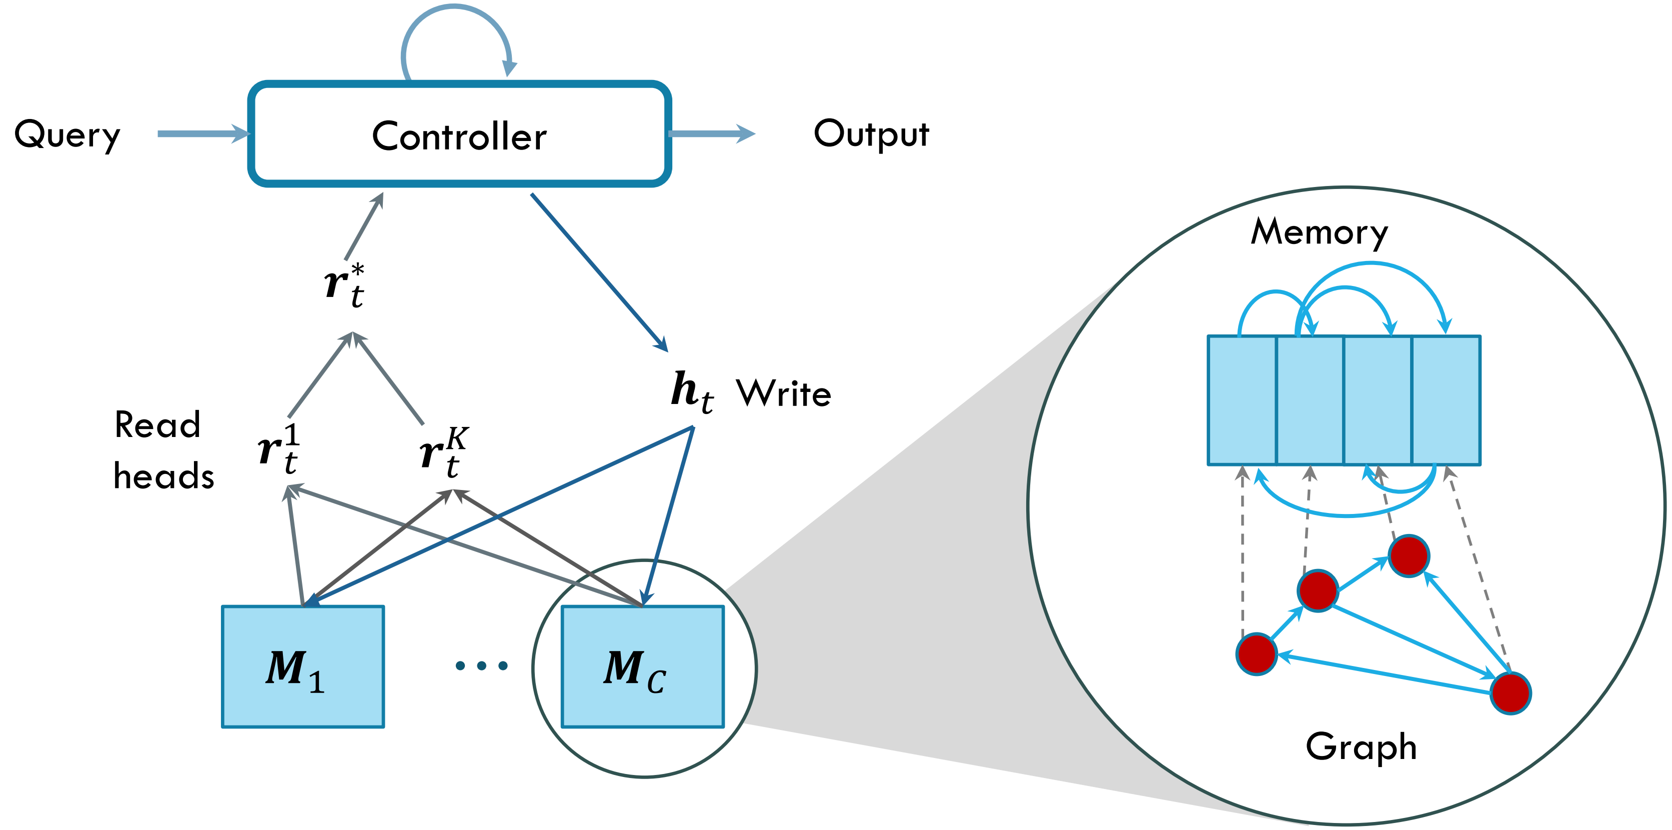

Pham, Trang, Truyen Tran, and Svetha Venkatesh. "Relational dynamic memory networks." *arXiv preprint arXiv:1808.04247*(2018).

VIDEO QA: OUR NEW STOTA RESULTS

Repetition Count

Question: How many times does the man kick his legs?

Preds: (GT) 6 (ST-TP) 4 (Ours) 6

Repeating Action

Question: What does the man do 5 times?

Preds: (GT) shake finger (ST-TP) turn (Ours) shake finger

State Transition

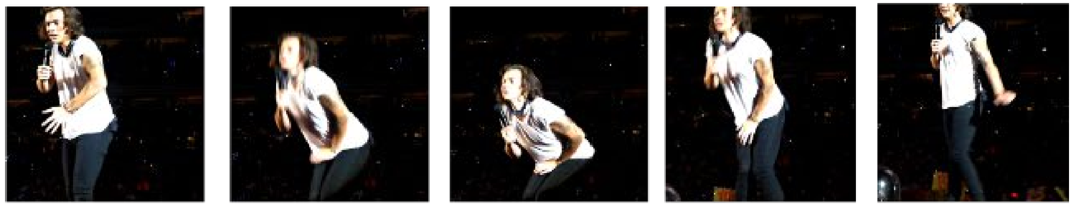

Question: What does the man do after squat?

Preds: (GT) stretch hand out (ST-TP) inhale (Ours) stretch hand out

Frame QA

Question: what is having its head rubbed ?

Preds: (GT) bird (ST-TP) dog (Ours) bird

INFERRING RELATIONS

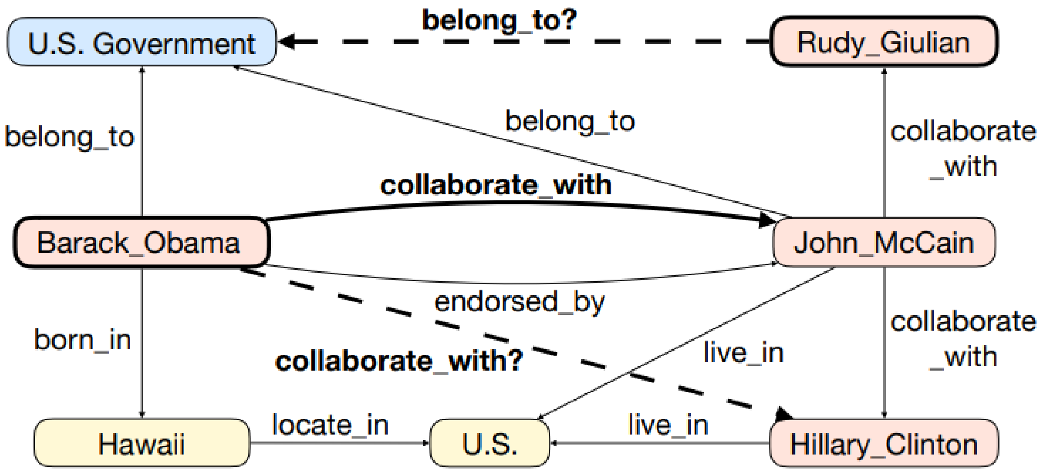

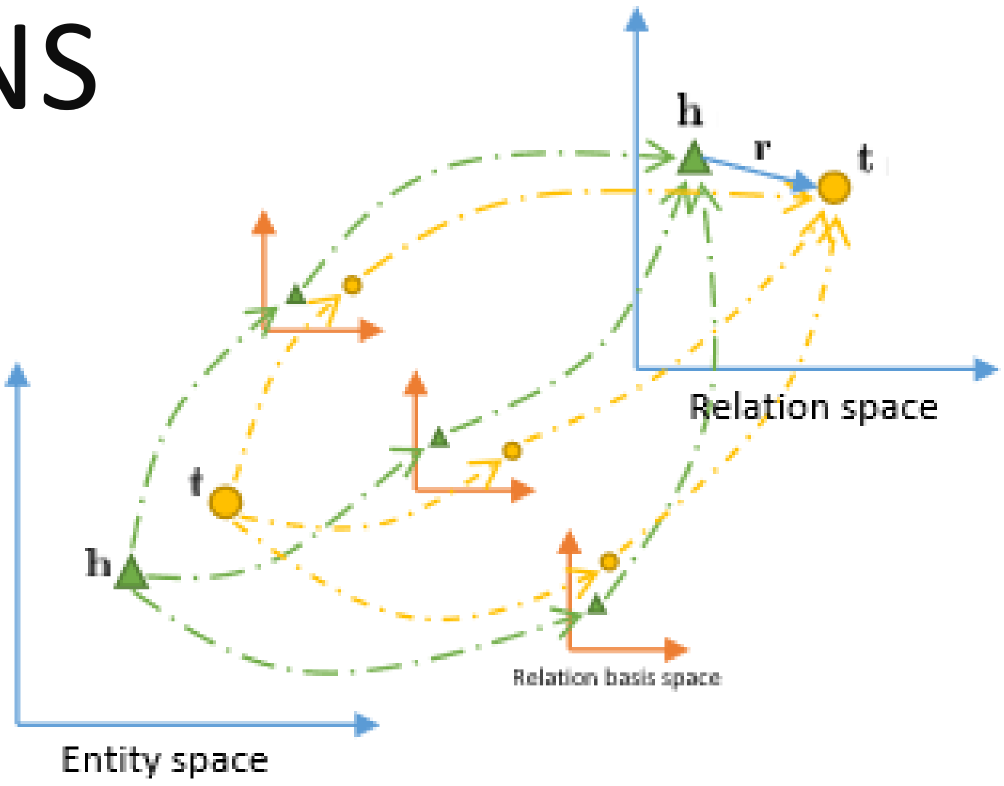

https://www.zdnet.com/article/salesforce-research-knowledge-graphs-and-machine-learning-to-power-einstein/

Do, Kien, Truyen Tran, and Svetha Venkatesh. "Knowledge graph embedding with multiple relation projections." *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2018.

LEARNING WITH LESS LABELS

- A hallmark of intelligence.
- Representation learning
- Deep generative models
- Continual learning

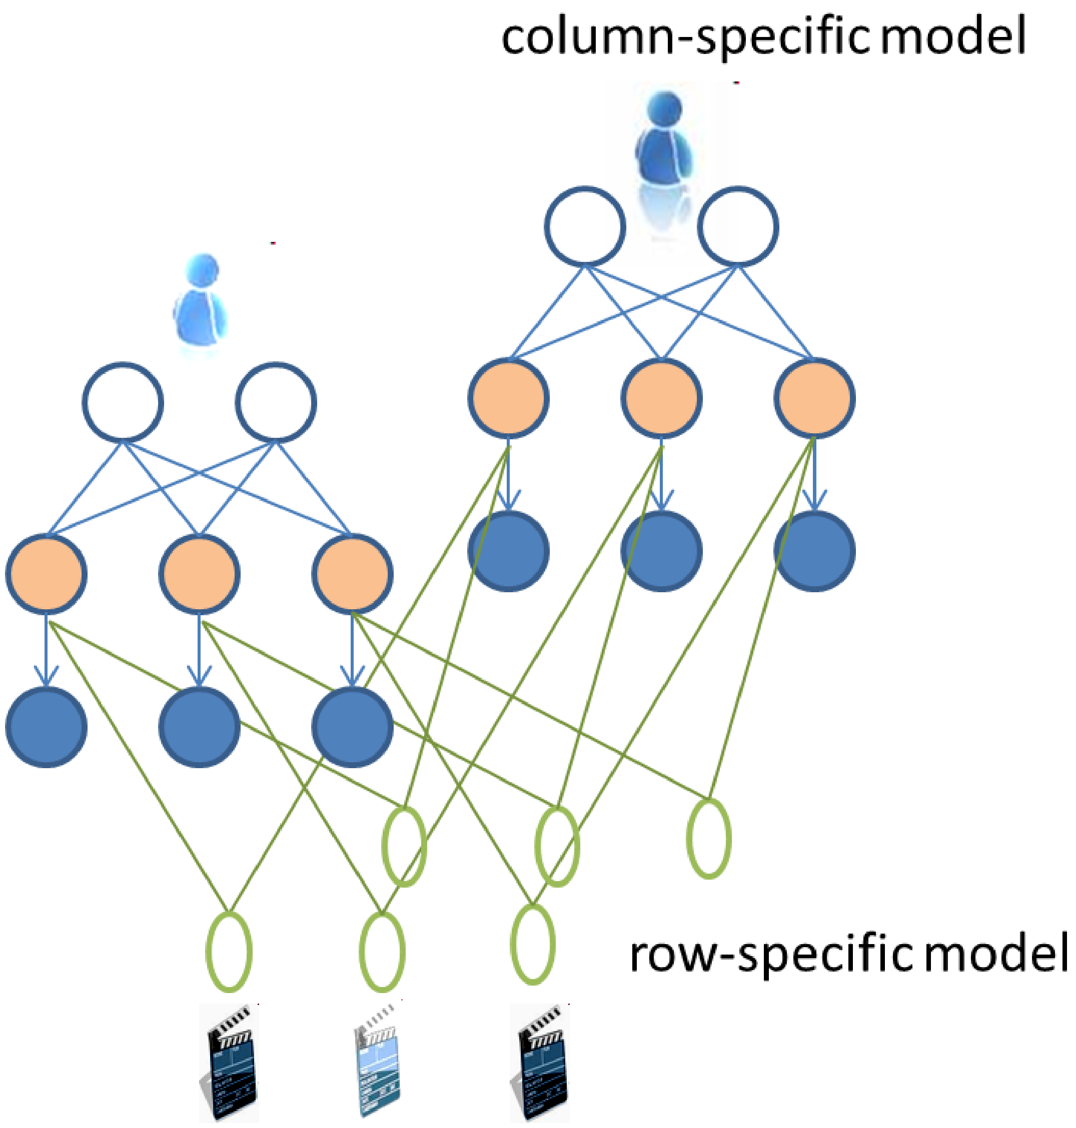

WORKS DONE

Restricted Boltzmann machines (2009-2016)

- Bimodal RBM (UAI'09)
- Mixed-variate RBM (ACML'11)
- Ordinal matrix RBM (ACML'12)
- Tensor-variate RBM (AAAI'15)

Continual learning, solve catastrophic forgetting (on going)

Deep generative models

- GAN as continual learning (ICML'18 workshop)
- Generalization and stability analysis of GAN (ICLR'19)
- Inverse design (SDM'19)
- **Learn within supported domain, and explore unsupported domain at the same time (on going)**

DEEP REINFORCEMENT LEARNING

New priors

RL equipped with memory

Multi-agent RL

Relational reasoning

Theory of AI/machine mind

Psychological games

RL for structured prediction

HUMAN BEHAVIOUR UNDERSTANDING

CVPR'19

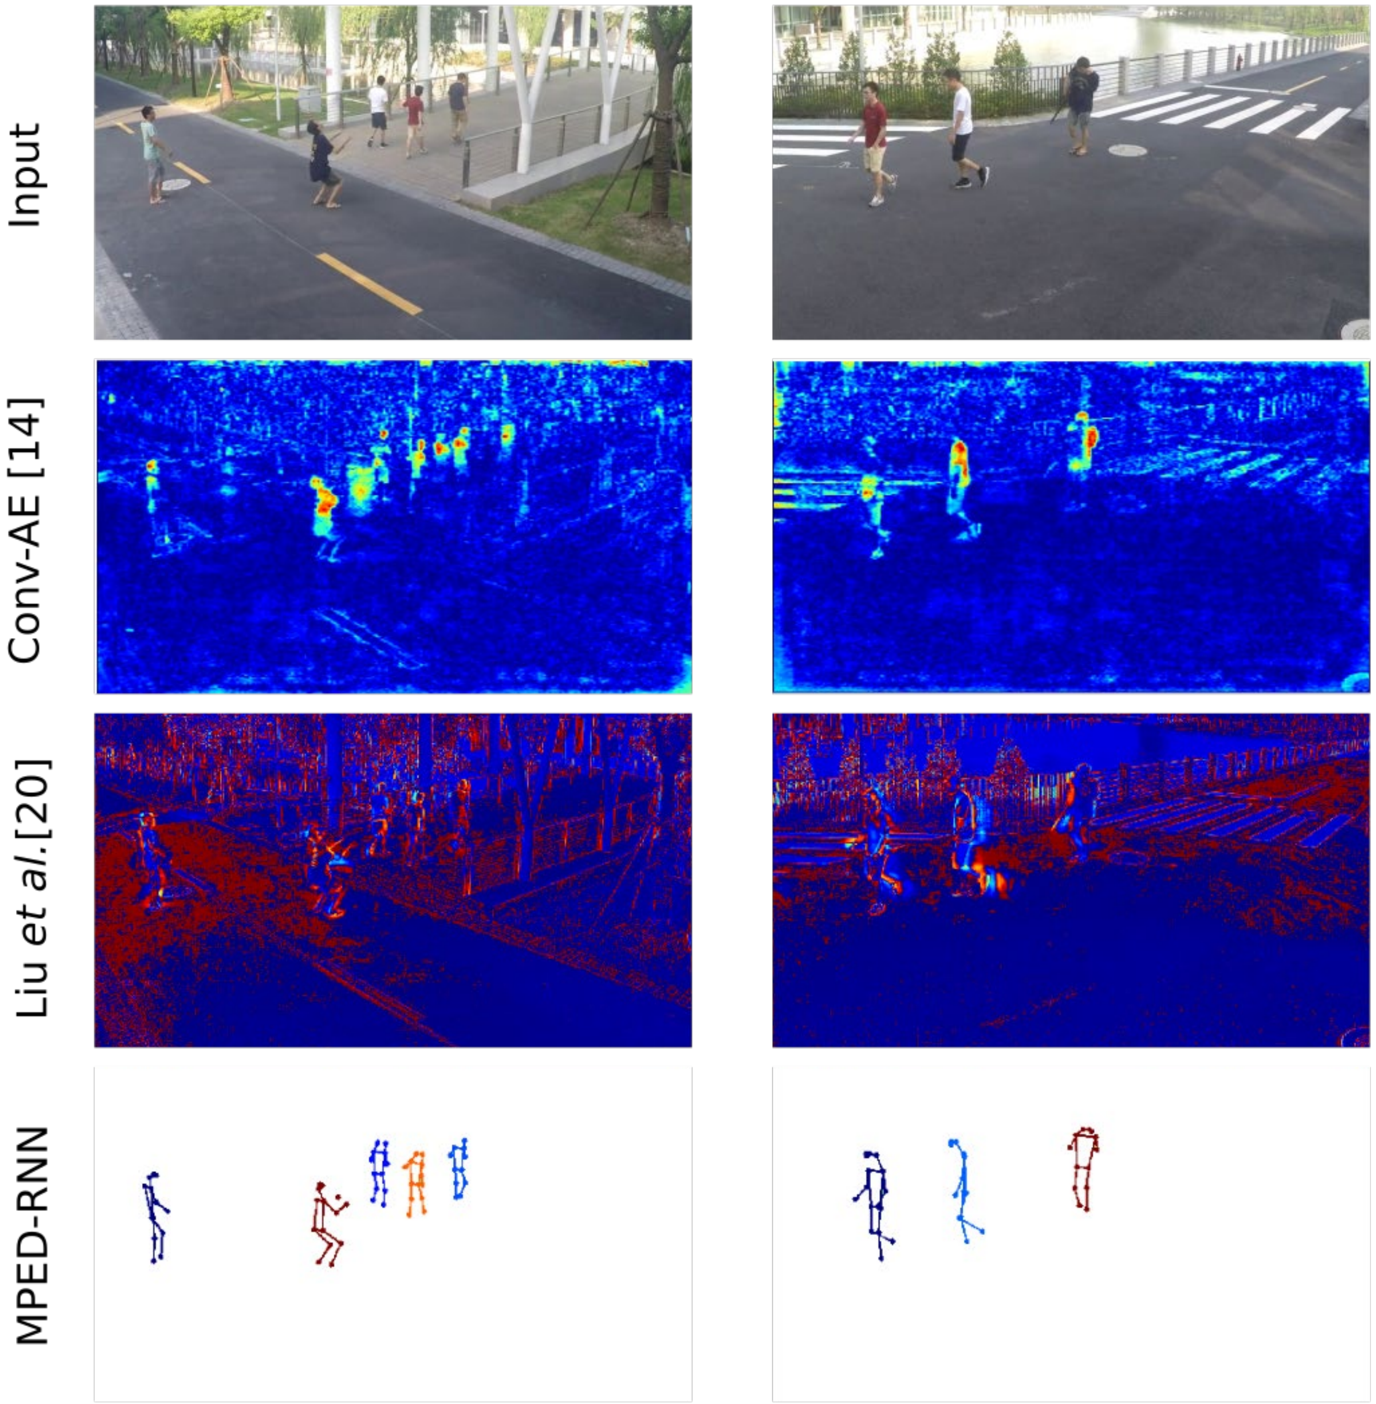

AI FOR AUTOMATED SOFTWARE ENGINEERING

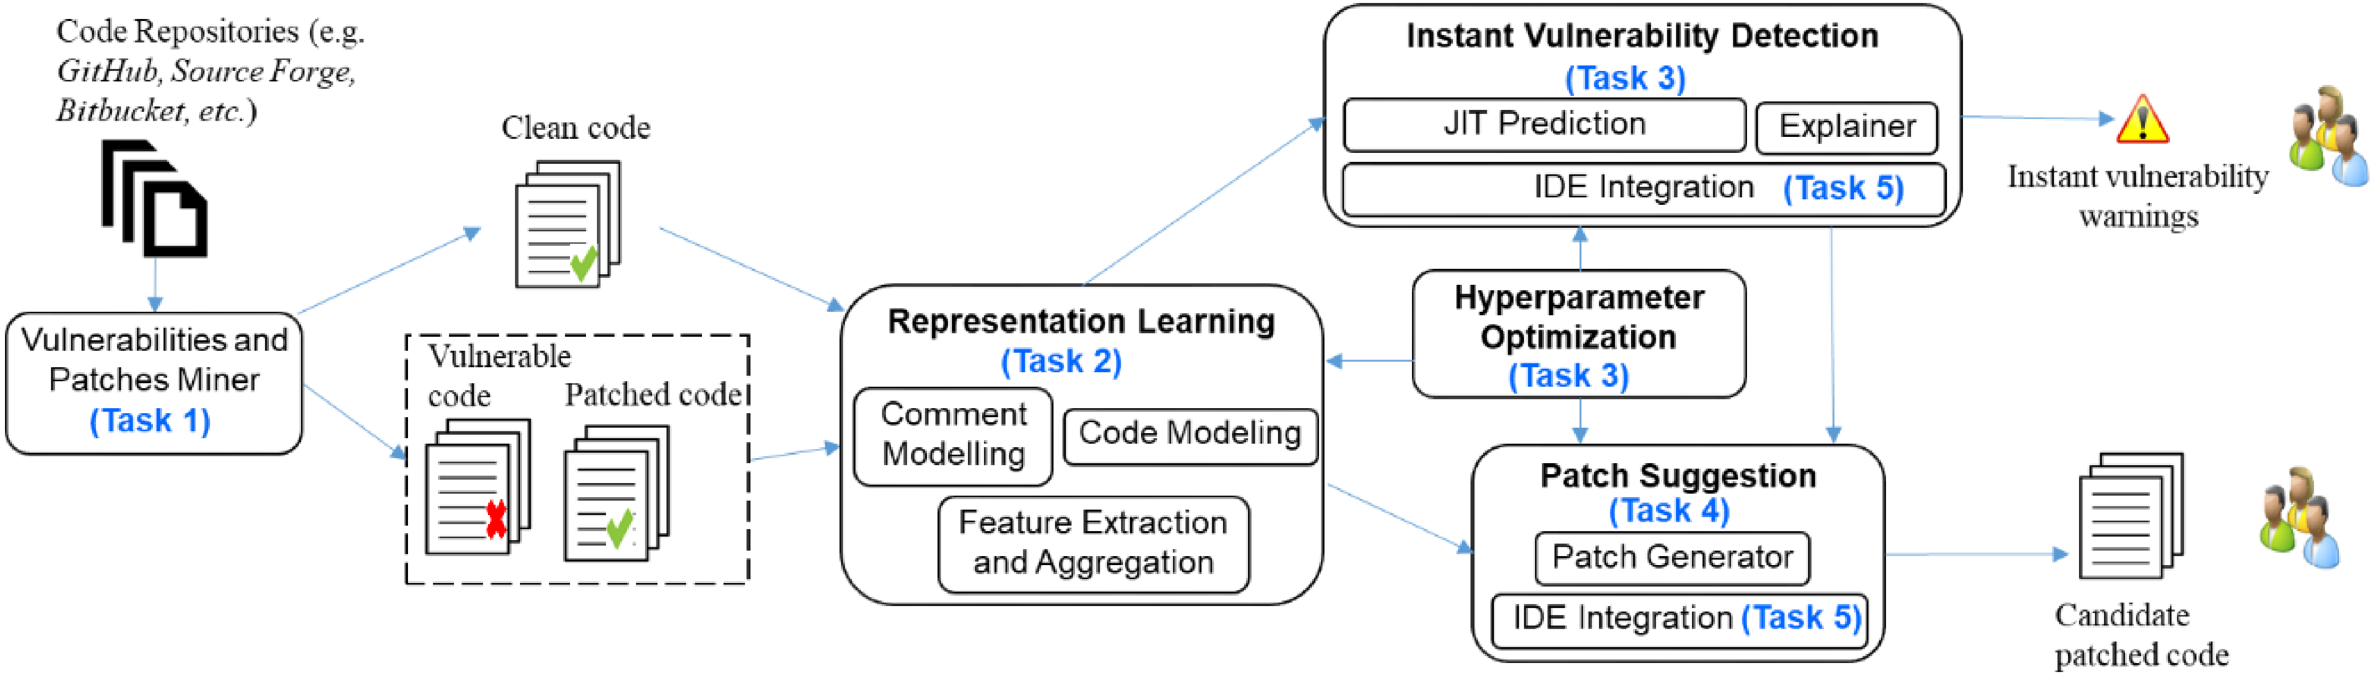

WORKS DONE

Project delay prediction with deep networked classification (ASE’15)

Code language model with LSTM (FSE’16 WS)

Defect prediction with Tree-LSTM (MSR’19)

Unsupervised feature learning for vulnerability detection (TSE, 2019)

Deep learning for story point estimation (TSE, 2018)

Graph with virtual nodes for software code modelling (IJCAI’17 WS)

Code repair (on going)

Code translation (on going)

Program synthesis (on going)

marketingland.com

AI FOR BIOMEDICINE

CONTEXT

2012-2016: Health focused

- Electronic Medical Records, ICU, EEG/ECG
- Social media health, wearable devices
- Intervention tools (Toby PlayPad)

2017-present: Genomics focused

- Cell functions modelling
- Drug bioactivity prediction
- Drug response against gene expressions
- BioNLP: Natural Language Processing for biology

2017-present: PAGI - Program in Advanced Genomic Investigation

- **Deakin + Garvan Partnership**
- "Big Data" machine learning approach to genomics
- Leveraging computing power to crunch data
- Whole-genome sequencing

GARVAN INSTITUTE OF MEDICAL RESEARCH

CONVERSATIONAL AI

Memory-augmented chatbot

Lifelong conversational agent

Chatbot with persona and ethics

Chatbot with theory of human mind

Multi-agent communication

Learning from instructions

Clinical dialog systems

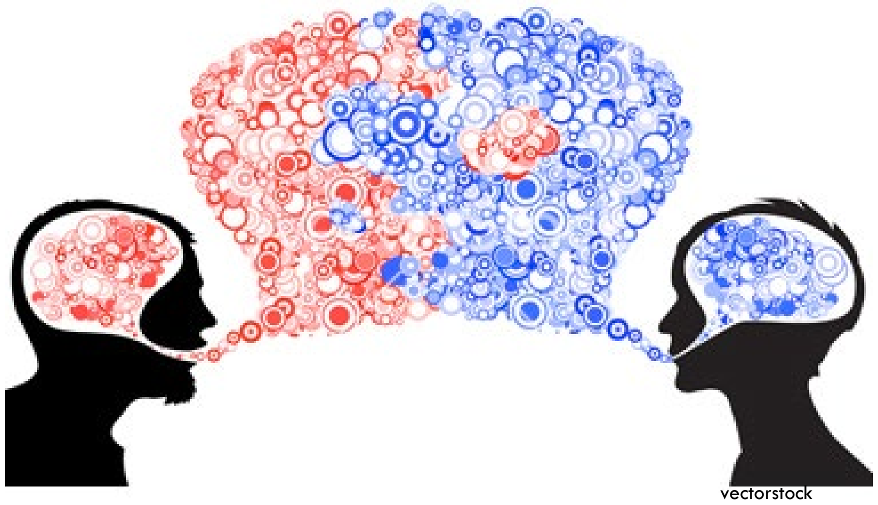

VARIATIONAL MEMORY ENCODER DECODER, NIPS'18

Input context	Response
Reddit comment: What is your favorite scene in film history ? Mine is the restaurant scene in the Godfather.	**Seq2Seq:** The scene in **Seq2Seq-att:** The final **DNC:** The scene in **CVAE:** Inception god! Not by a shark /*/ Amour great /*/ Pro thing you know 3 dead **VMED (K=3):** The opening scene from history movie /*/ The scene in a shot nights! Robin movie /*/ The psycho scene in fight from
Reddit comment: What actor will win an Oscar in the next 10 years ?	**Seq2Seq:** Colin **Seq2Seq-att:** Liam Neeson **DNC:** Tom Gyllenhaal **CVAE:** Daryl and Aaron /*/ Carefully count Alfred Deniro /*/ Ponyo Joker posible **VMED (K=3):** Edward or Leo Dicaprio goes on /*/ Dicaprio will /*/ Dicaprio Tom has actually in jack on road

AI for molecules

- Represent atom/molecular space
- Predict molecular properties
- Estimate chem-chem interaction
- Predict chemical reaction
- Fast search for new molecules
- Plan chemical synthesis

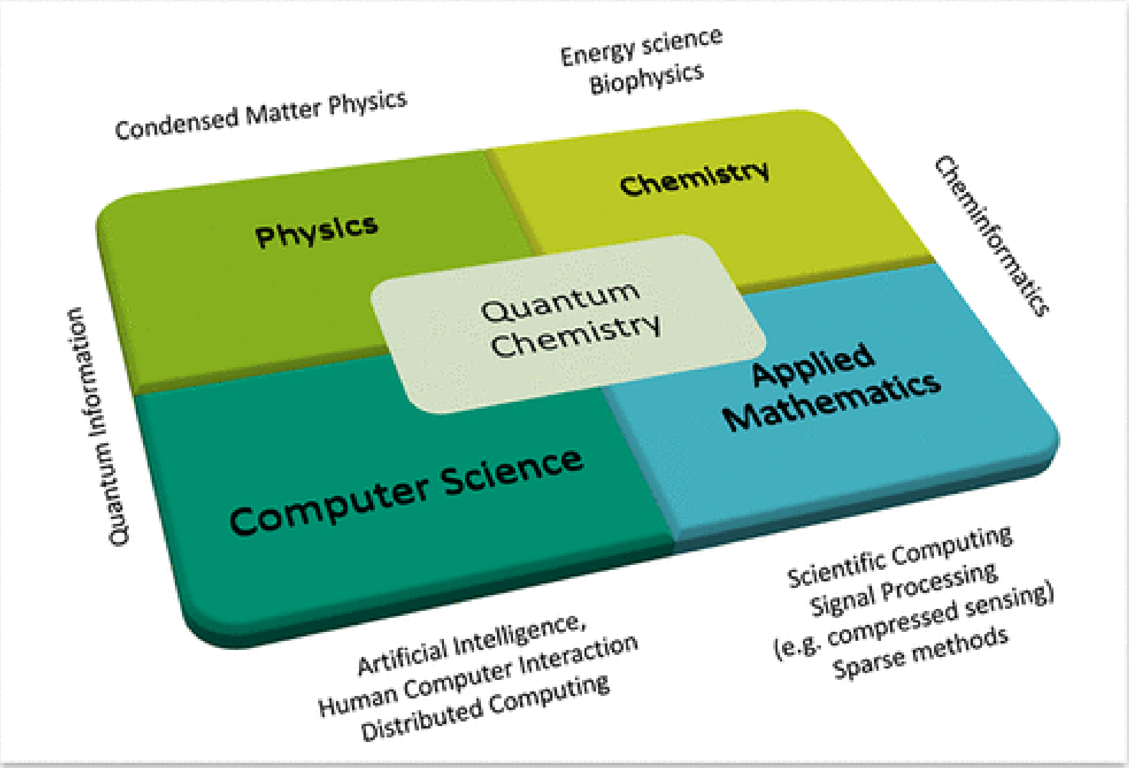

https://pubs.acs.org/doi/full/10.1021/acscentsci.7b00550

AI for materials

- Characterise the space of materials
- Represent crystals
- Map alloy composition → phase diagram
- Inverse design: Map phase diagram → alloy composition
- Generate alloys
- Optimize processing parameters

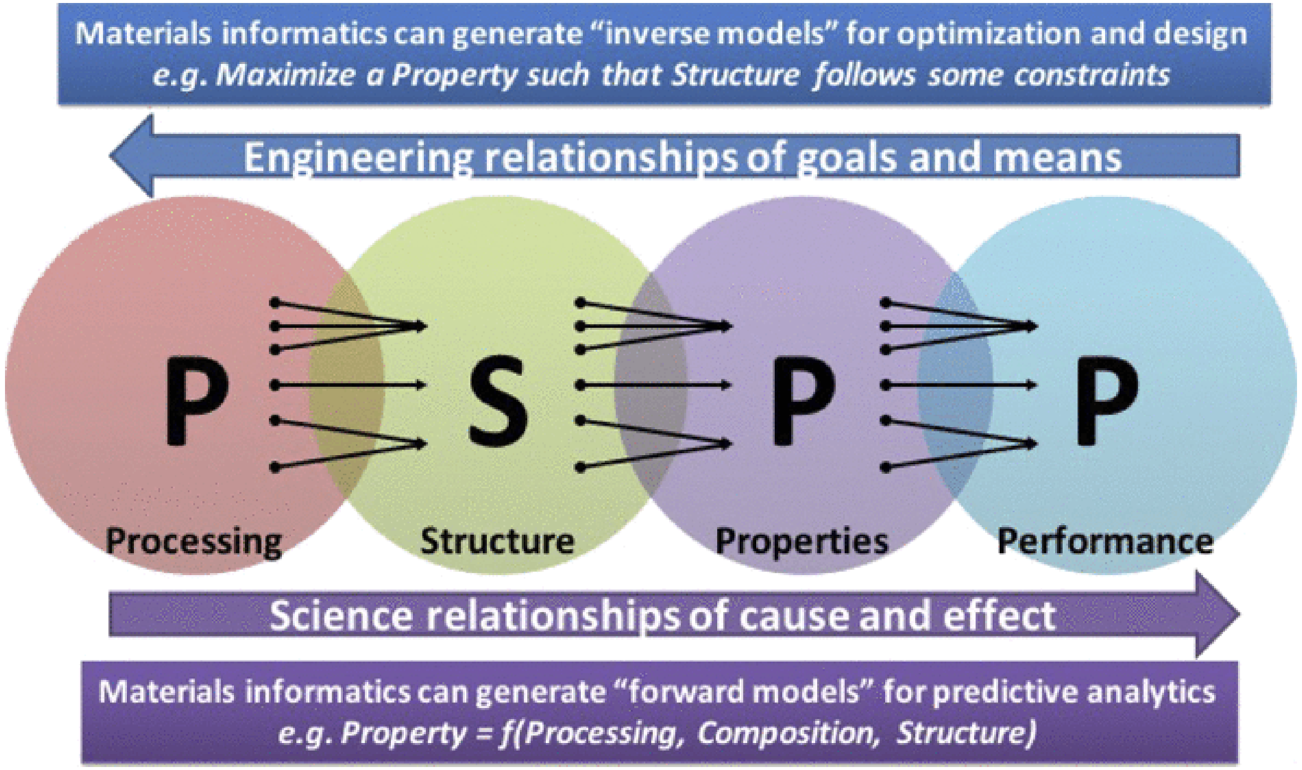

Agrawal, A., & Choudhary, A. (2016). Perspective: Materials informatics and big data: Realization of the "fourth paradigm" of science in materials science. *Apl Materials*, *4*(5), 053208.

Alloy space exploration

- Scientific innovations are expensive
- One search per specific target
- Availability of growing data

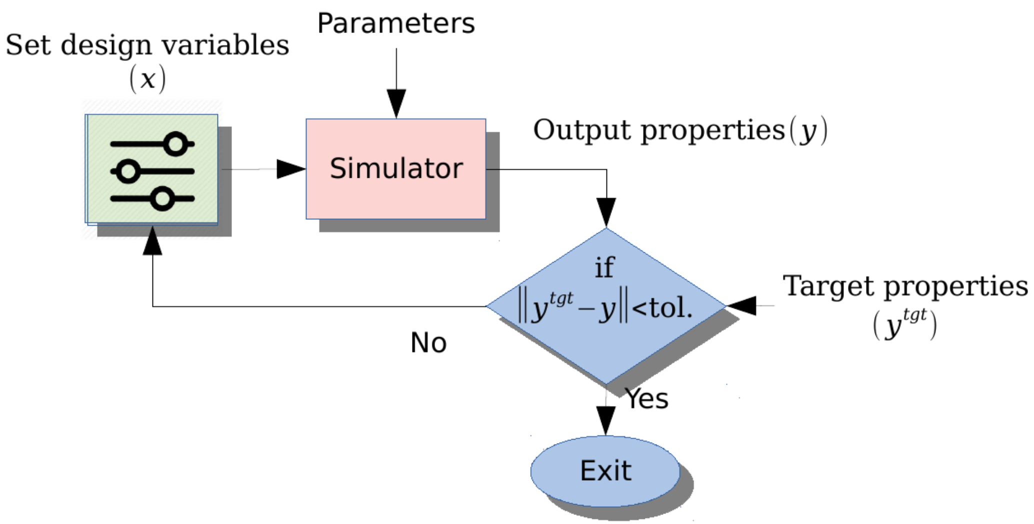

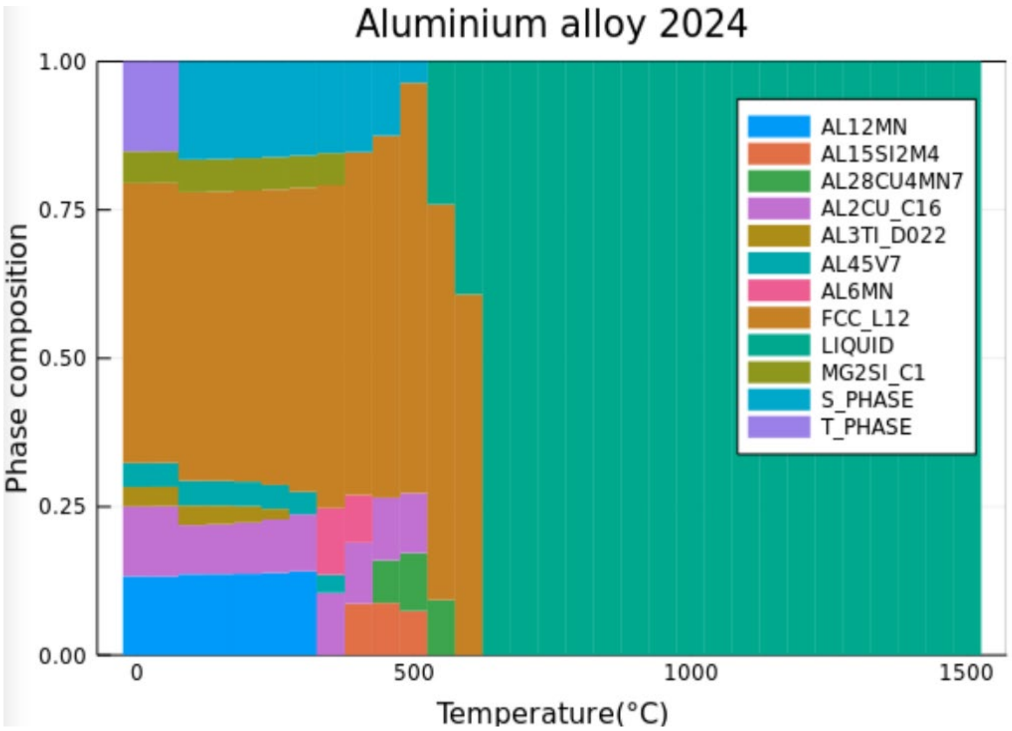

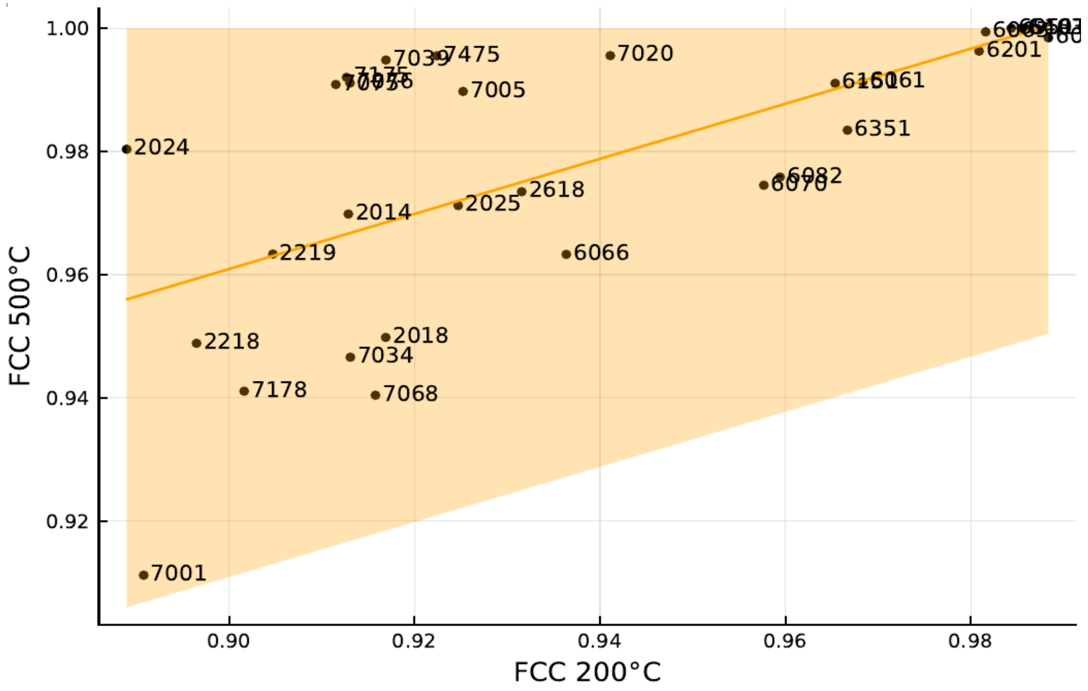

BUILDING A SMARTER HOME

Short-term: Anomaly detection

- Trajectories modelling with spatio-temporal models under irregular timing
- Fast adapting to changing condition
- Robust against noise

Mid-term: A conversational agent that understands context and history

- Agents with episodic memory
- Contextual reasoning
- Knowledge and commonsene

Long-term: Developing a lifelong digital companion

- Lifelong learning
- Theory of human mind

ARC Research Hub for DIGITAL ENHANCED LIVING

A NOVEL, CONTEXT-AWARE, UNSUPERVISED LEARNING PLATFORM TO LEVERAGE SENSING AND ANOMALY DATA

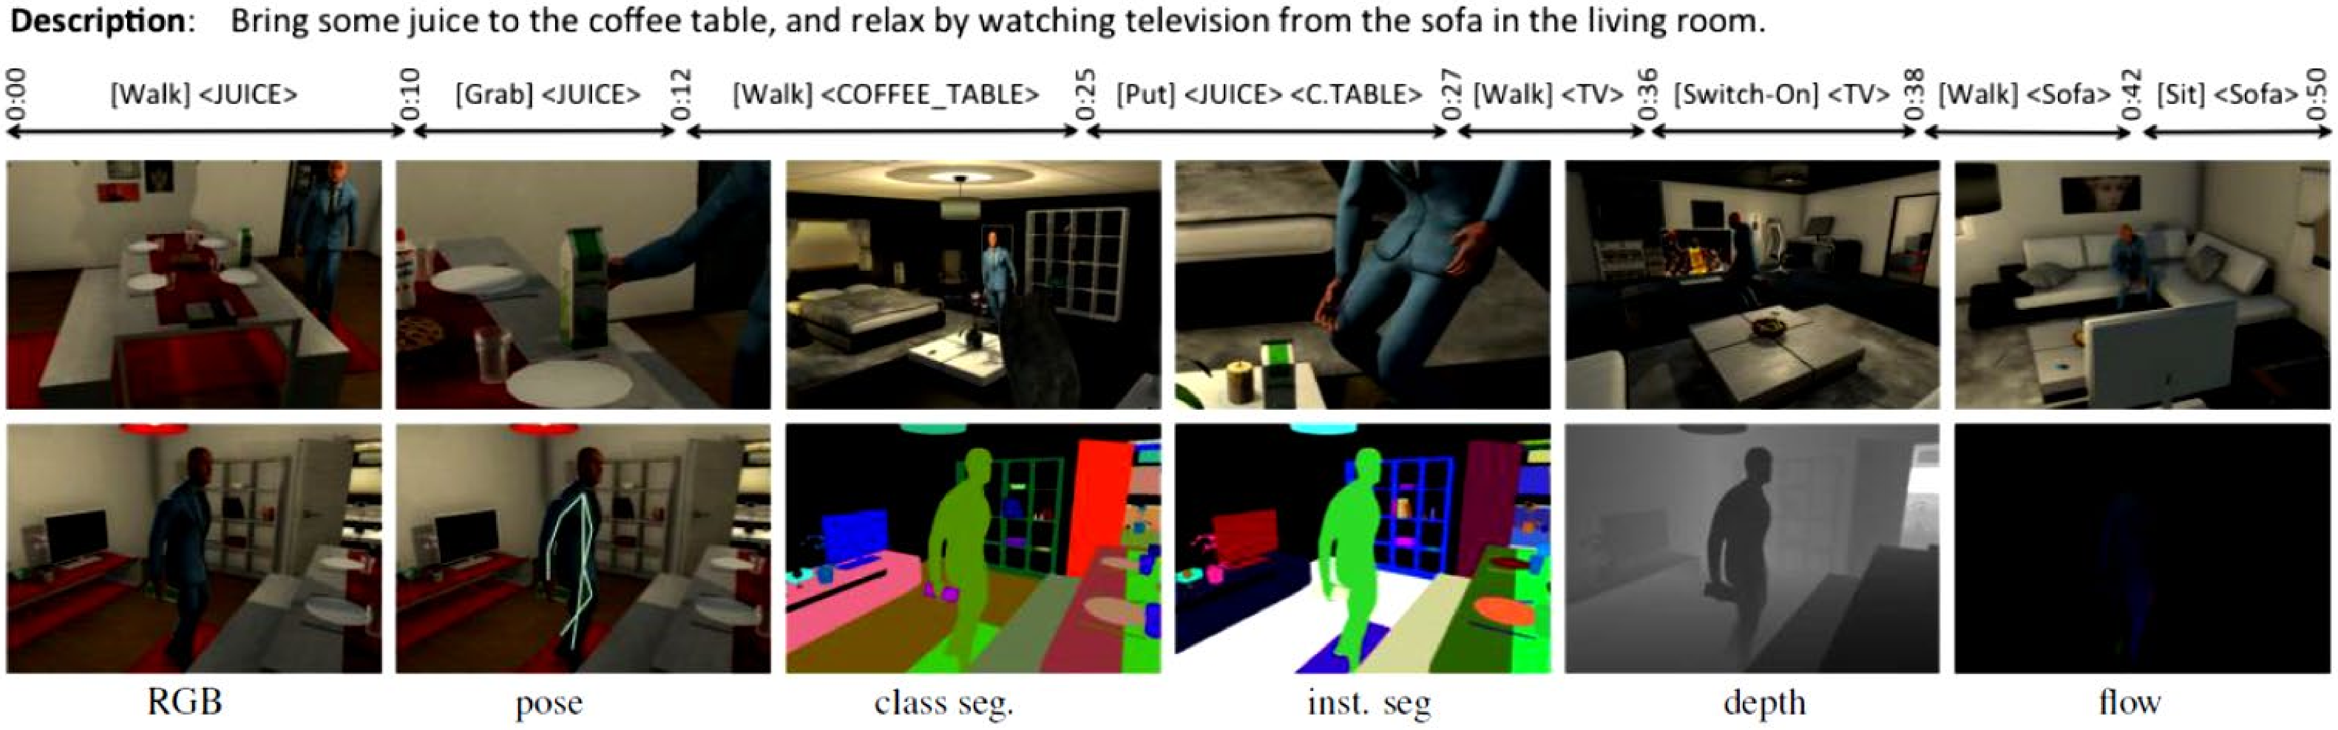

#REF: Puig, Xavier, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. "VirtualHome: Simulating Household Activities via Programs." In *CVPR*. 2018.

NEW DEEP GENERATIVE MODELS FOR SMART-HOME SENSORY DATA

Represent, learn and infer about human activities:

- Infer person's locations
- Count and disambiguate multiple persons (residents and visitors)
- Recognise interwoven and concurrent activities
- Detect anomalies

Model multi-resolution activities from multi-modal, noisy sensory data.

- Integrate temporal resolutions (minutes to months) & spatial resolutions (person bounding box to entire home).

Handle (dynamic) contexts in a residential care.

- Adapt automatically without human intervention.

KEY TECHNICAL CHALLENGES

- Learning from minimal, unobtrusive sensing;
- Incremental, lifelong unsupervised learning with sparse, delayed feedbacks;
- Robustly handling spurious anomalies that arise from routine environmental signals (e.g. pets, outside light changes, wind, etc.);
- Transfer from simulated home to real home.

VALUE-ALIGNED MACHINE LEARNING

Value regularisation in the context of conversational AI

Personalised values

Value-centric multi-agent communication

Theory of others' values

Learning to reason legally with values

Thank you!

We’re hiring

PhD & Postdocs

truyen.tran@deakin.edu.au

https://truyentran.github.io/scholarship.html

http://ahsanqawl.com/2015/10/qa/