Deep Learning - the Key Technology of Artificial Intelligence

Sepp Hochreiter

JKU LIT AI LAB

Audi | JKU deep learning center

IARAI institute of advanced research in artificial intelligence

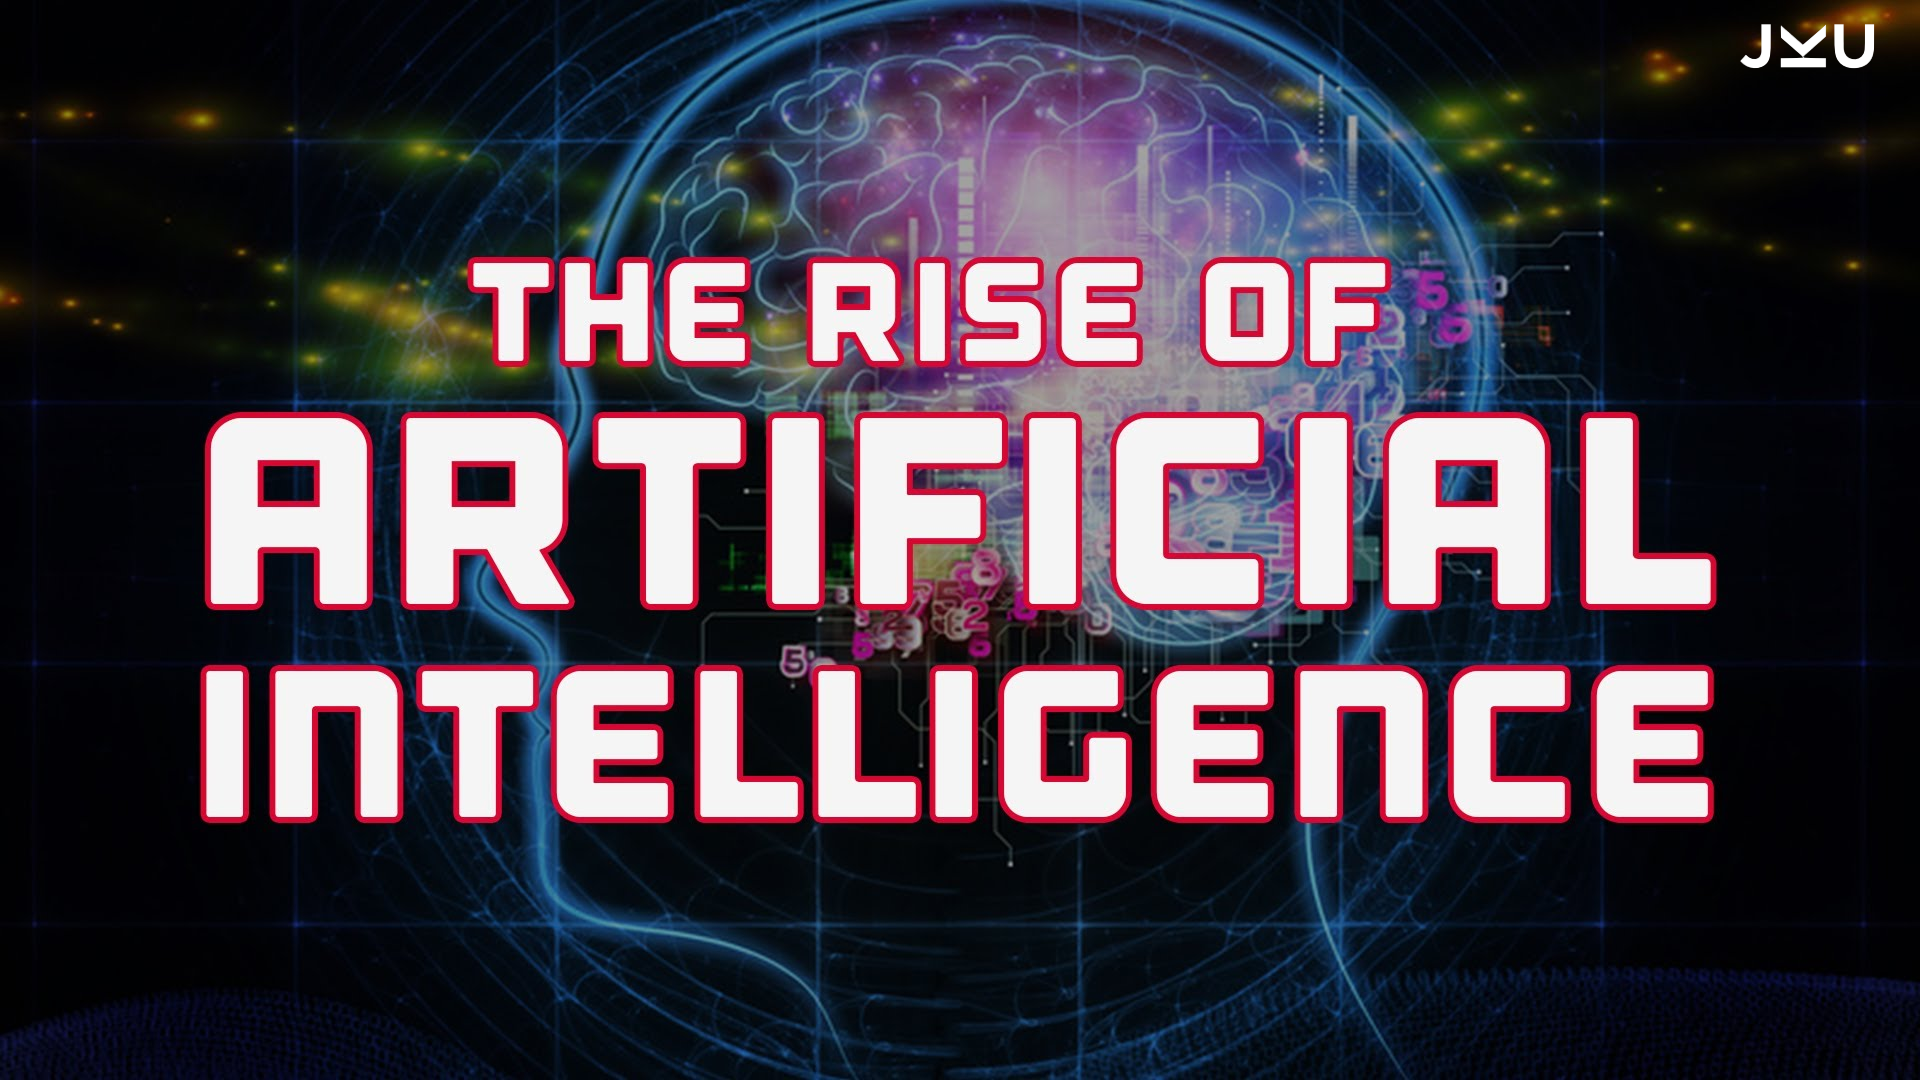
JKU
THE RISE OF
ARTIFICIAL
INTELLIGENCE

WHAT IS A.I.?

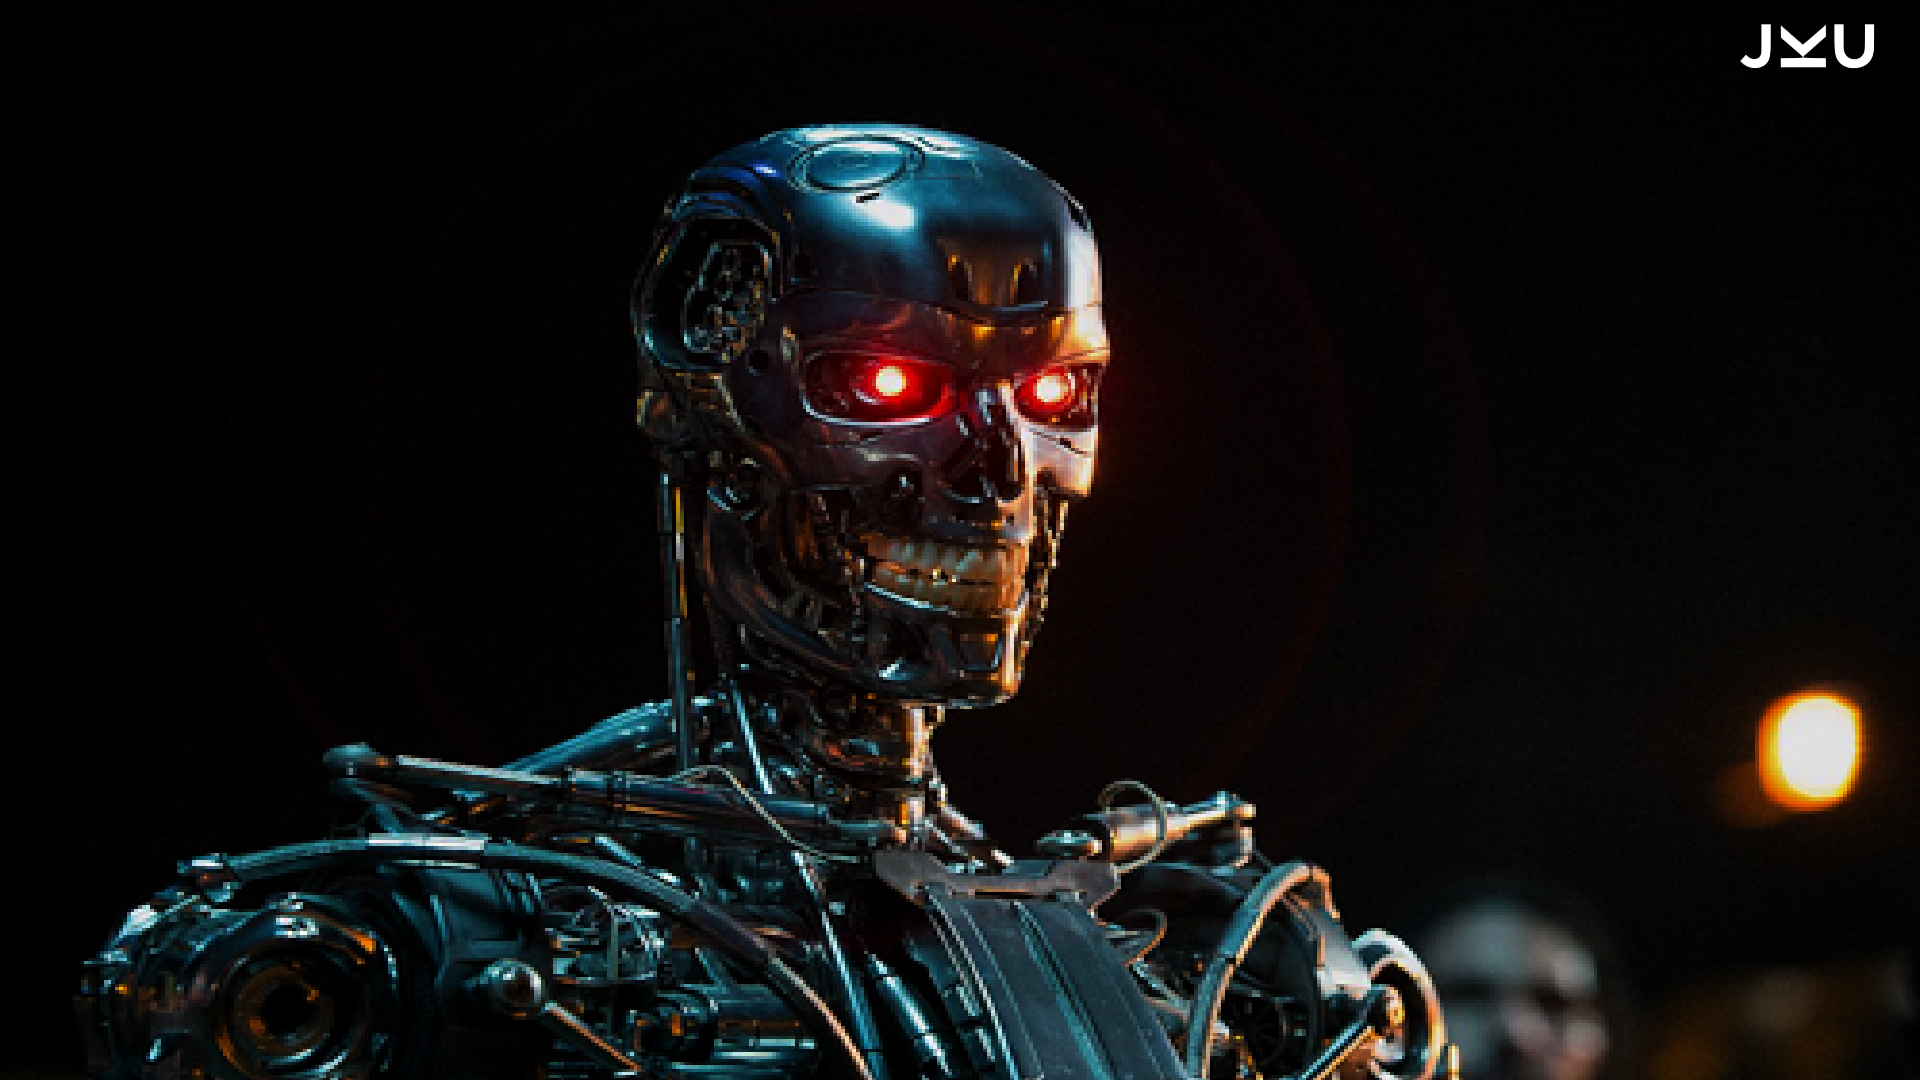

Artificial Intelligence (AI)

- In 1956 coined by John McCarthy
- A machine that **exhibits cognitive functions like humans** (learning, planning, reasoning, problem solving)
- **Deep Learning** successes drive the current AI boom: revolutionized **vision**, **speech**, and **text processing**


JKU
Deep Learning
Intelligence
from Big Data

ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.

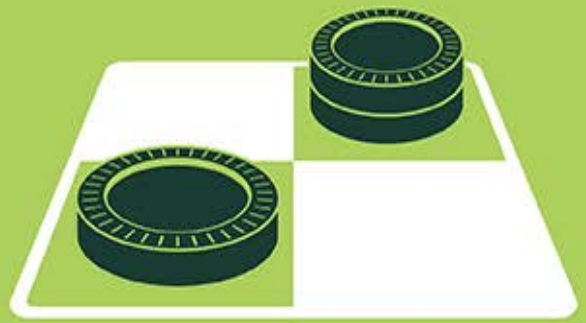

MACHINE LEARNING

Machine learning begins to flourish.

DEEP LEARNING

Deep learning breakthroughs drive AI boom.

Intelligent Machines & Programs

Ability to learn

Neural Networks

1950's 1960's 1970's 1980's 1990's 2000's 2010's

Deep Learning

- Neural networks & very fast computers & massive data sets

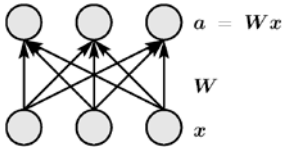

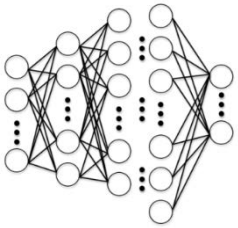

- Multiple levels of sparse representations → higher levels code abstract concepts

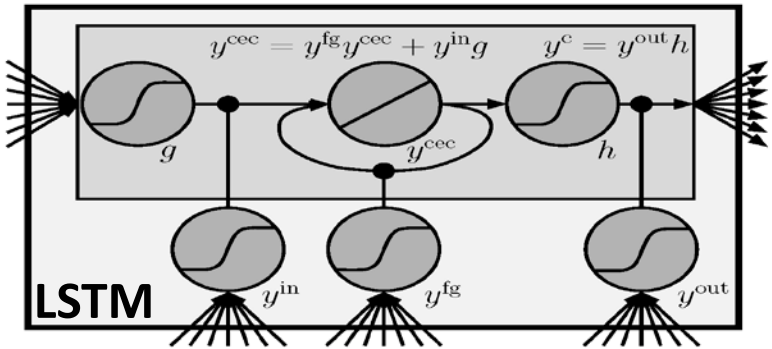

- Started already in 1991 with LSTM by Hochreiter and Schmidhuber

Deep Learning learns layers of features

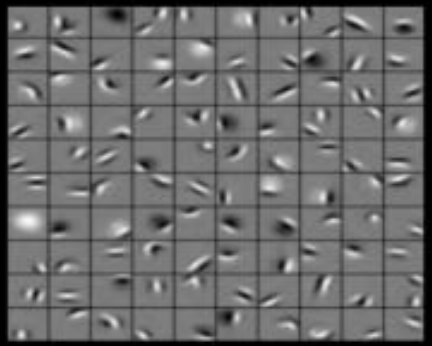

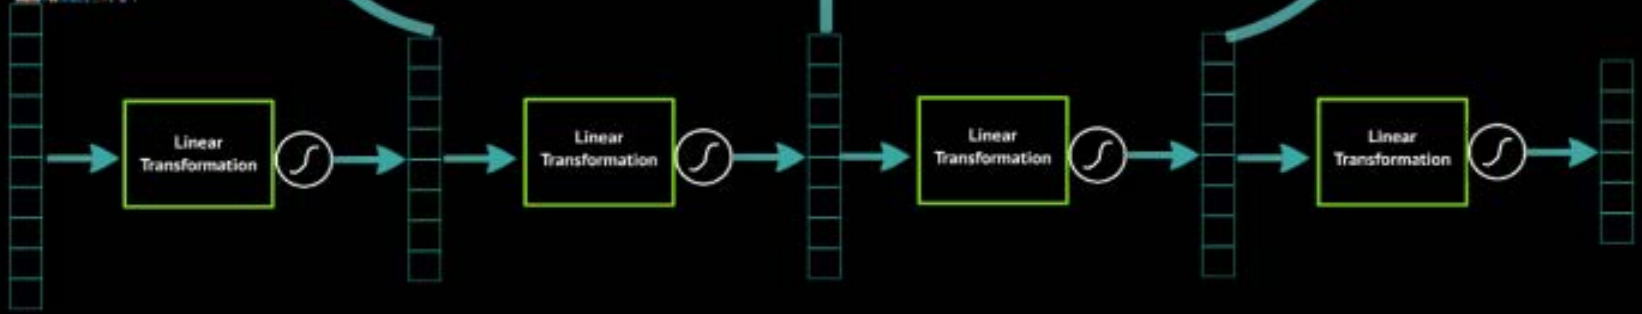

Records of Deep Learning

2011 speech recognition: on an industry benchmark: **16%** error for Deep Learning vs. **24%** by second

2013 vision: on the ImageNet competition with 1.2M images: **16%** error for Deep Learning vs. **26%** by second

Audio

TIMIT Phone classification	Accuracy
Prior art (Clarkson et al.,1999)	79.6%
Stanford Feature learning	**80.3%**

TIMIT Speaker identification	Accuracy
Prior art (Reynolds, 1995)	99.7%
Stanford Feature learning	**100.0%**

Images

CIFAR Object classification	Accuracy
Prior art (Krizhevsky, 2010)	78.9%
Stanford Feature learning	**81.5%**

NORB Object classification	Accuracy
Prior art (Ranzato et al., 2009)	94.4%
Stanford Feature learning	**97.3%**

Video

Hollywood2 Classification	Accuracy
Prior art (Laptev et al., 2004)	48%
Stanford Feature learning	**53%**

YouTube	Accuracy
Prior art (Liu et al., 2009)	71.2%
Stanford Feature learning	**75.8%**

KTH	Accuracy
Prior art (Wang et al., 2010)	92.1%
Stanford Feature learning	**93.9%**

UCF	Accuracy
Prior art (Wang et al., 2010)	85.6%
Stanford Feature learning	**86.5%**

Multimodal (audio/video)

AVLetters Lip reading	Accuracy
Prior art (Zhao et al., 2009)	58.9%
Stanford Feature learning	**65.8%**

Other records:

- **Pedestrian detection**
- Different phone recognition task
- PASCAL VOC **object classification**

Deep Learning

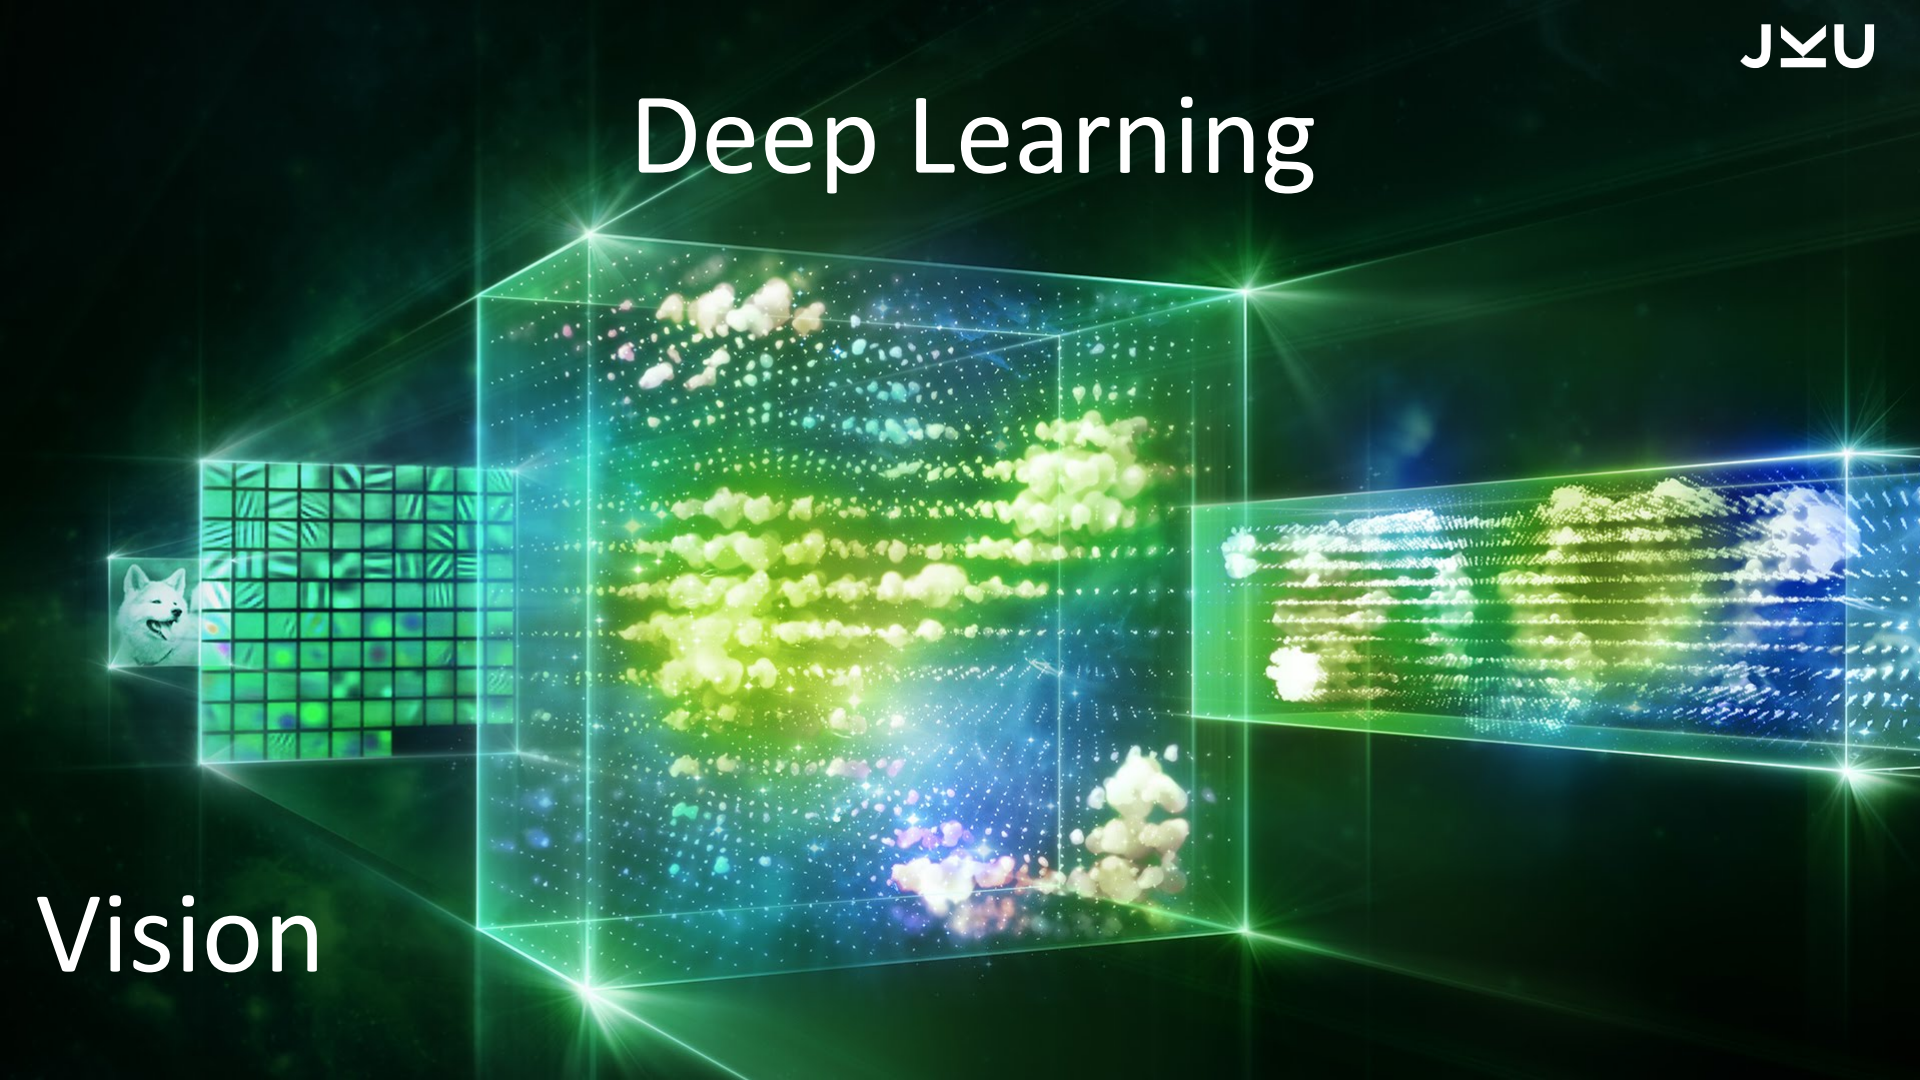

Vision

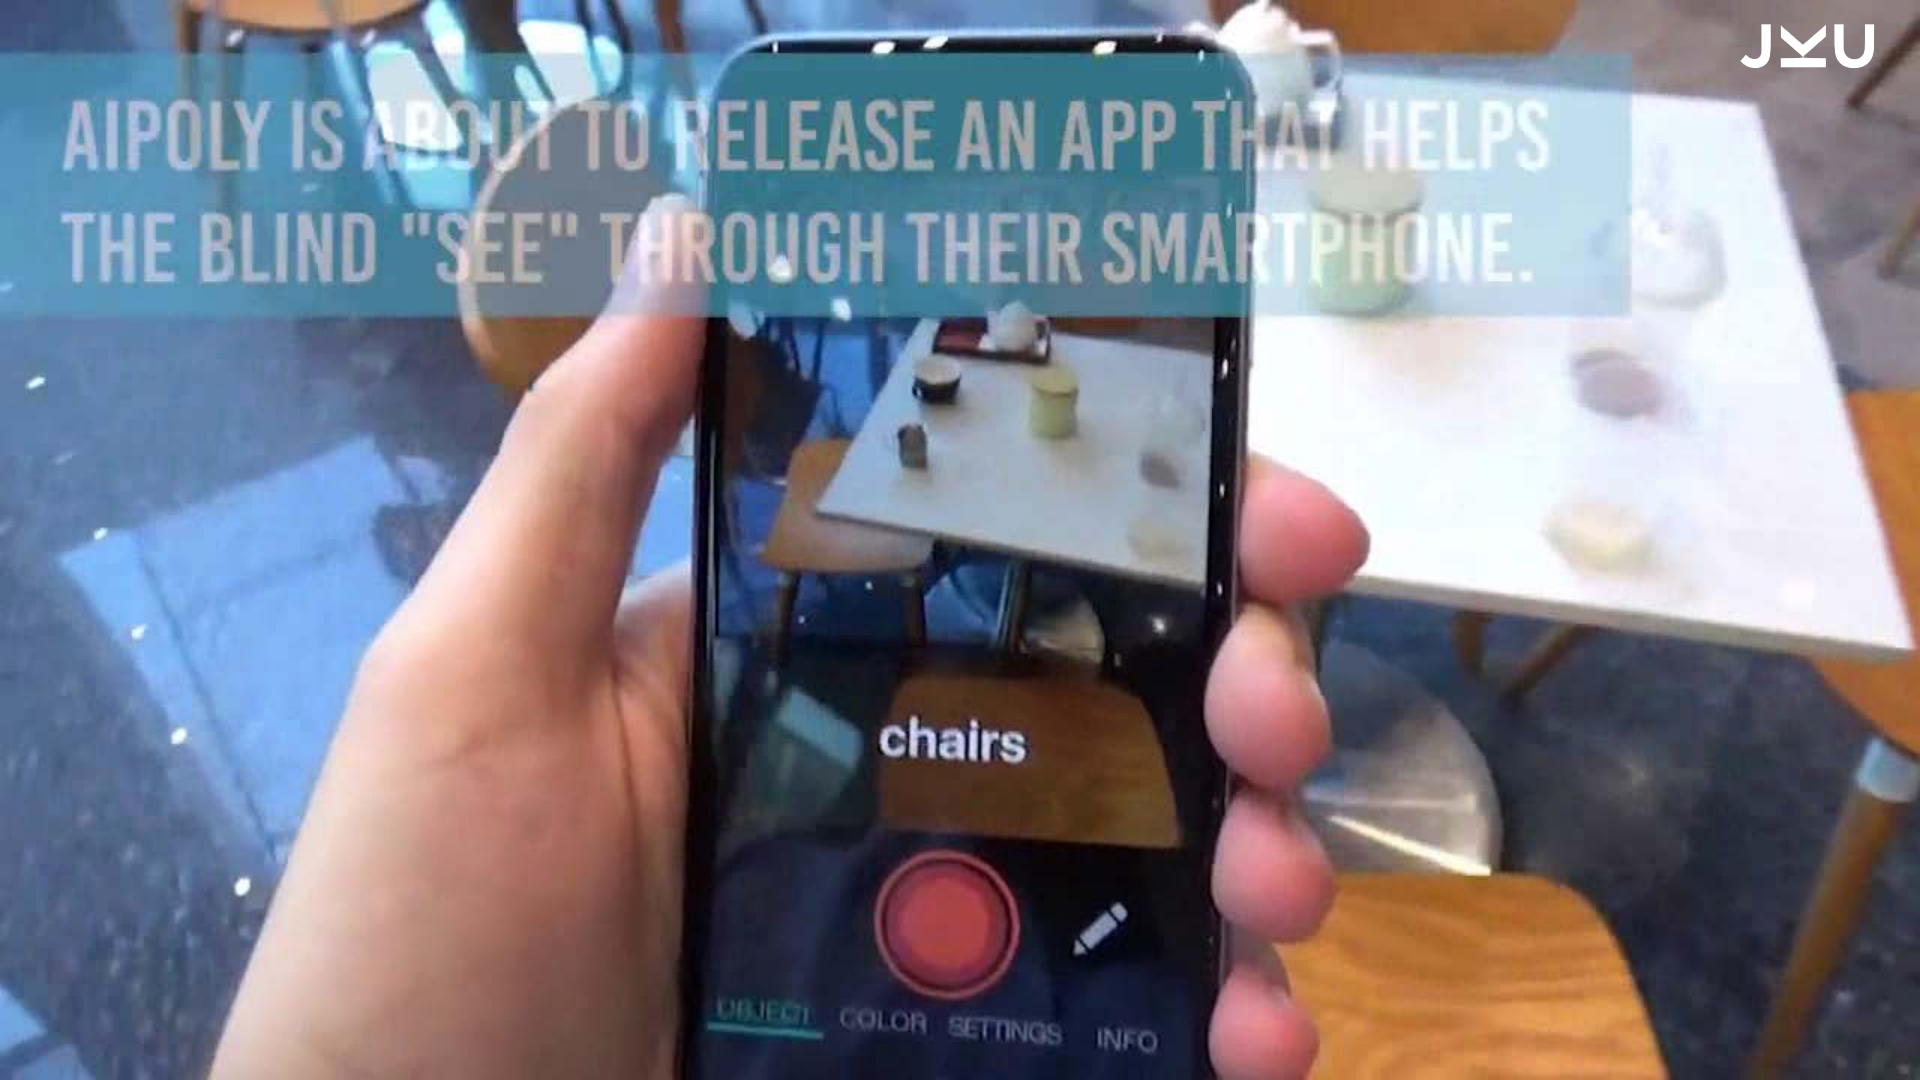
AIPOLY IS ABOUT TO RELEASE AN APP THAT HELPS
THE BLIND "SEE" THROUGH THEIR SMARTPHONE.
chairs
OBJECT
COLOR
SETTINGS
INFO

Source Actor

Target Actor

Real-time Reenactment

Reenactment Result






Original Photo

Example Photo

Result

JKU

Zebras ⇄ Horses

zebra → horse

horse → zebra

Summer ⇄ Winter

summer → winter

winter → summer


SIMULATION — THE PATH TO BILLIONS OF MILES
Volta
1/2 MILE
1/4 MILE
World drives trillions o
U.S. has 770 accident
A fleet of 20 test ca
per year.

Simulation

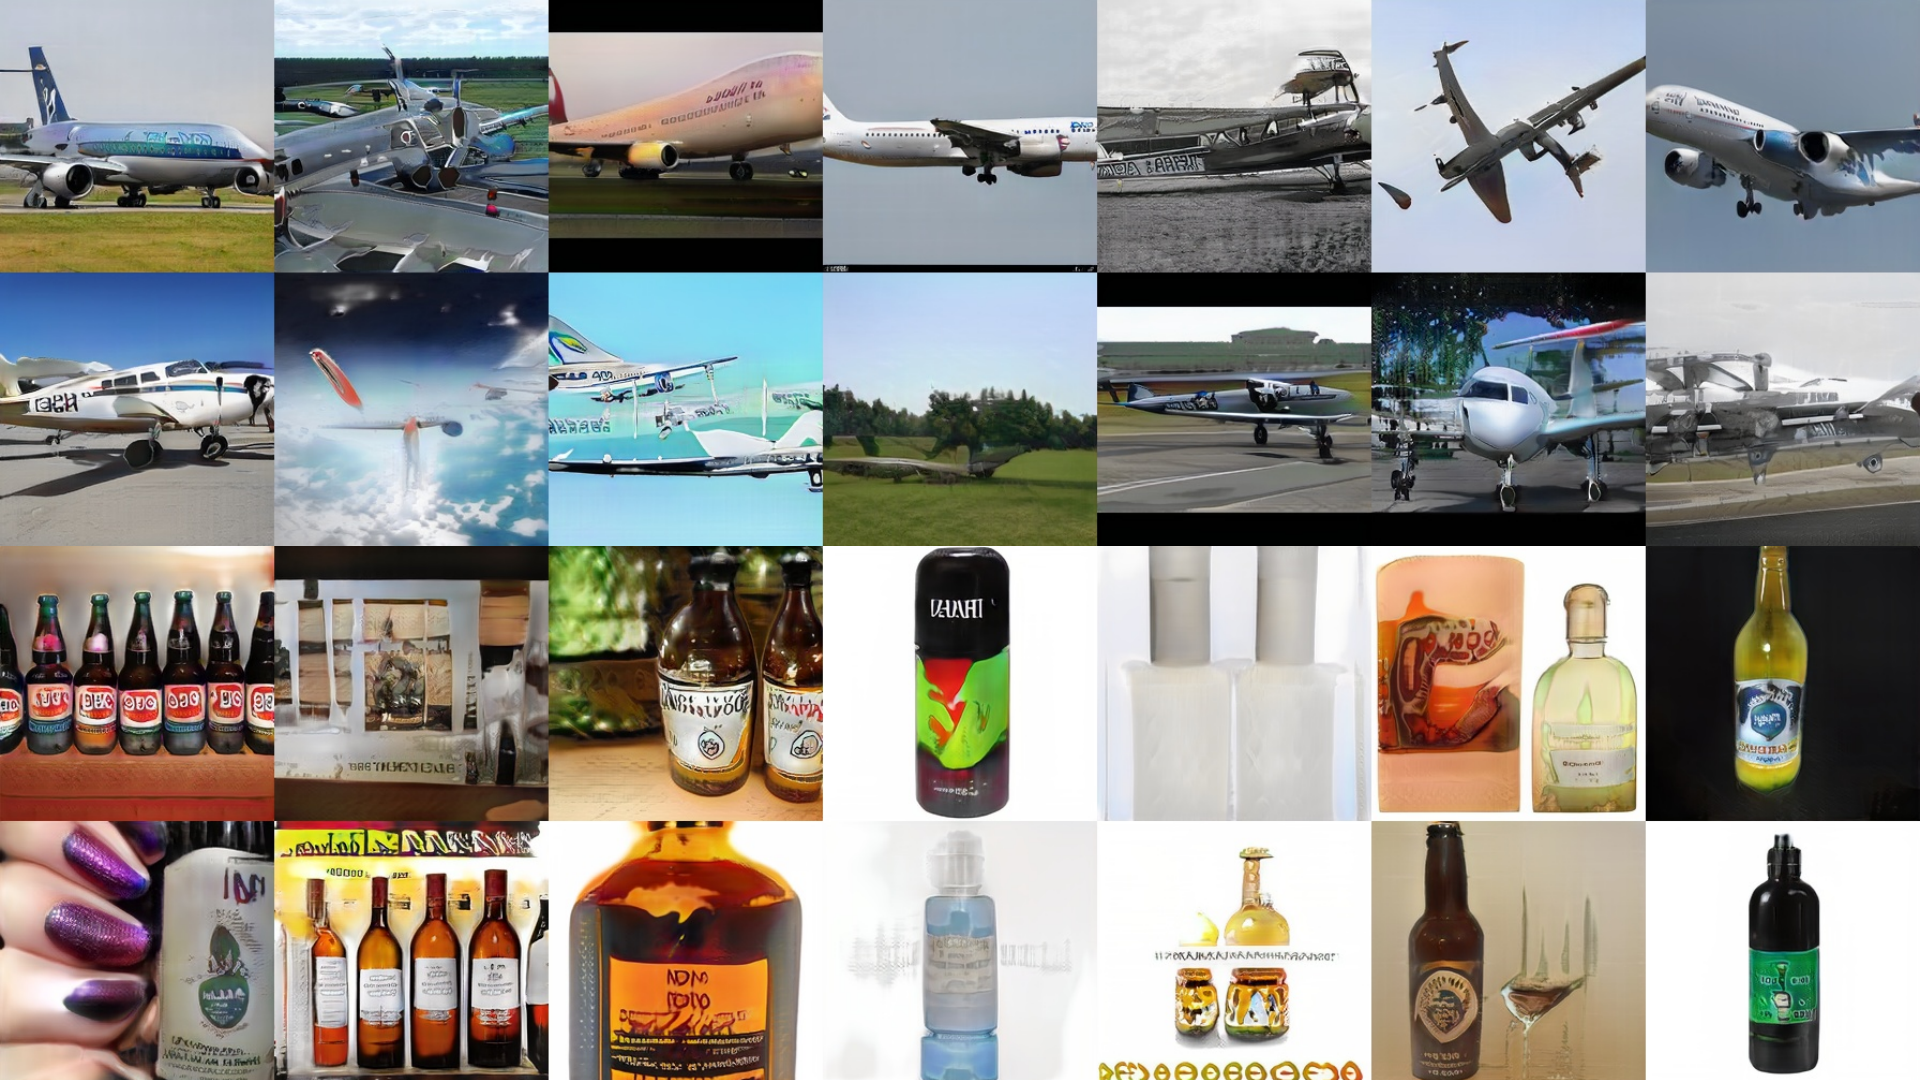

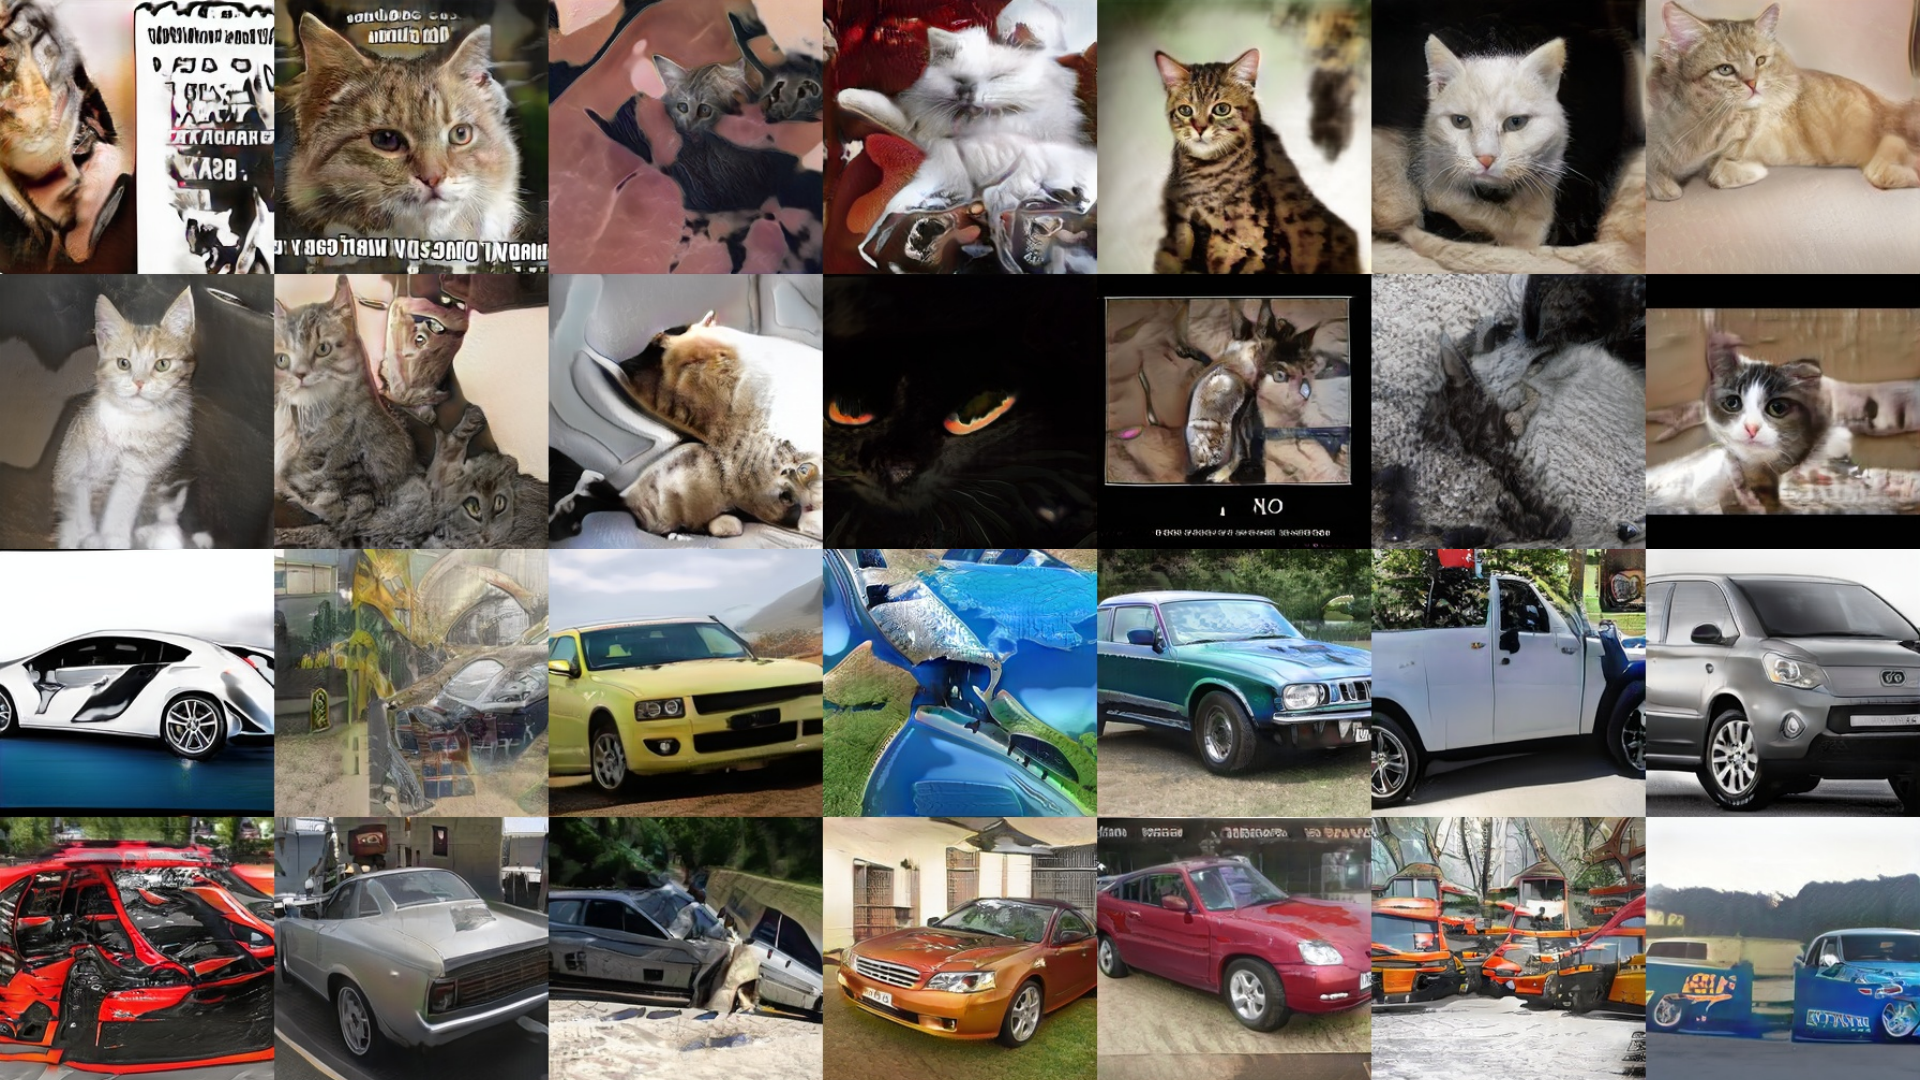

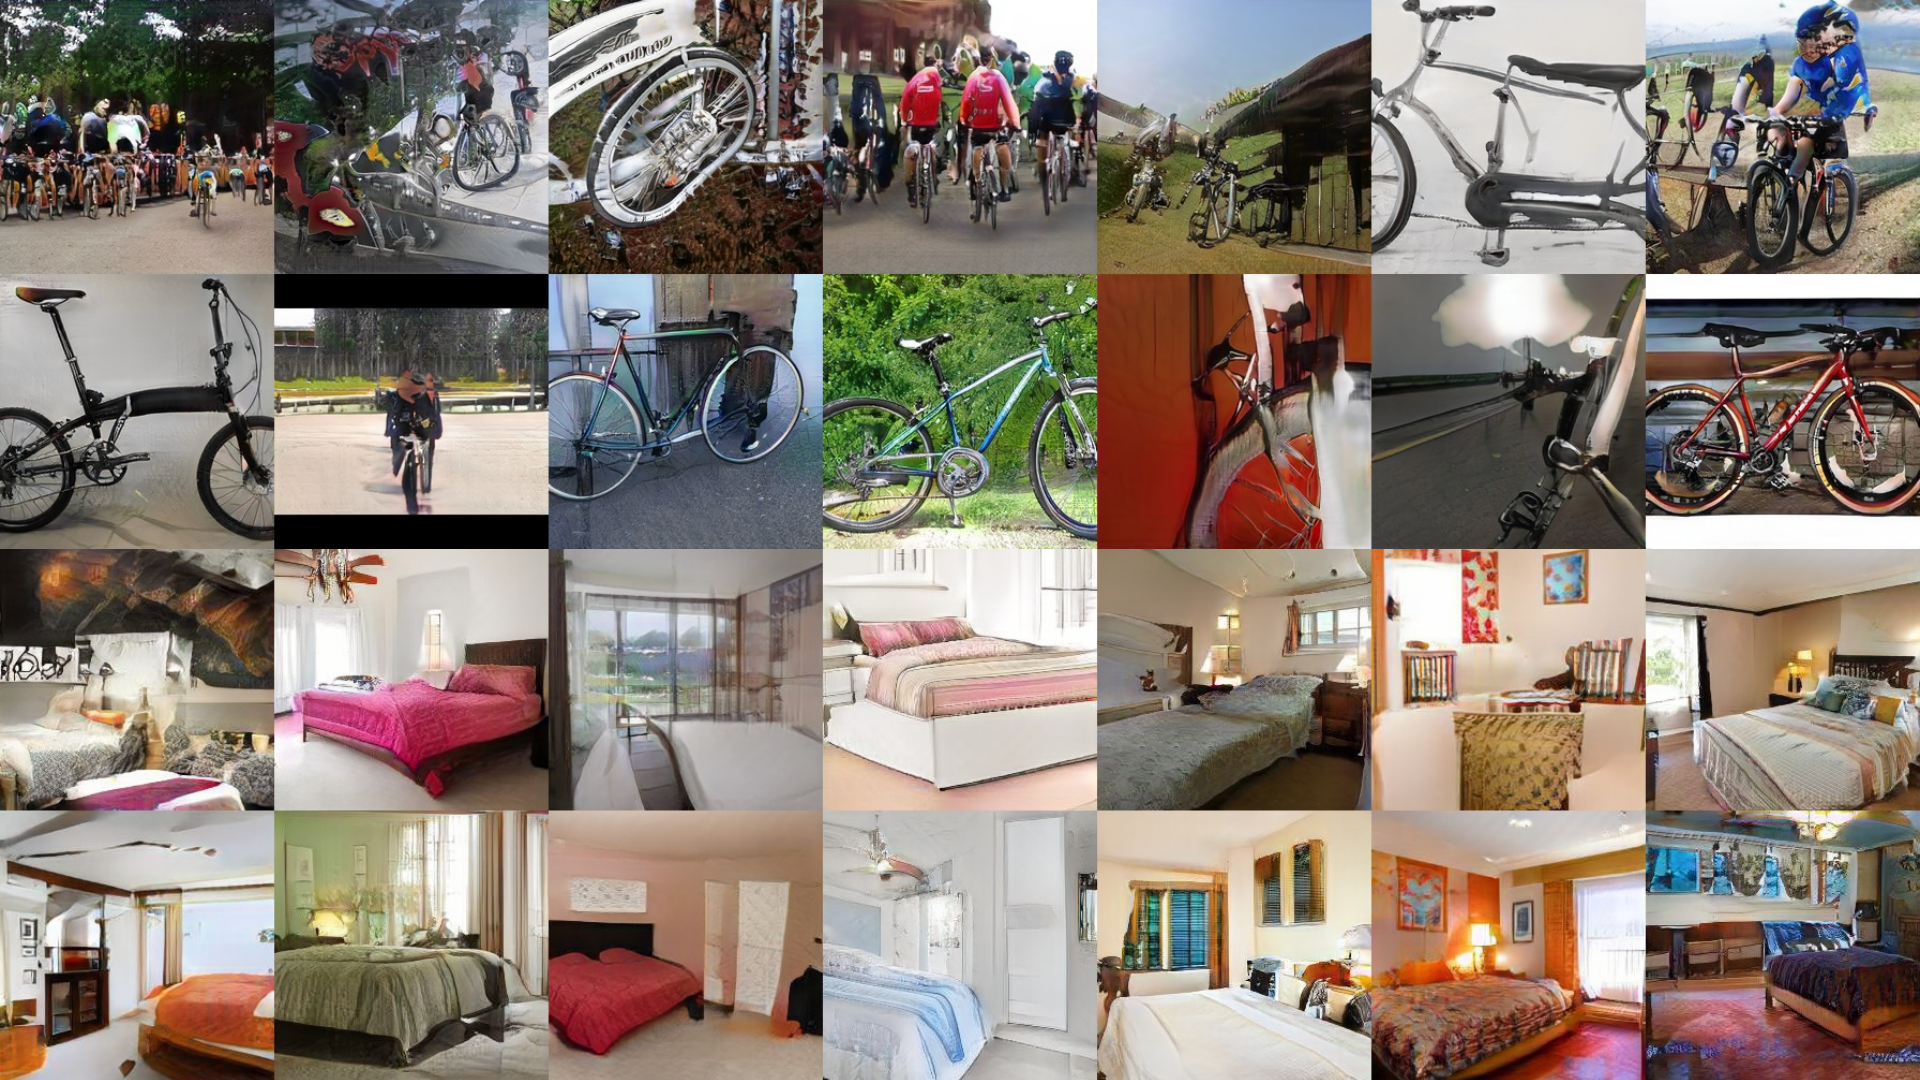

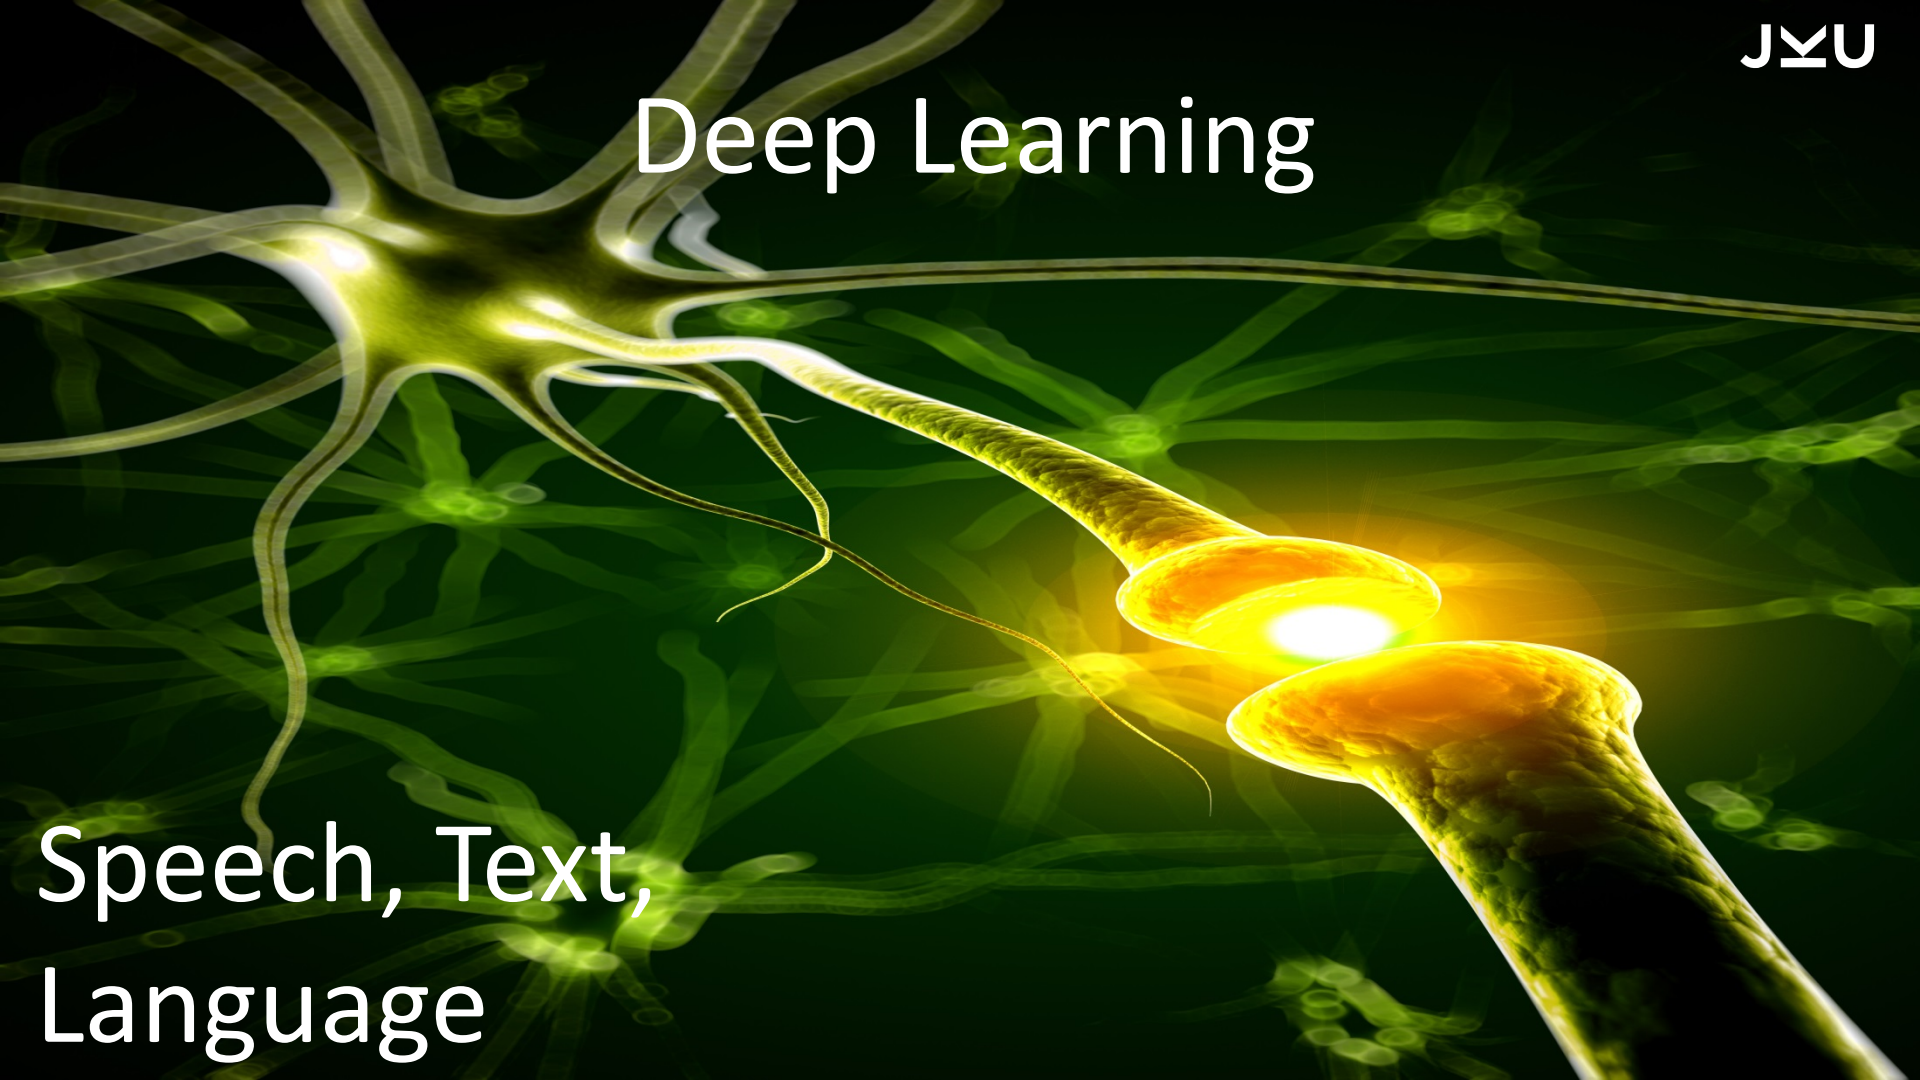
JKU
Deep Learning
Speech, Text,
Language

Long Short-Term Memory

- 1991: invented by Hochreiter
- 1997: publication Hochreiter&Schmidhuber
- 2009: wins French & Arabic handwriting
- 2011: wins offline Chinese handwriting
- 2012: **Google's Android speech recognizer**
- 2015: **Google's Voice transcription**
- 2016: **Apple's iOS 10 → Quicktype**
- 2016: **Google's Translate**
- 2016: **Amazon's Alexa**
- 2017: **Facebook's Translation**

KEY MOMENTS IN DEEP-LEARNING HISTORY 1989-1997

1989
French researcher Yann LeCun, then at Bell Labs, begins foundational work on a type of neural net that becomes crucial for image recognition.

1991
German researchers Sepp Hochreiter and Jürgen Schmidhuber pioneer a neural net with memory features, which eventually proves superior for natural-language processing.

1997
IBM's Deep Blue beats **world champion Garry Kasparov** (right) in chess using traditional AI techniques.

Fortune 28.09.2016, "The AI Revolution: Why Deep Learning is suddenly changing your life"

STAN HONDA—AFP/GETTY IMAGES


facebook
Artificial Intelligence

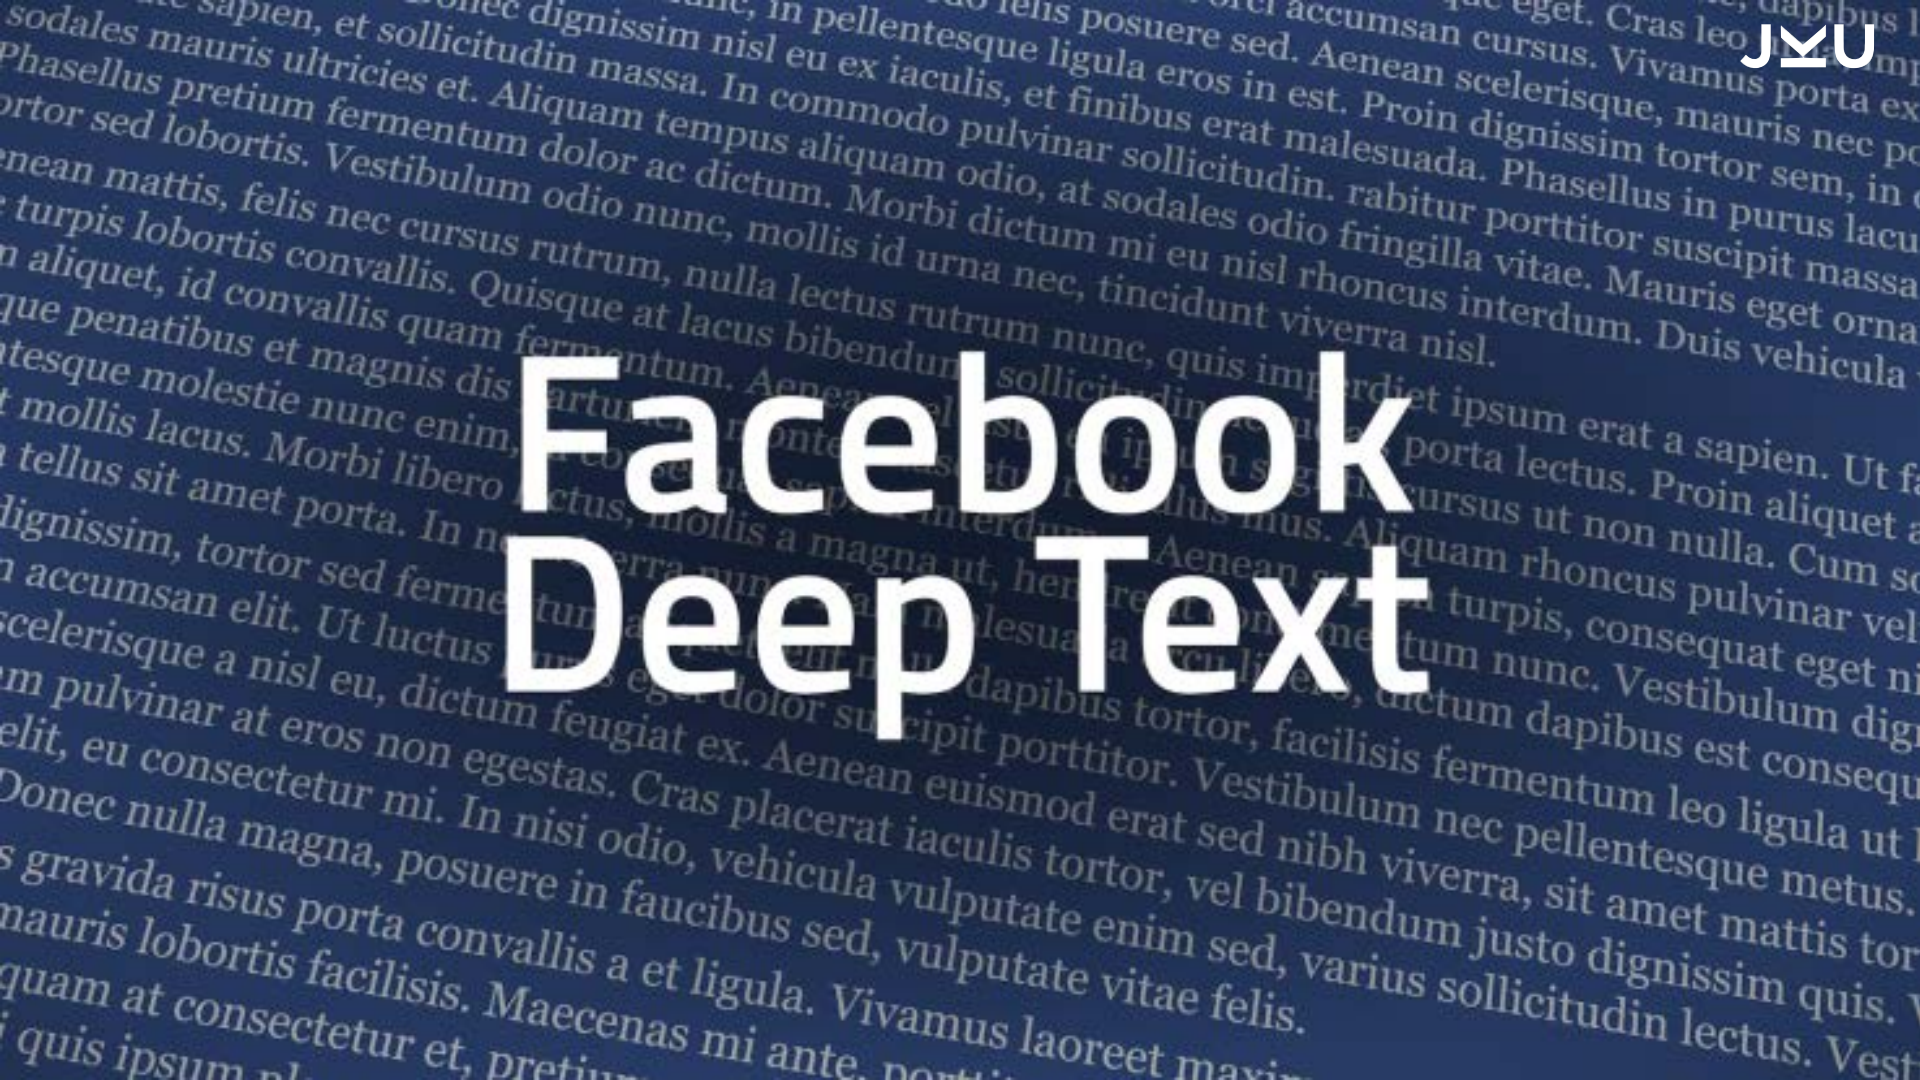
Facebook
Deep Text


JKU
Google
Artificial Intelligence

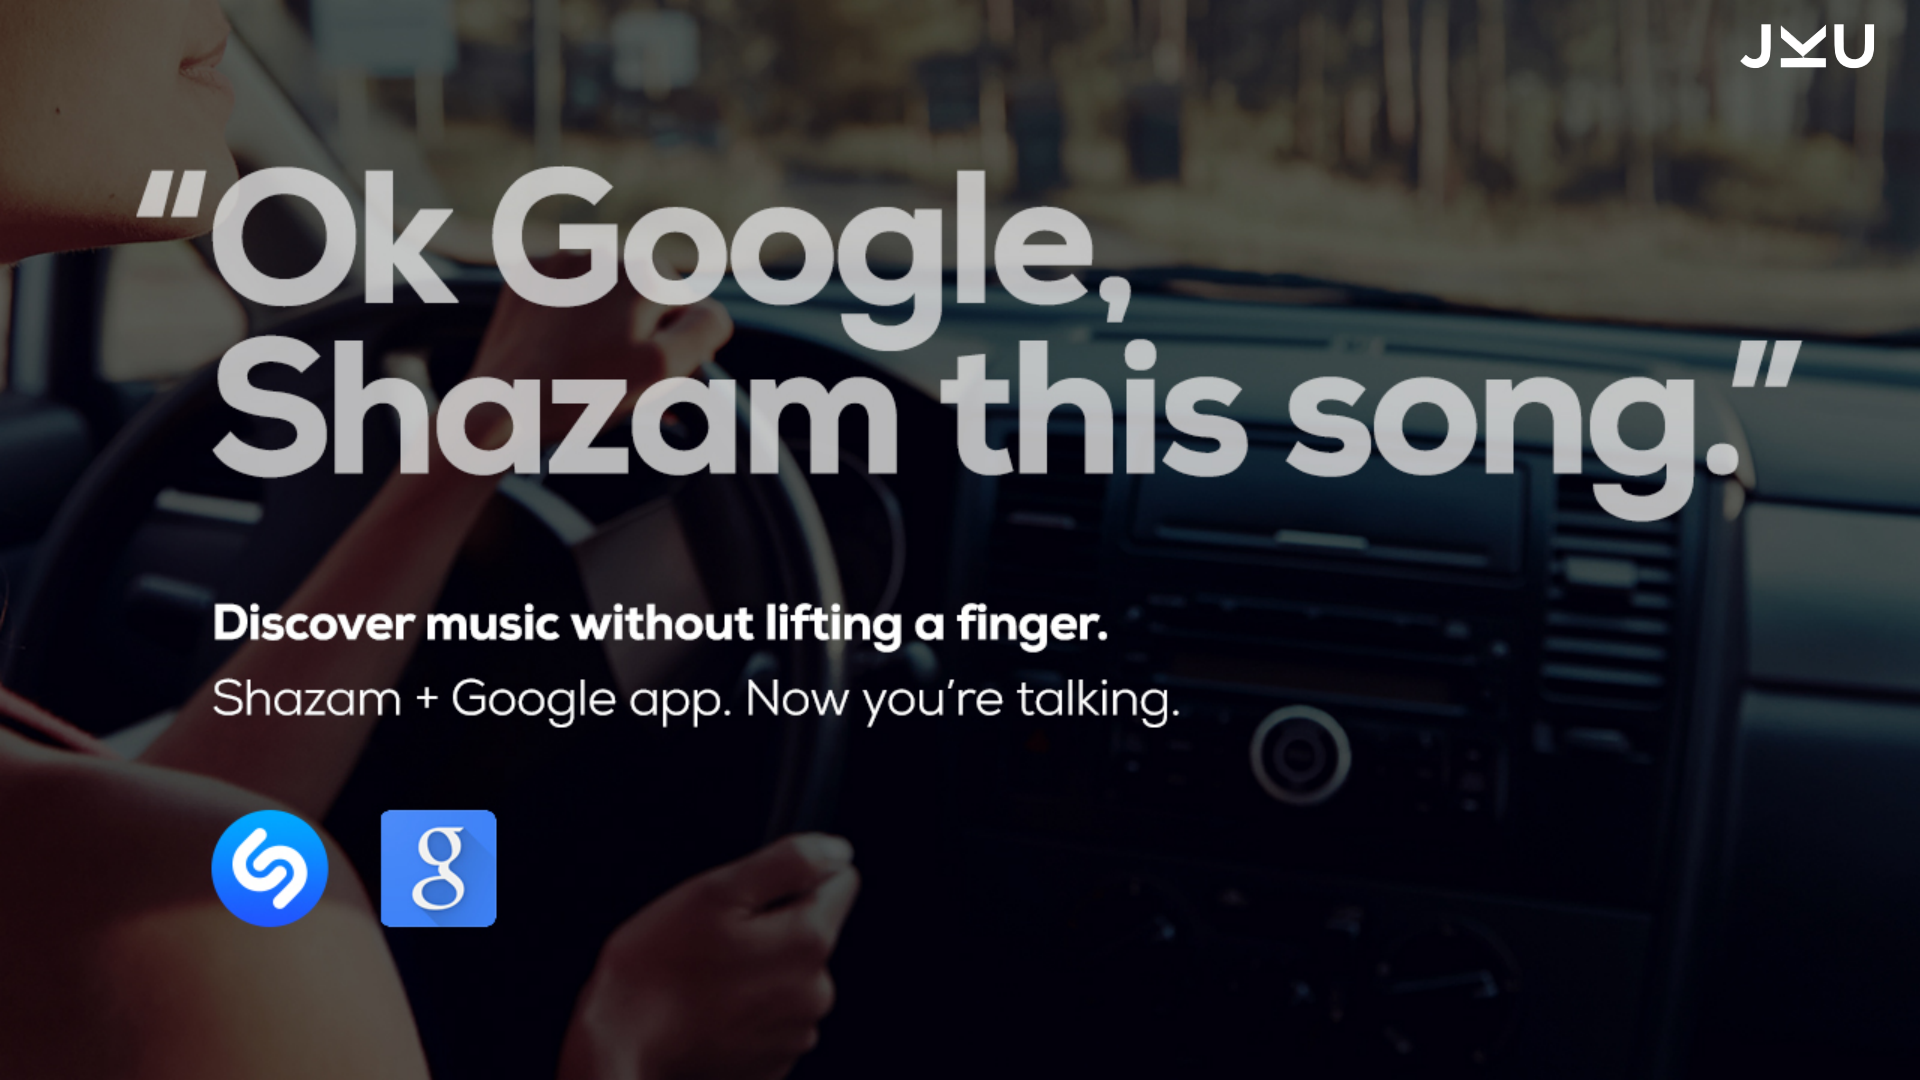
"Ok Google,
Shazam this song."
Discover music without lifting a finger.
Shazam + Google app. Now you're talking.


JKU
Amazon AI
Bringing powerful artificial intelligence to all developers


amazon
“Alexa....please tell Baxter the Robot to machine this part.”


Siri


JKU
“Hey Siri, what’s the best sushi place in town?”

Long Short-Term Memory

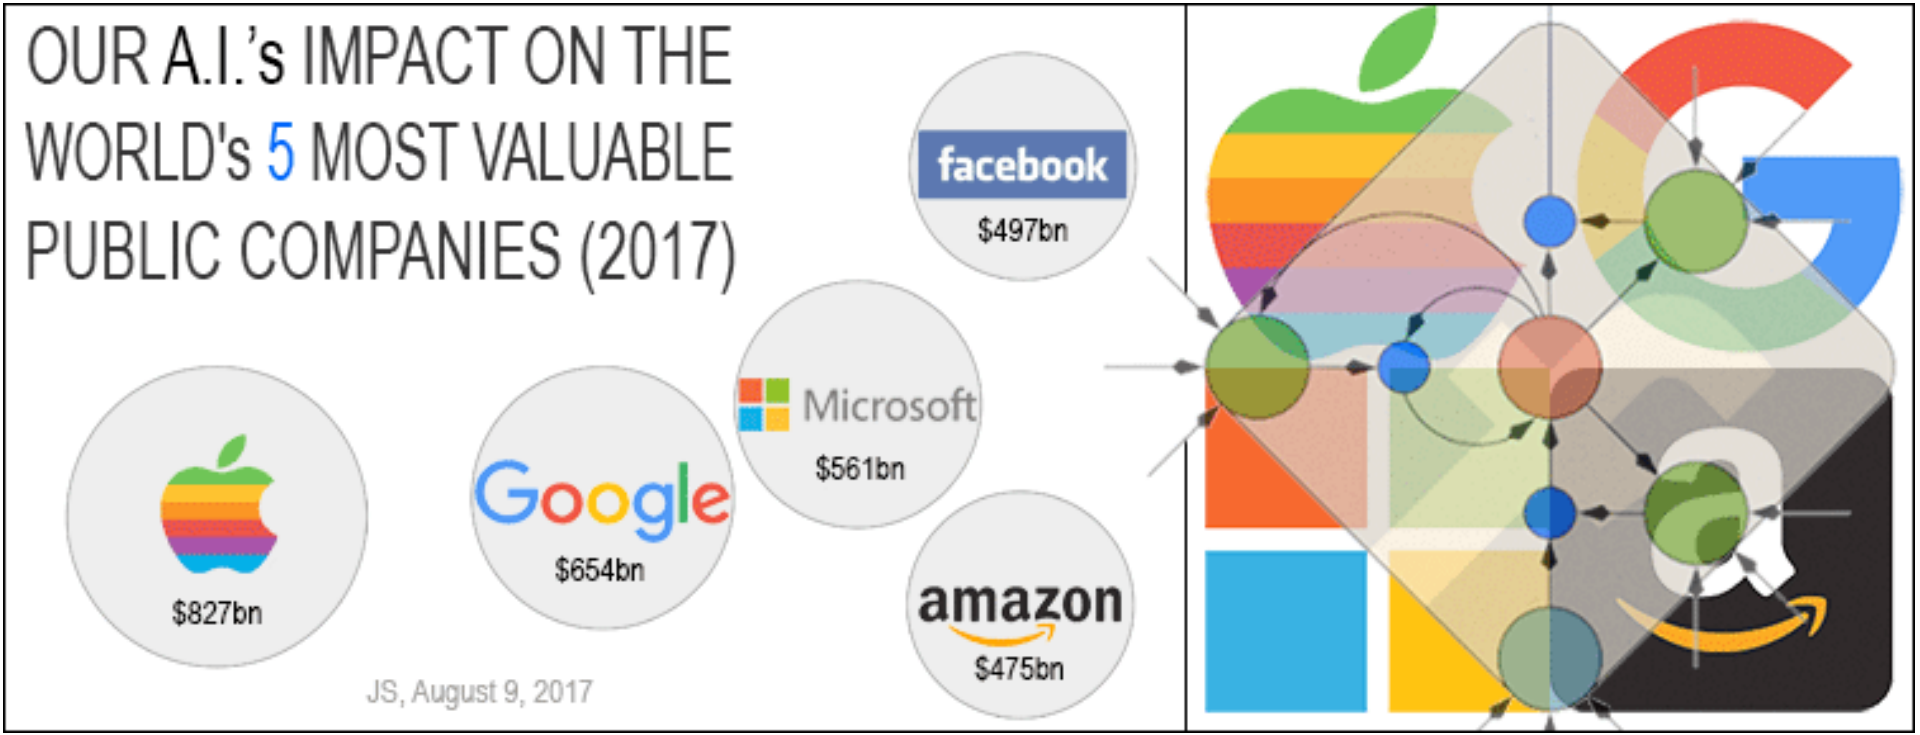

From Jürgen Schmidhuber: http://people.idsia.ch/~juergen/impact-on-most-valuable-companies.html

Self-Normalizing Networks (2017)

1 bn $

Self-Normalizing Networks (2017)

From: Zhen Huang <zhen_huang@apple.com>
Subject: Thank you for inventing SELU.
Date: **Tue, 19 Nov 2019 14:19:36 -0500**
To: hochreit@bioinf.jku.at, klambauer@bioinf.jku.at, unterthiner@bioinf.jku.at

Dear Authors of SELU,

I want to express my sincere thanks to you for inventing SELU.

As in the paper:
[arxiv.org]https://arxiv.org/abs/1910.01992
SELU helps us deploy the 50 layer CNN for Siri speech recognition.

I can not represent Apple, but as the first author of the paper, I personally sincerely appreciate the great contributions you made to the field of machine learning!

Thanks again!

Zhen Huang
Ph.D., Senior machine learning scientist in Siri

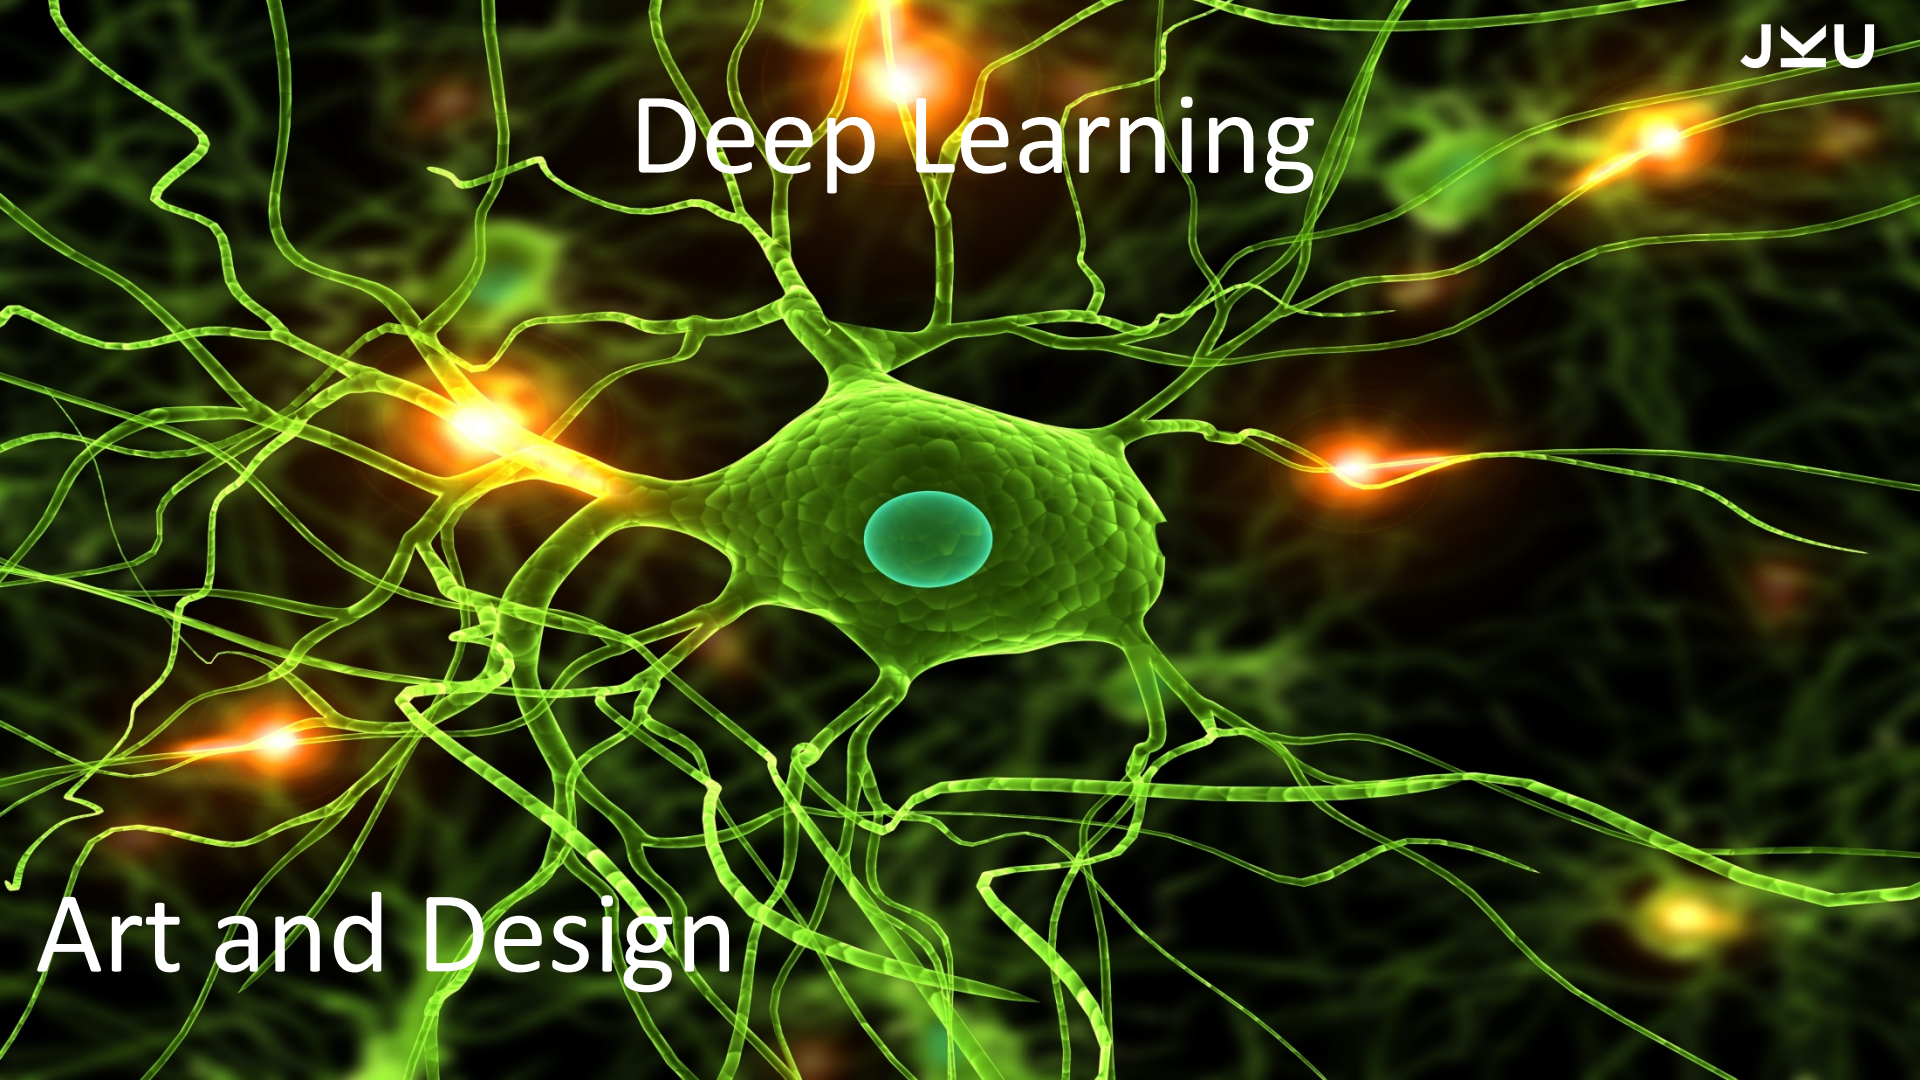
JKU
Deep Learning
Art and Design


A
B
C
D
E
F
JKU

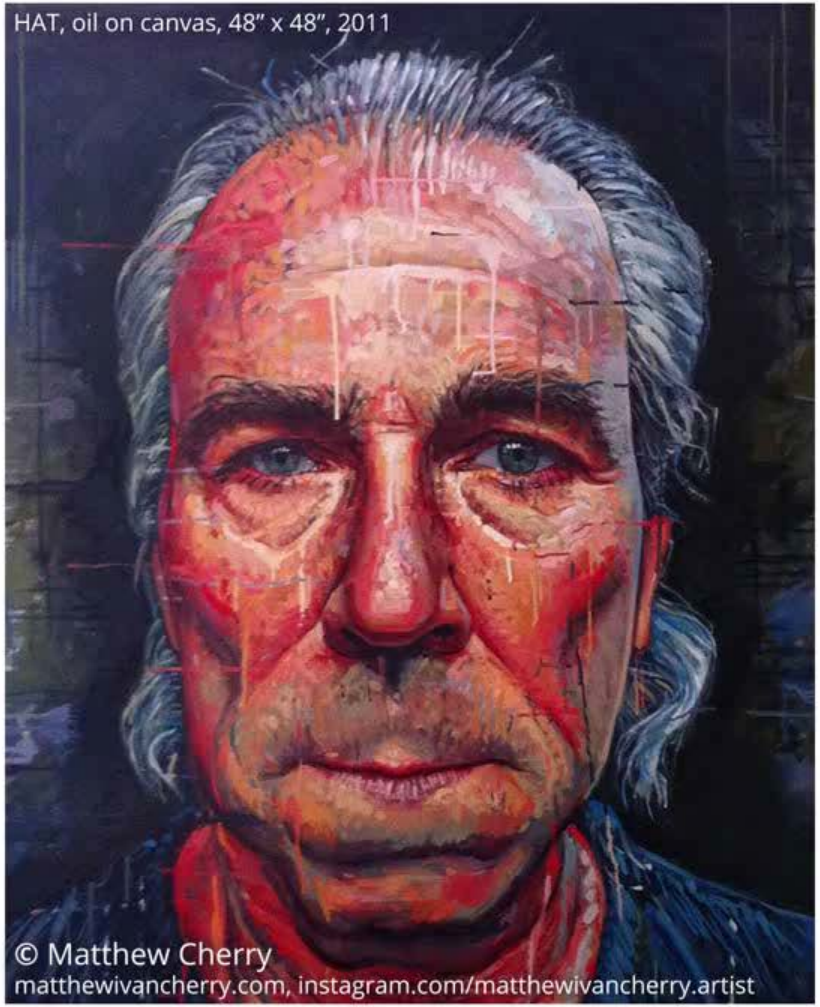
HAT, oil on canvas, 48” x 48”, 2011
© Matthew Cherry
matthewivancherry.com, instagram.com/matthewivancherry.artist

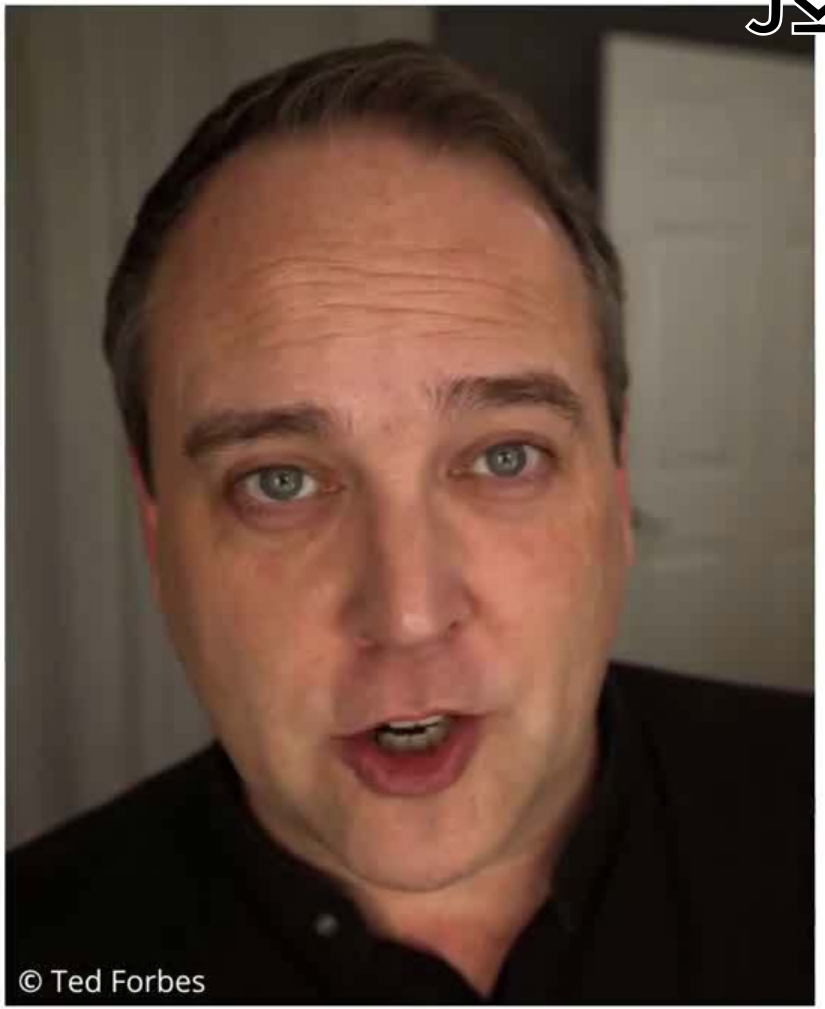
© Ted Forbes

JKU

Input	Ground truth	Output	Input	Ground truth	Output
					
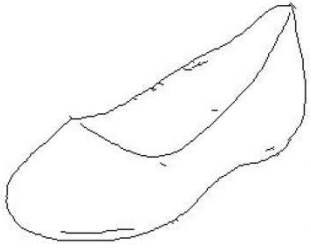					
			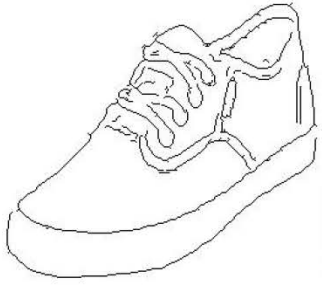		

Input	Ground truth	Output	Input	Ground truth	Output
			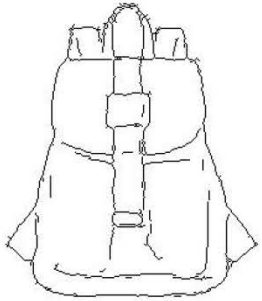		
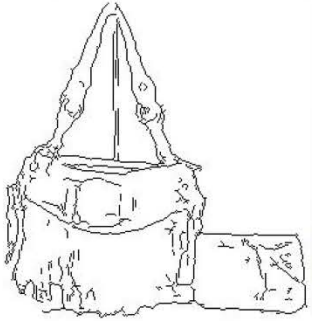			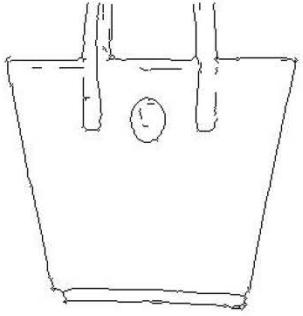		
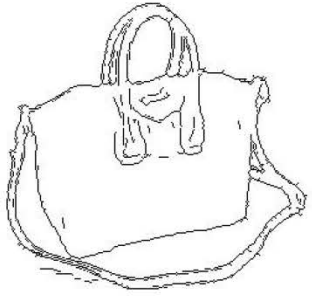					

"Ode to Joy" harmonized in the style learned from:

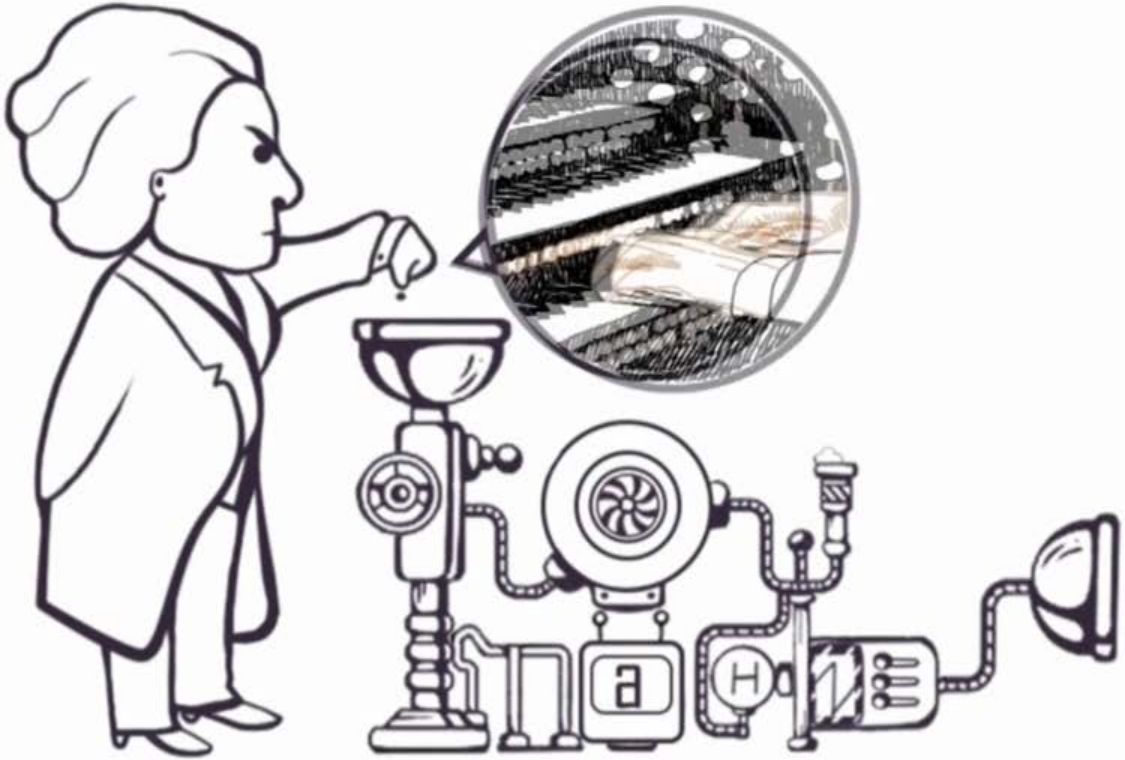

Instrument to Instrument

Sample #1

Source Piano Sample

Generated Harpsichord Sample

LSTM composed music piece

Deep Learning Against COVID19

Coronavirus at Different Scales

		ML/AI landscape
population scale		Tracking, epidemiological modeling, ...
macroscopic scale		Healthcare management, patient surveillance,...
microscopic scale	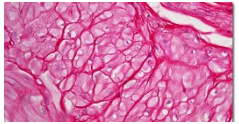	Diagnosis, ...
molecular scale	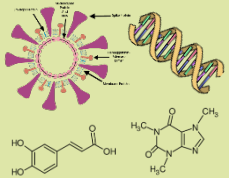	**Drug discovery,** genome sequencing, ...

Drug Discovery Pipeline

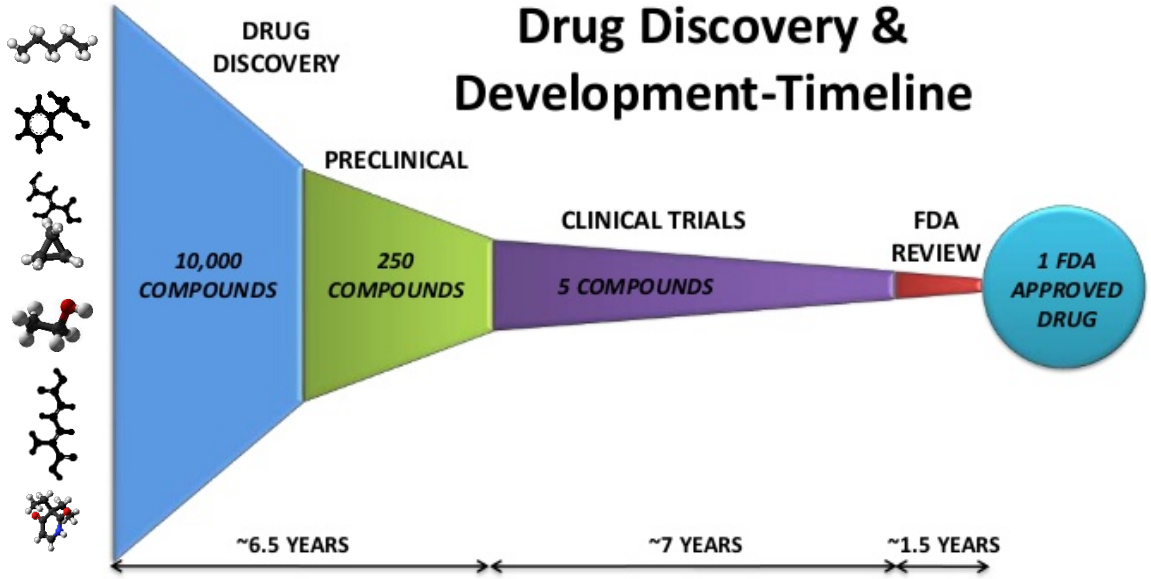

Inhibiting the SARS-Cov-2 Virus

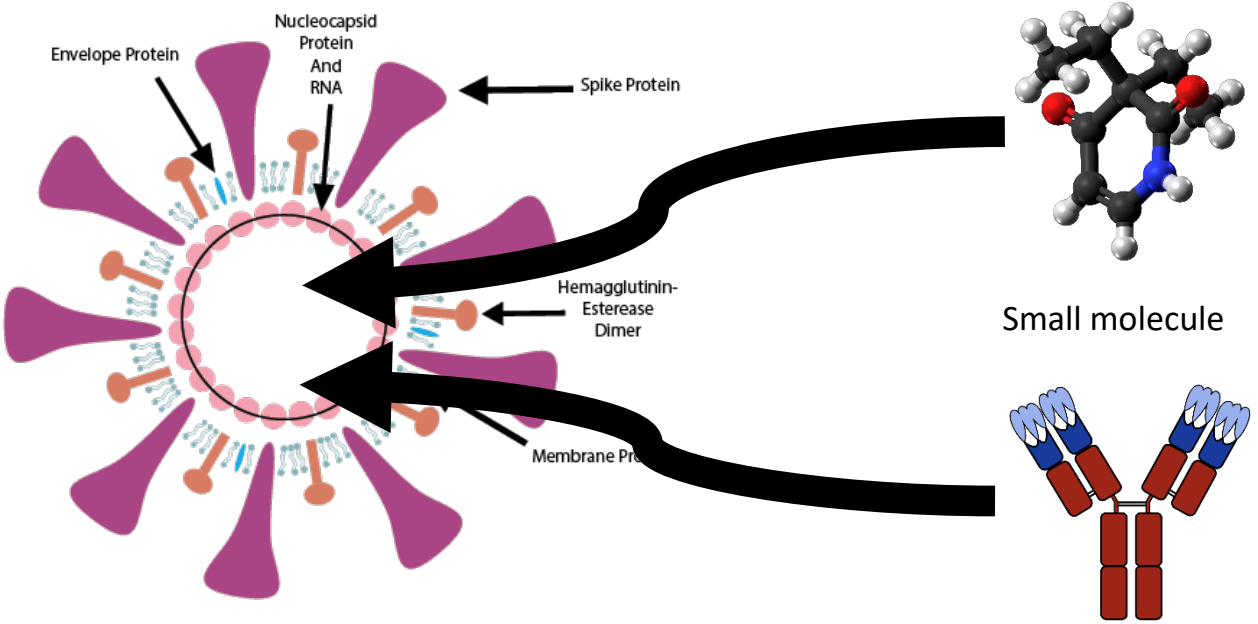

Small molecule

Schematic representation of the Corona virus

Antibodies

Computer-Aided Drug Discovery

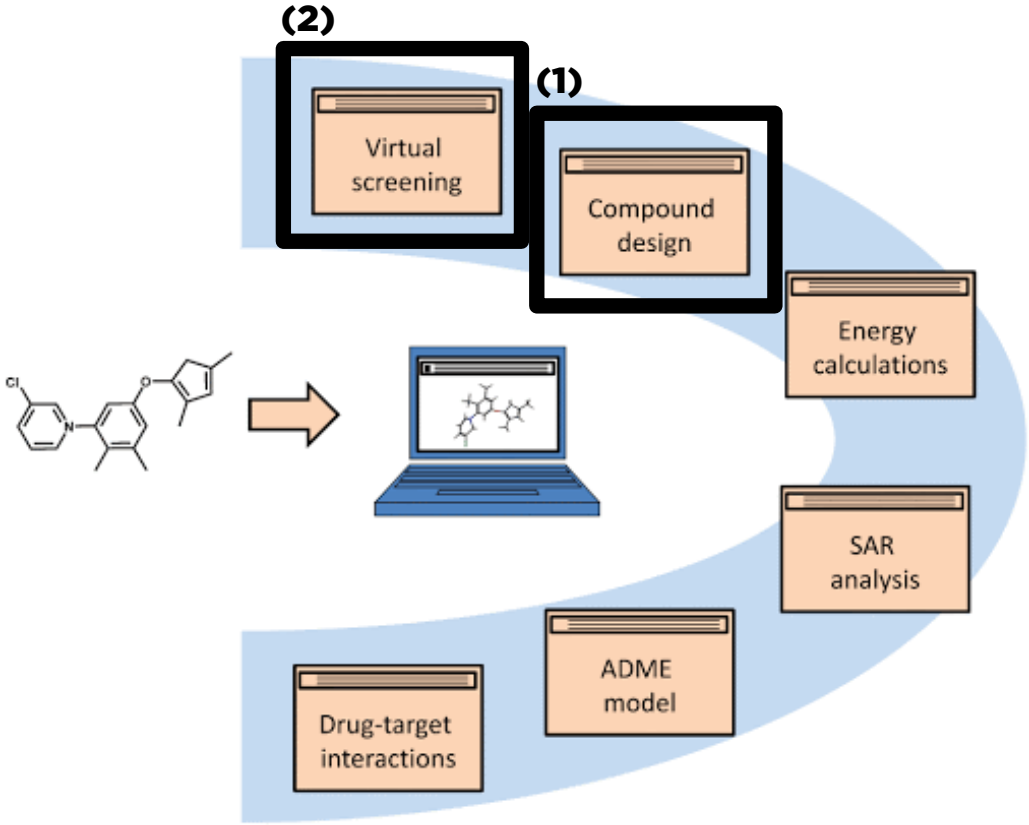

Compound Design Using ML

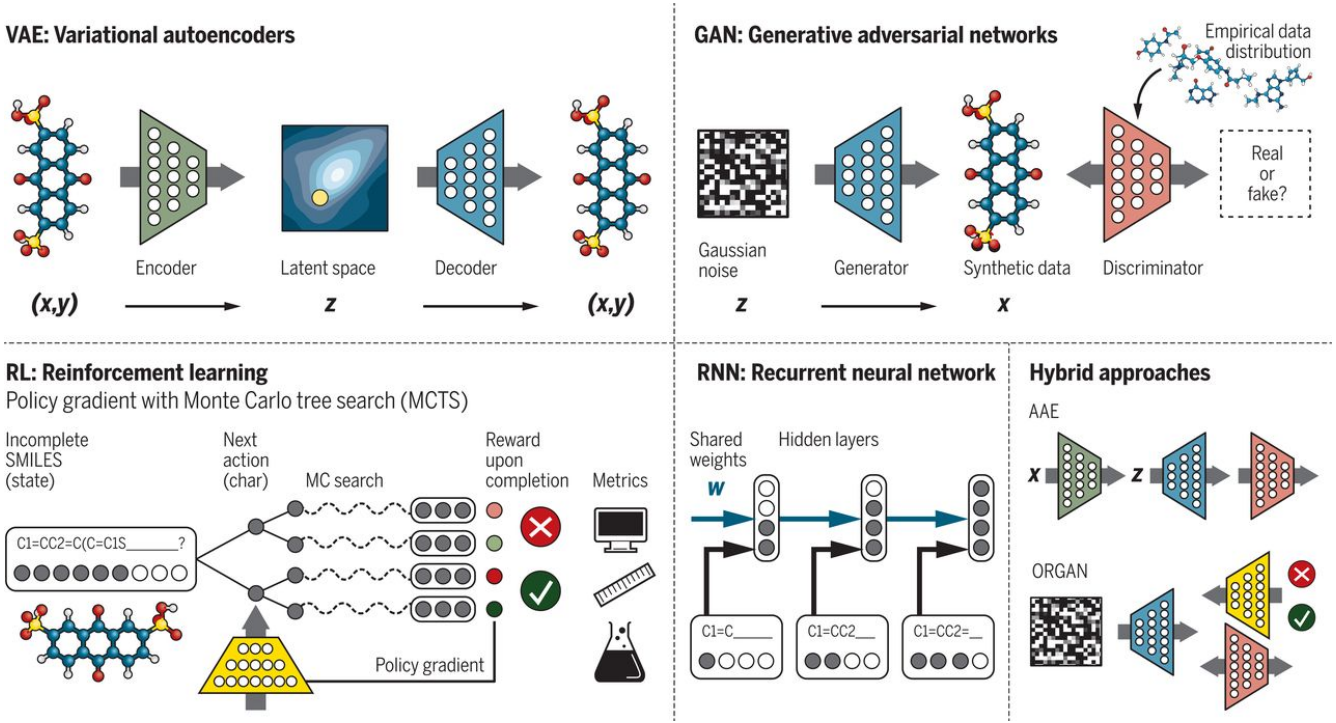

Drug Target Prediction

- **Ligand-based**
  - Substructure search
  - Similarity search
  - Model-based search
- **Structure-based**
  - Docking
- Hybrid approaches
  - Proteochemometrics models

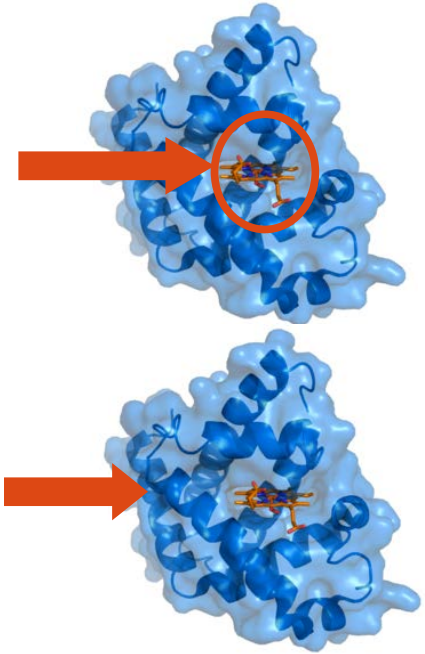

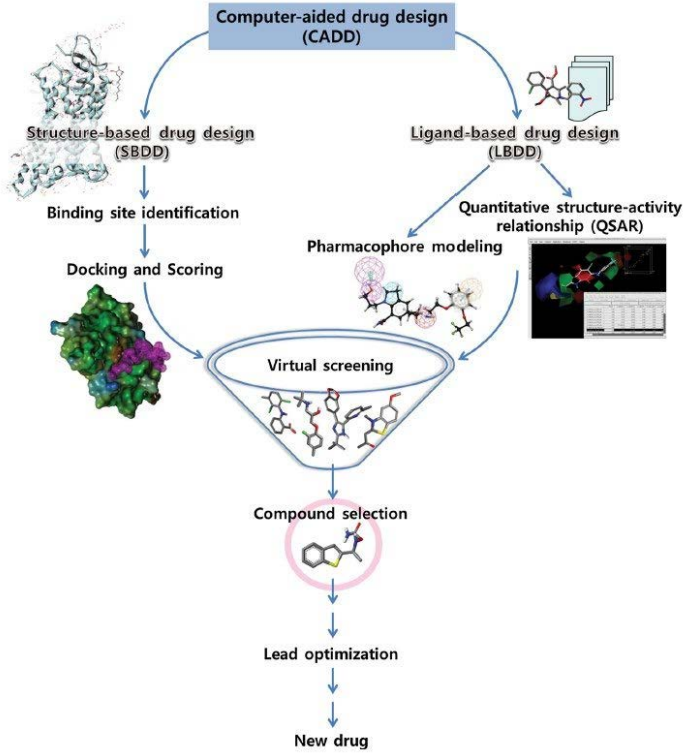

Drug target prediction

$f_{\boldsymbol{\theta}}($ 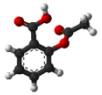 $) =$ 

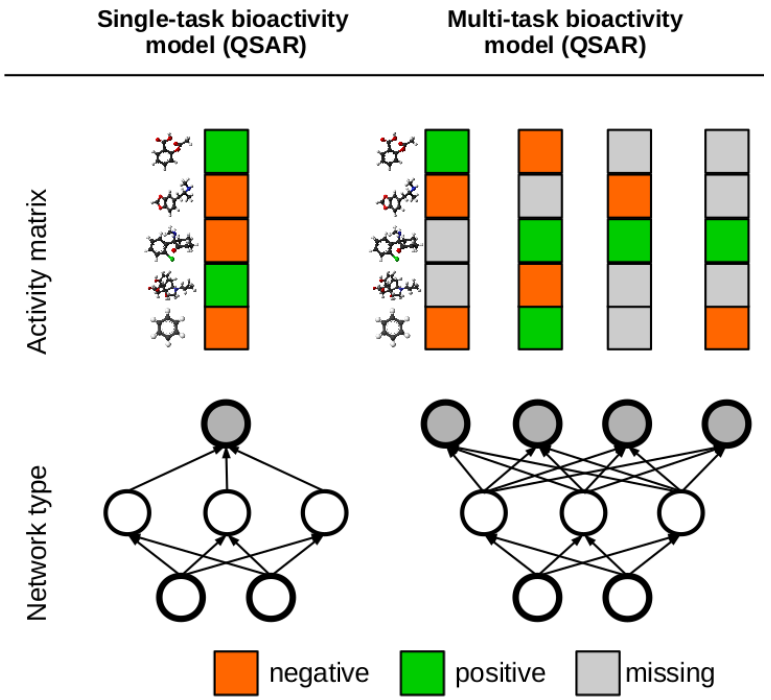

Typically:
>99% missing

Large Databases for Learning

- Large databases of small molecules are available
- **Triplets**: molecule – target (assay) – activity
- Databases

PubChem	268 million triplets
ChEMBL	16 million triplets
DRUGBANK	13,000 drugs
ZINC DATABASE	~1 billion molecules

Screening with Trained Models

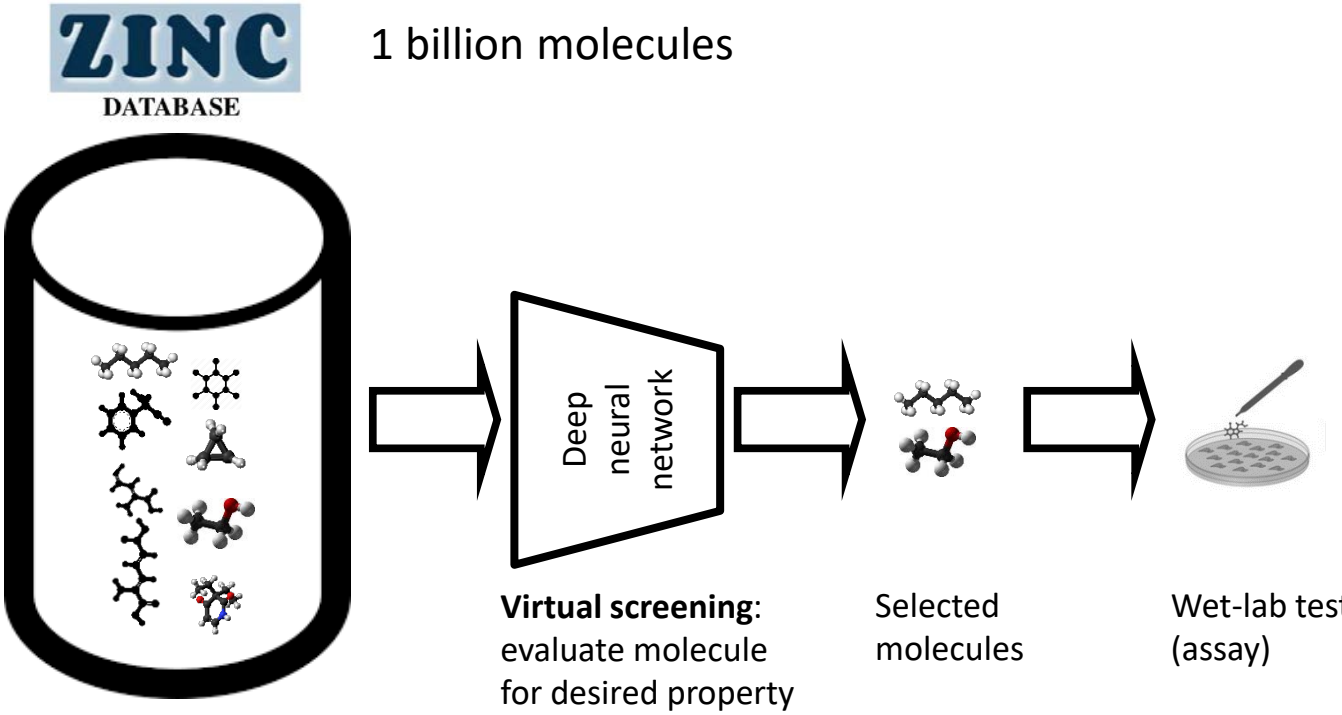

SARS-Cov Inhibitors

- Deep neural network with SMILES as input
  Type: SmilesLSTM (Mayr et al., 2018)
- Trained on 3.6 M molecules
- ~220M drug-assay-activity triplets (labels)
- Multi-task network with ~6,000 output unit
- 4 output units related to SARS-Cov

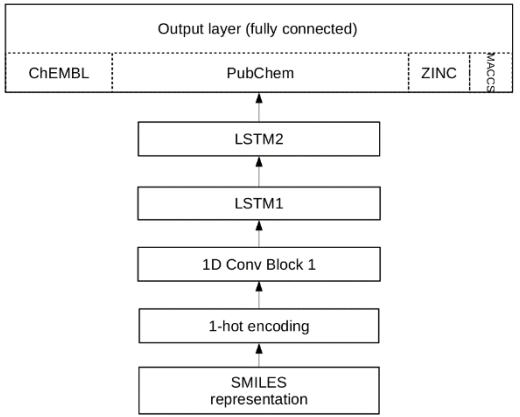

Assay ID	Source	#inact	#act	Description
1706	PubChem	193637	269	QFRET-based assay for SARS-CoV 3C-like Protease
1879	PubChem	167	86	QFRET-based assay for SARS-CoV 3C-like Protease (confirmation)
485353	PubChem	215030	390	Yeast-based Assay for SARS-CoV PLP
652038	PubChem	493	135	Yeast-based Assay for SARS-CoV PLP (validation)

Mayr, A., Klambauer, G., Unterthiner, T., Steijaert, M., Wegner, J. K., Ceulemans, H., ... & Hochreiter, S. (2018). Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. Chemical science, 9(24), 5441-5451.

SARS-Cov Inhibitors

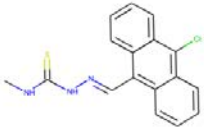

ZINC000254565785

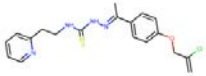

ZINC000726422572

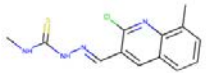

ZINC000916265995

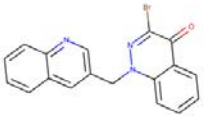

ZINC000806591744

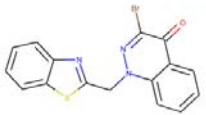

ZINC000178971373

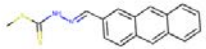

ZINC000000155607

- Top-ranked molecules (Figure)
- List of 30,000 top-ranked molecules provided to scientific community via github

Hofmarcher, M., Mayr, A., Rumetshofer, E., Ruch, P., Renz, P., Schimunek, J., ... & Klambauer, G. (2020). Large-Scale Ligand-Based Virtual Screening for SARS-CoV-2 Inhibitors Using Deep Neural Networks. Available at SSRN 3561442.

Drug Repurposing: Screening of Drugs

DRUGBANK

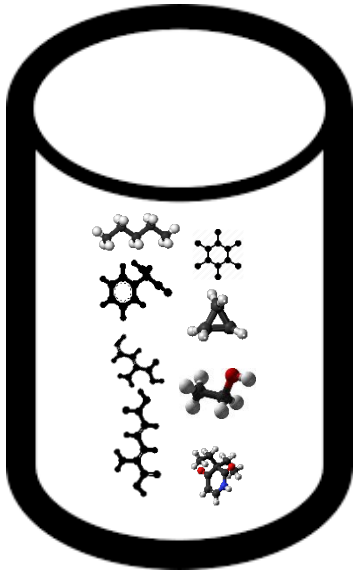

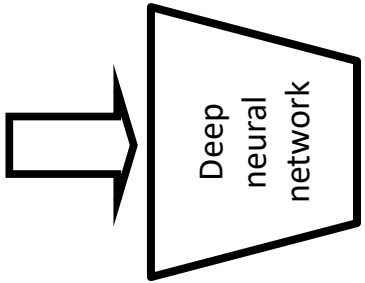

Virtual screening:
evaluate molecule
for desired property

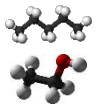

Selected
molecules

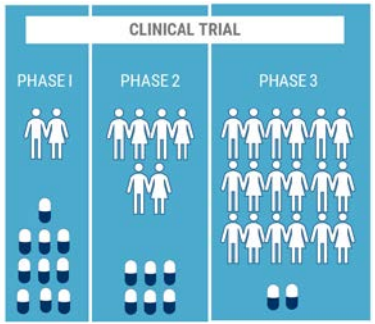

Reinforcement Learning
The New Frontier

Reinforcement Learning

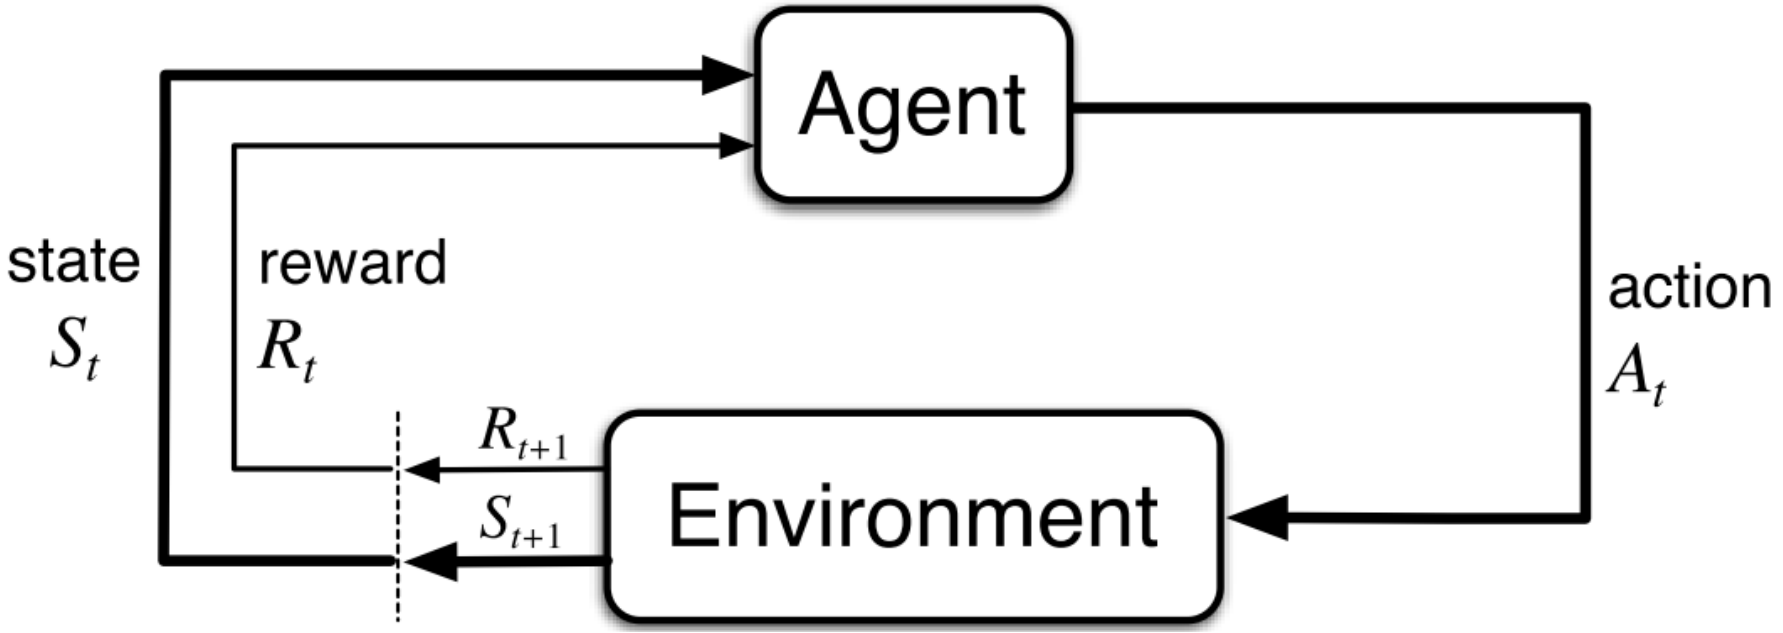

Reinforcement Learning

Model-based reinforcement learning

- Model given → optimal control
- Model parameters learned → adaptive optimal control
- Model learned → model-based reinforcement learning

Reinforcement Learning

Model-based reinforcement learning

- Model generates samples
- Model used for planning
- Model propagated through

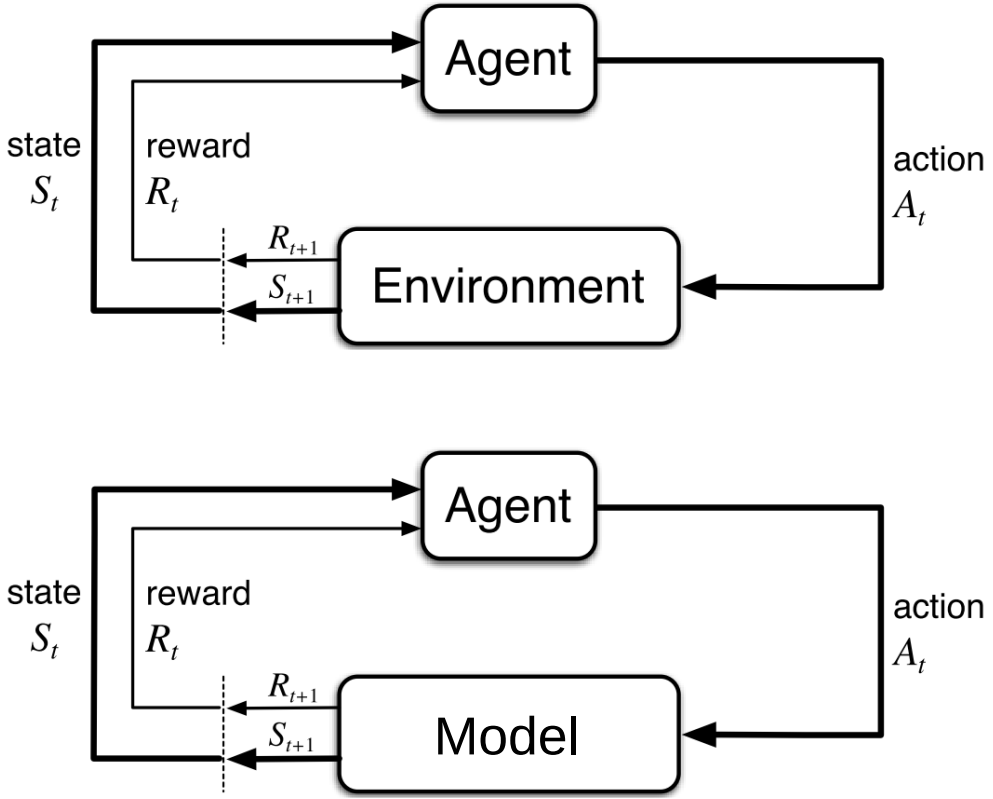

Reinforcement Learning

Model-based:

Monte Carlo Tree Search

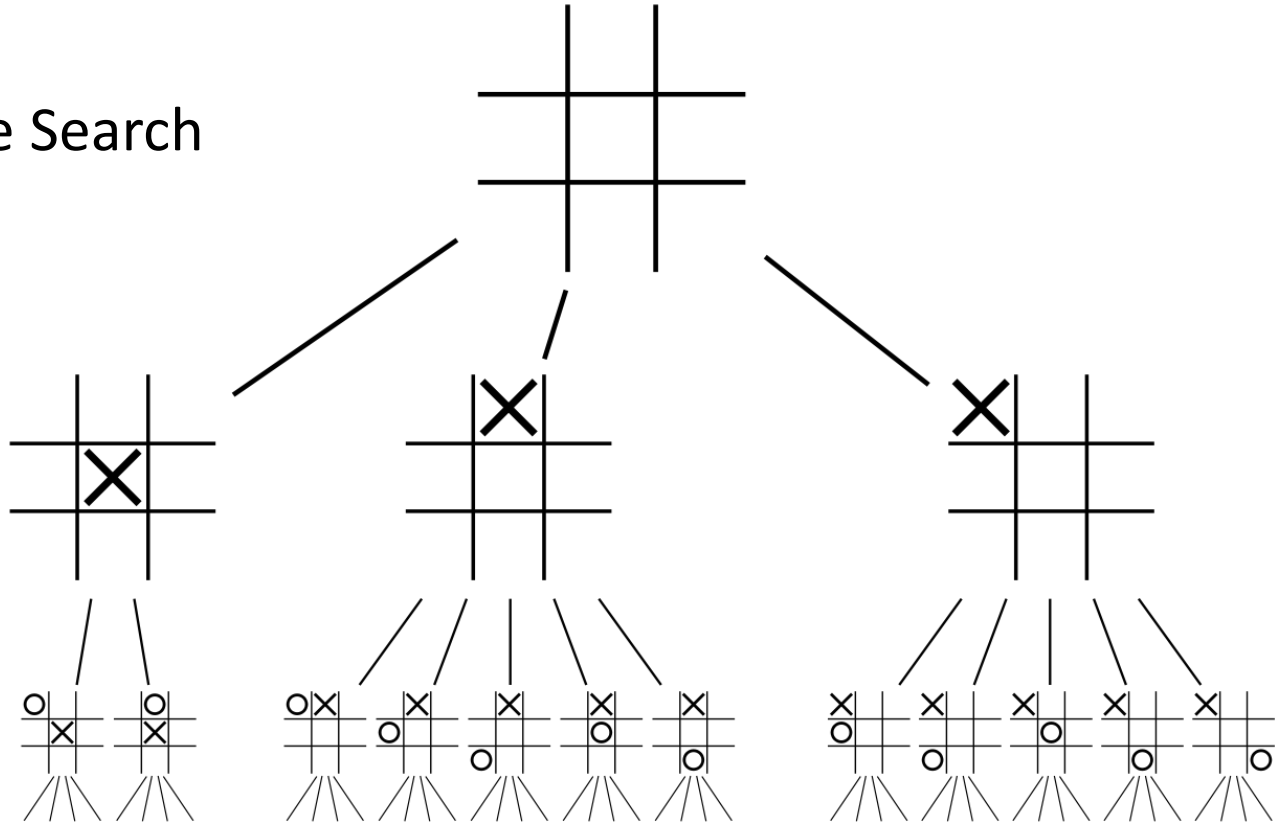

Reinforcement Learning

Model-based: Monte Carlo Tree Search

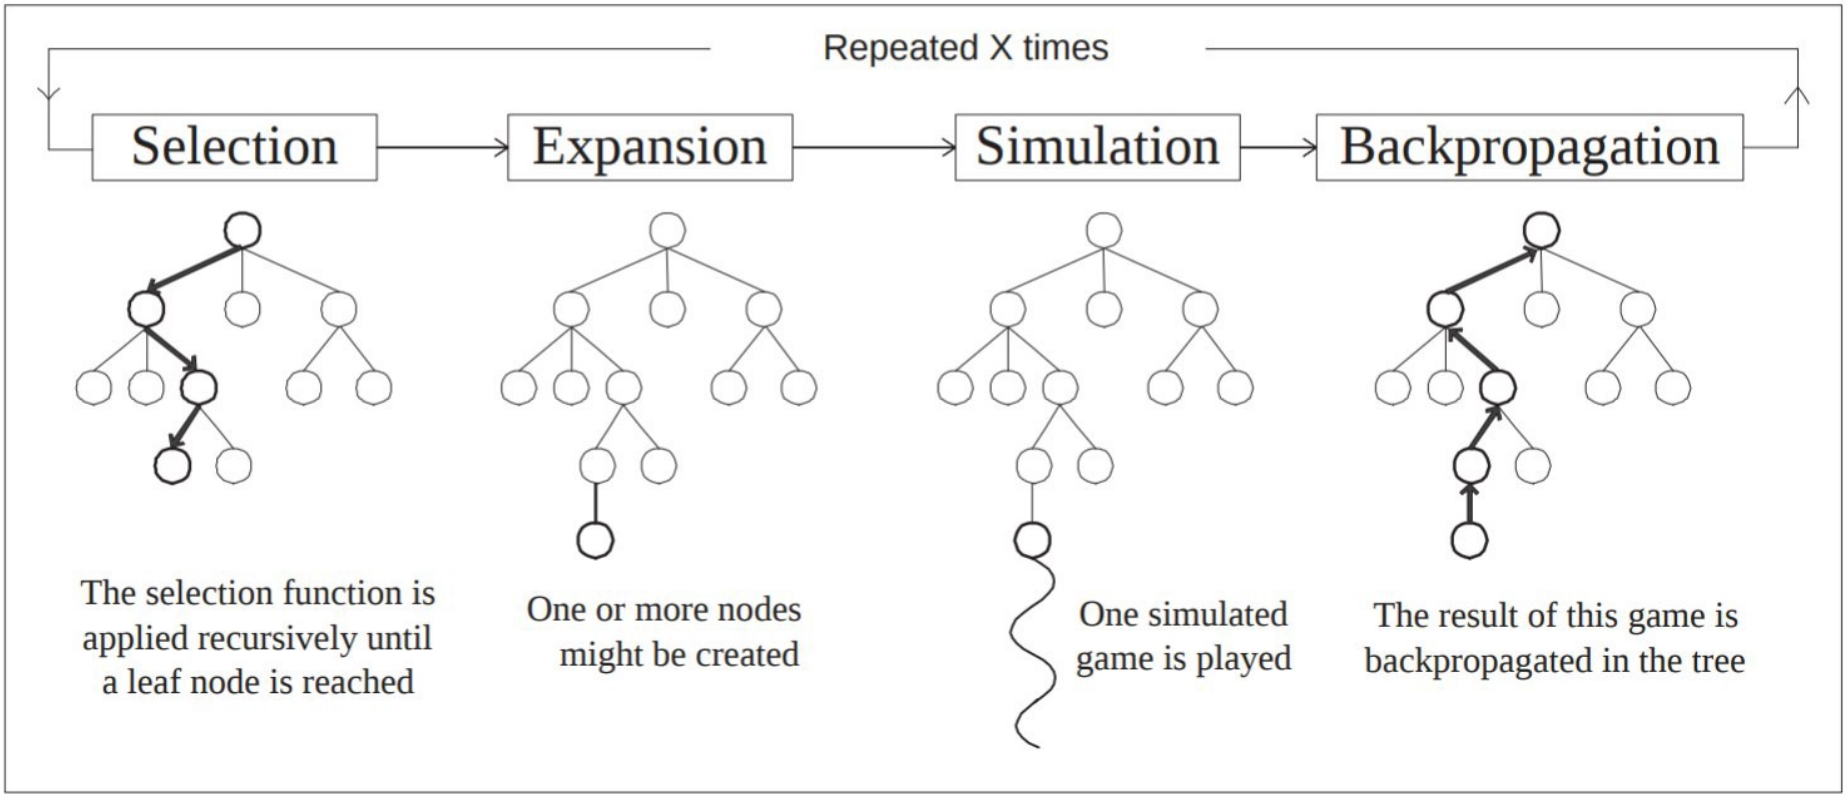

AlphaGo Beats Lee Sedol

DeepMind's AlphaGo beats one of the best Go players

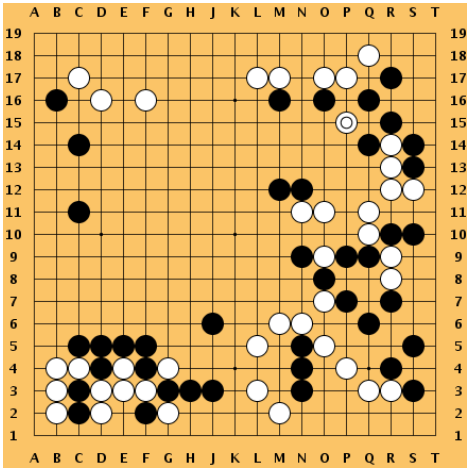

December 2017:
Alpha Zero Teaches Itself Chess 4 Hours
→ beats Stockfish

Reinforcement Learning

Model-free reinforcement learning in real world

- self-driving cars
- robots
- optimization of traffic
- smart city
- environmental optimization (air pollution)
- climate change


DOTA 2
OpenAI


StarCraft II
WINGS OF LIBERTY
JKU
StarCraft II
DeepMind

Nature
14. November 2019

DeepMind's AlphaStar is big LSTM network that plays starcraft.

AlphaStar was rated at Grandmaster level for all three StarCraft races and above 99.8% of officially ranked human players.

RUDDER

Delayed Rewards

Current weakness of reinforcement learning: strategic decisions.

Strategic decisions lead to **delayed rewards**:

- actions cause reward that is obtained much later
- many distracting rewards
- credit assignment problem: what action was responsible

Problem

probabilistic transitions

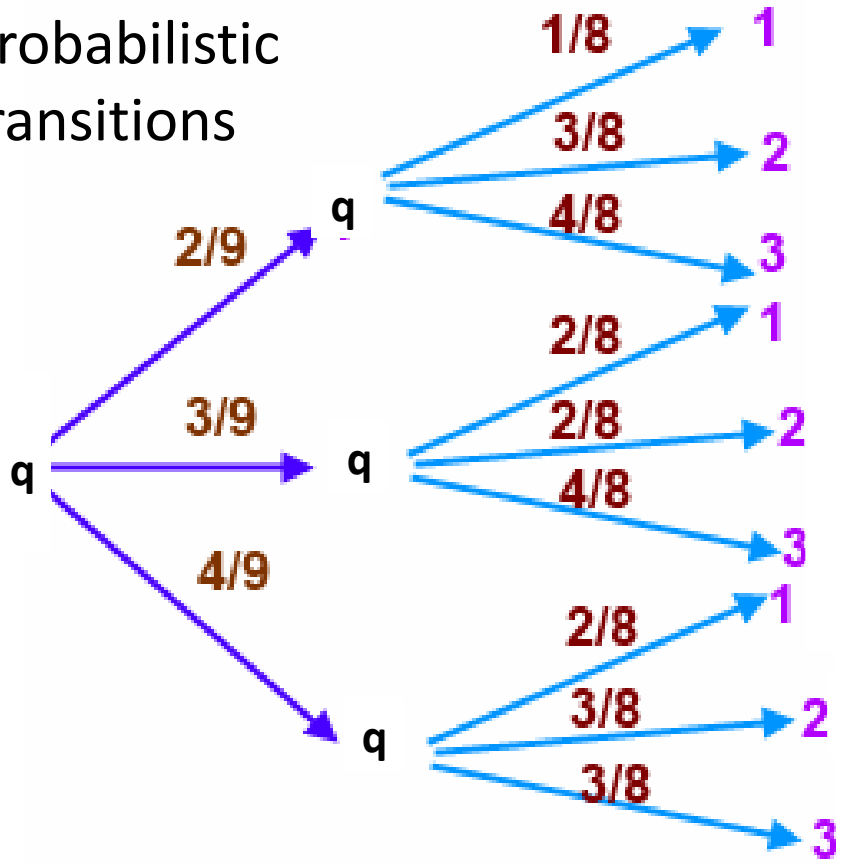

high branching factor

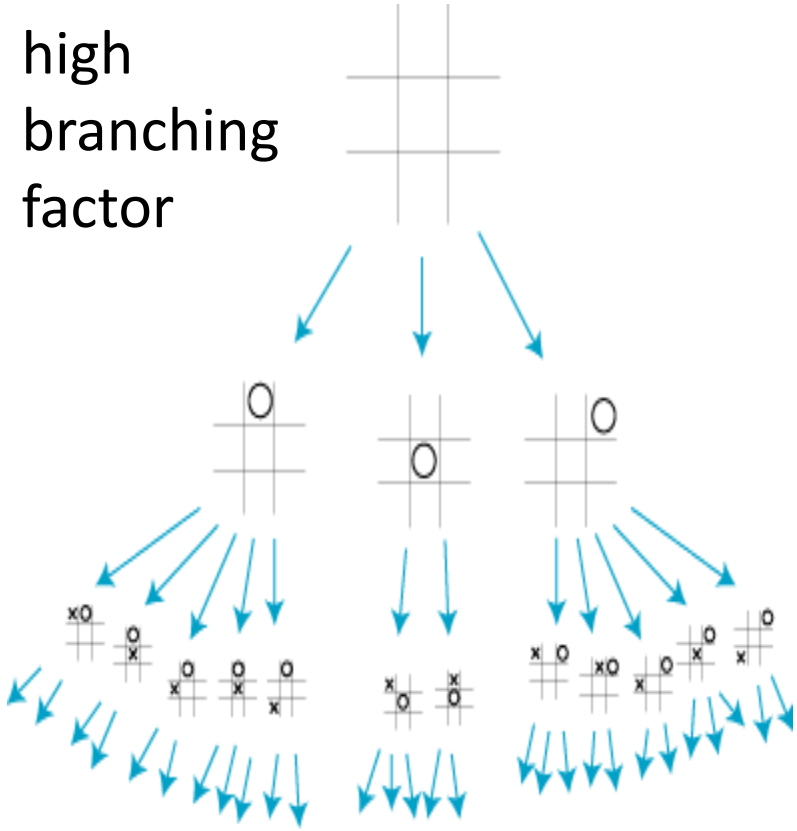

Our Goal

All future reward is zero since it is given immediately.

Learning simplifies to estimating the expected immediate reward, e.g. by the arithmetic mean.

Learning is simplified to estimating the expected immediate feedback.

Our Goal

- immediate feedback
- immediately adjust the return expectation

Reward Redistribution

A **reward redistribution** is a procedure that redistributes for each sequence the realization or the expectation of its return variable along the sequence.

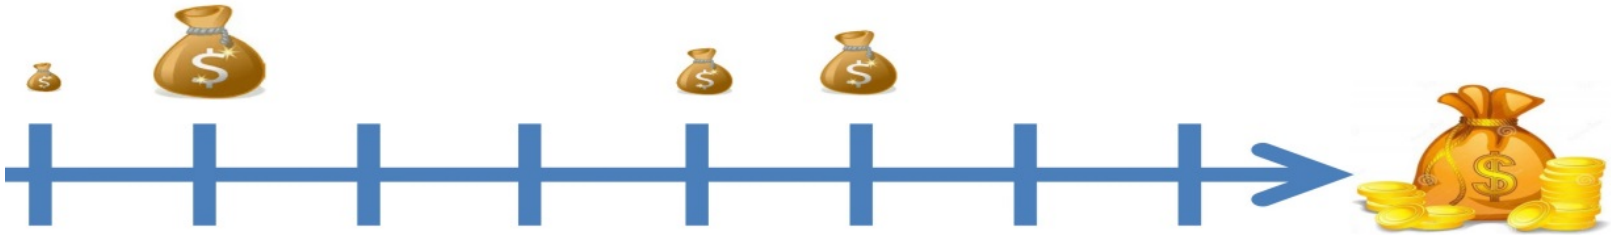

Reward redistributions do not change optimal policies:

Reward Redistribution

- Reward is the **change in the expected return**
  positive reward → increase of expected return
  negative reward → decrease of expected return

- Immediately adjust the return expectation

Reward Redistribution

Complex tasks are hierarchical with sub-tasks or sub-goals.
A step in the value function is a change in return expectation: amount or probability to obtain.

Example task below: both getting the key and opening the door increases the probability of obtaining the treasure.

Reward Redistribution

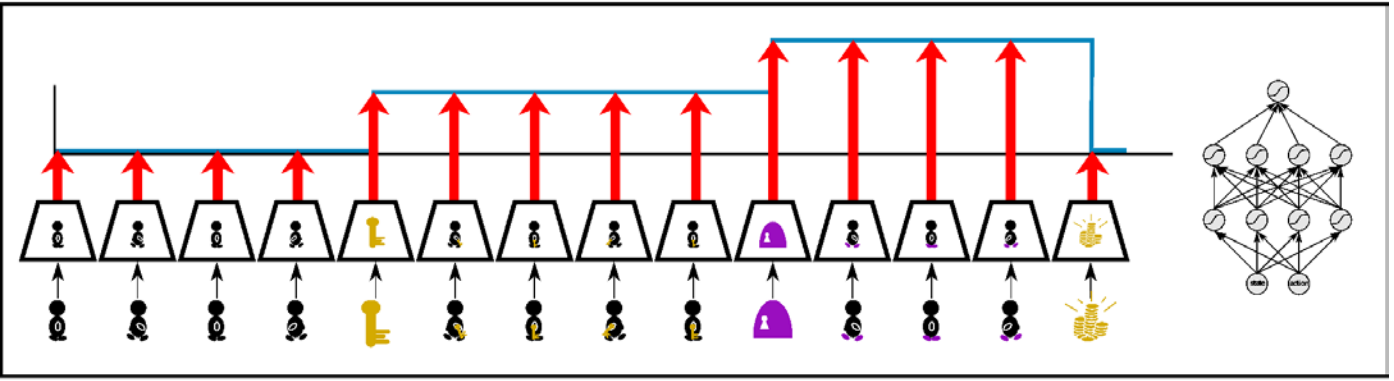

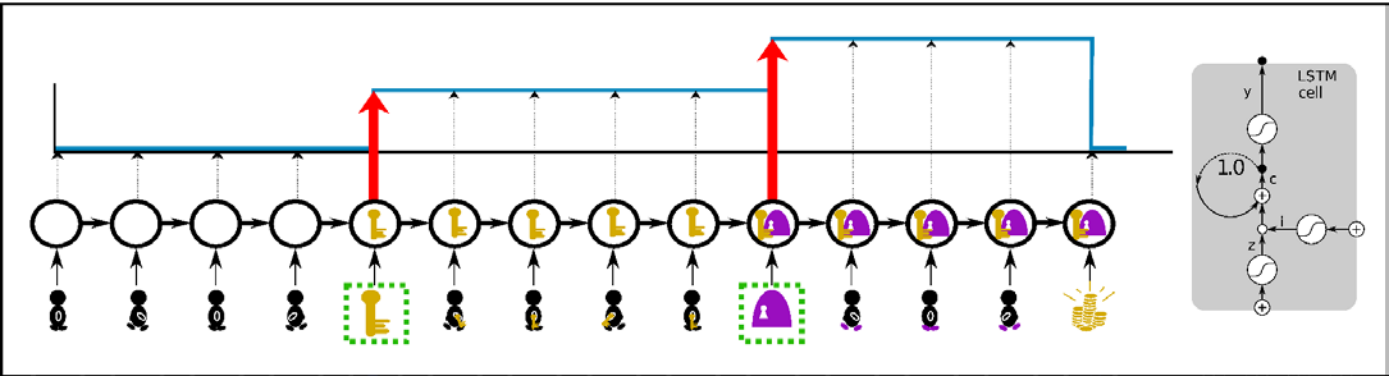

Learning step functions by fully connected networks requires to extract the expected return from every state-action.

Learning step functions by memorizing (→ LSTM) is much more sample efficient.

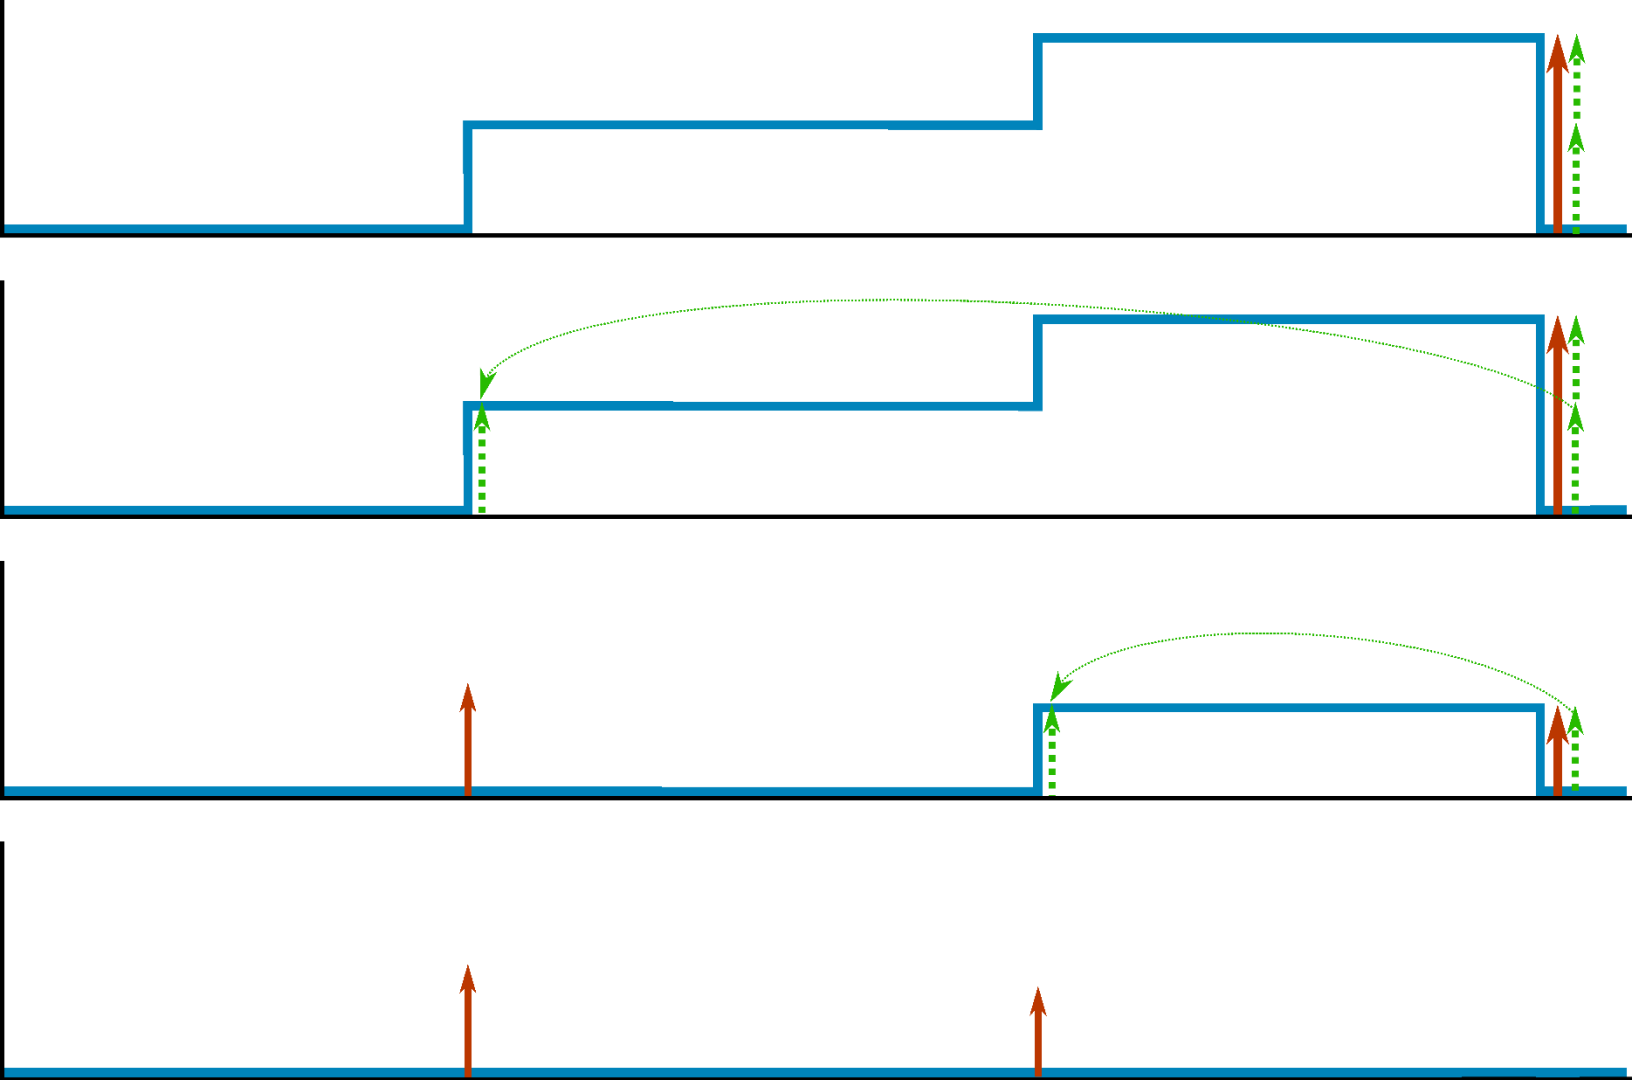

Reward Redistribution

Reward is redistributed to the steps which are identified by LSTM.

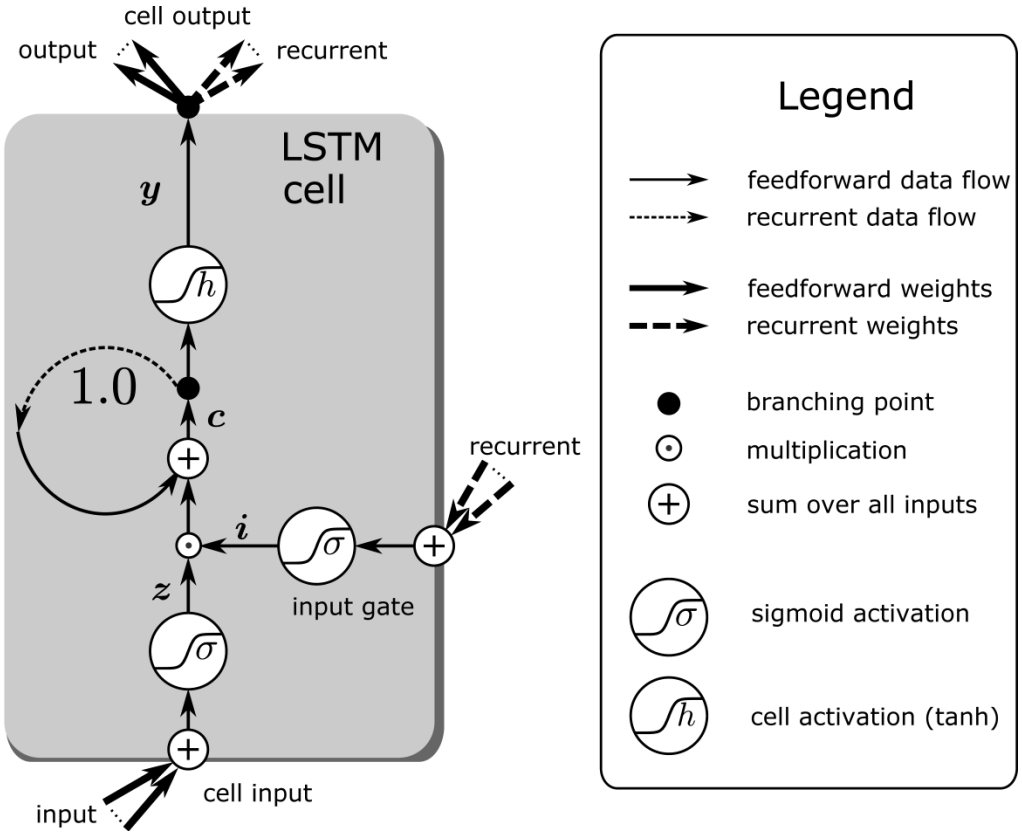

RUDDER for ATARI game BOWLING

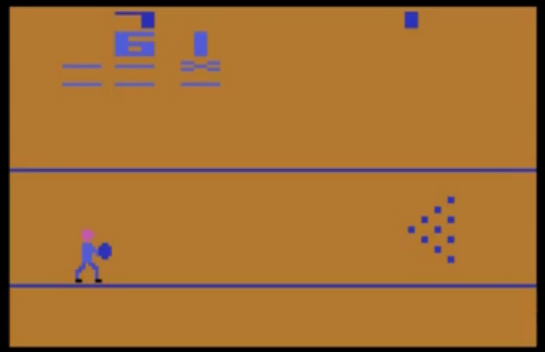

Agent has to
1.) Throw bowling ball
2.) Curve ball while it is rolling
3.) Clear all pins with ball

Agent has 2 rolls to knock down pins.
Rolling a strike (clearing all pins) will give 2 more rolls.
Reward is given at the end of all rolls.

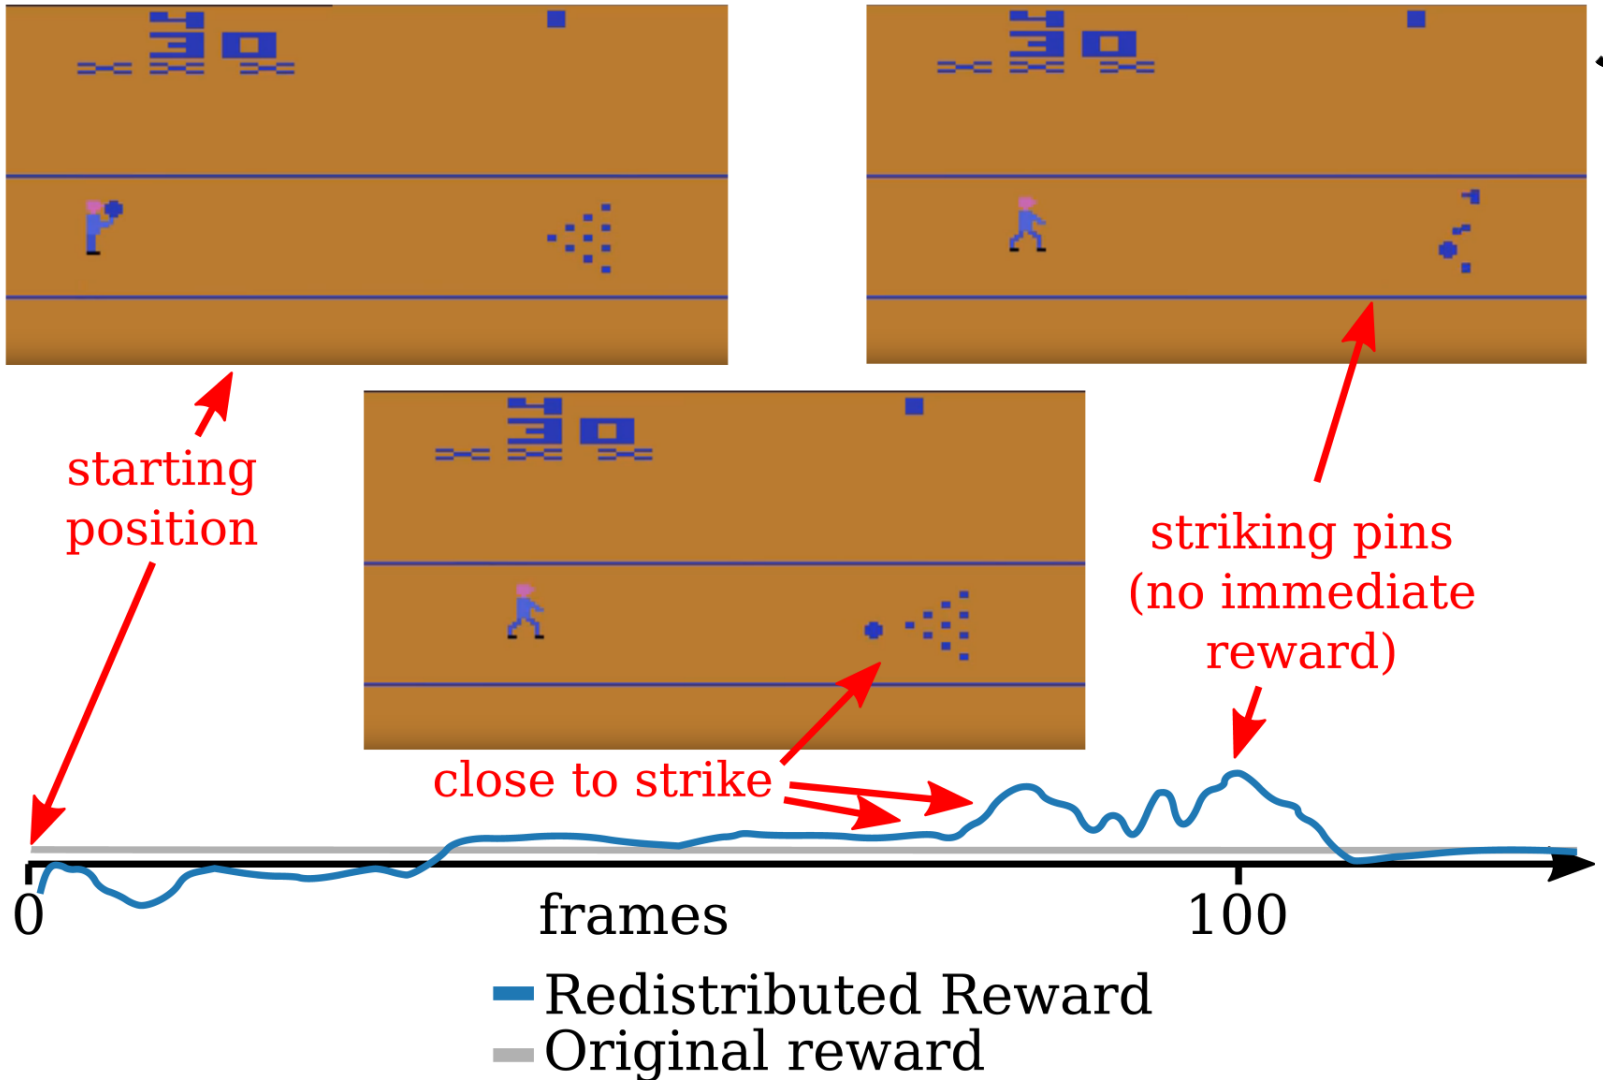
starting position
close to strike
striking pins (no immediate reward)
0
frames
100
Redistributed Reward
Original reward

RUDDER for ATARI game VENTURE

Agent has to
1.) Enter a room
2.) Collect the treasure 
3.) Shoot monsters 

Shooting monsters only gives reward after collecting treasure.
The hallmonsters  can not be killed.

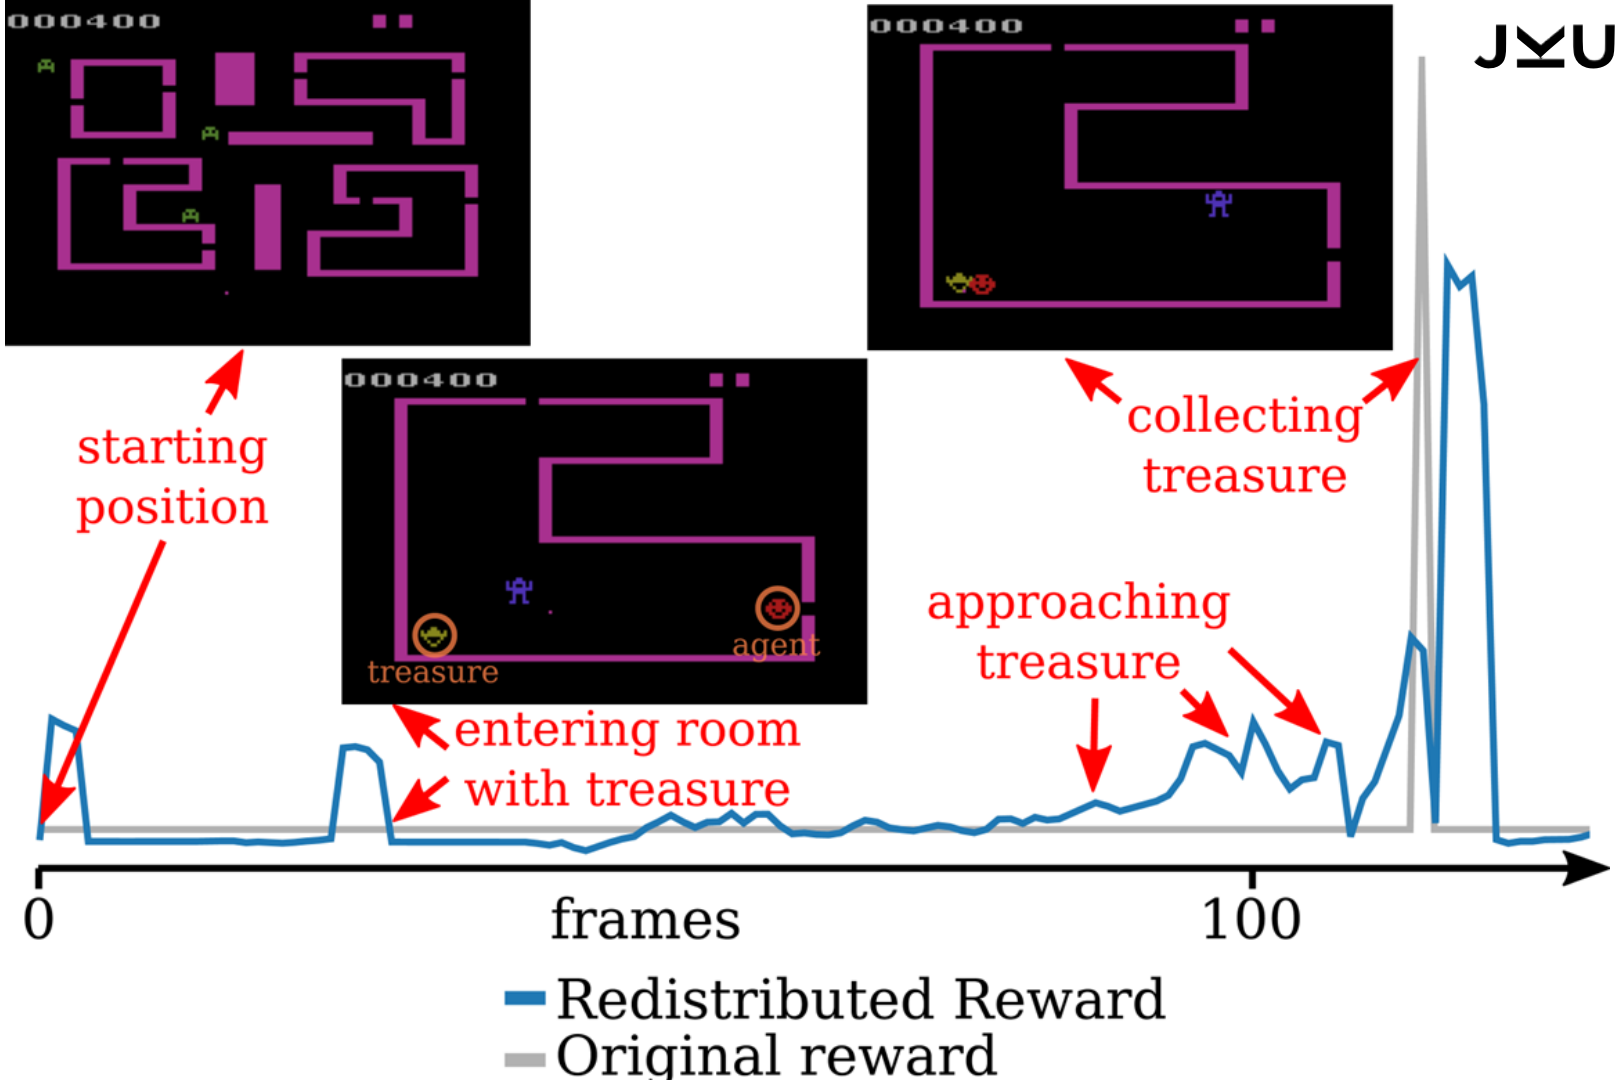
000400
000400
000400
starting position
collecting treasure
treasure
agent
entering room with treasure
approaching treasure
0
frames
100
Redistributed Reward
Original reward

Align-RUDDER

MineCraft

Align-RUDDER: Few Demonstrations

We have only **few demonstrations from humans**

- LSTM cannot learn on few demonstrations
- Overfitting

Solution

- Use **alignment** from bioinformatics to align demonstrations
- Underlying strategy is found: **profile model** and PSSM
- Works for as few as two demonstrations

Align-RUDDER: Alignment

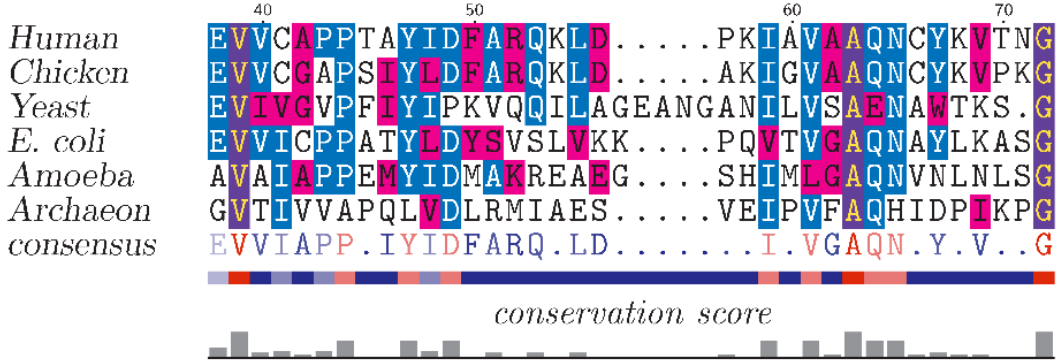

MineCraft (MineRL task "ObtainDiamond")

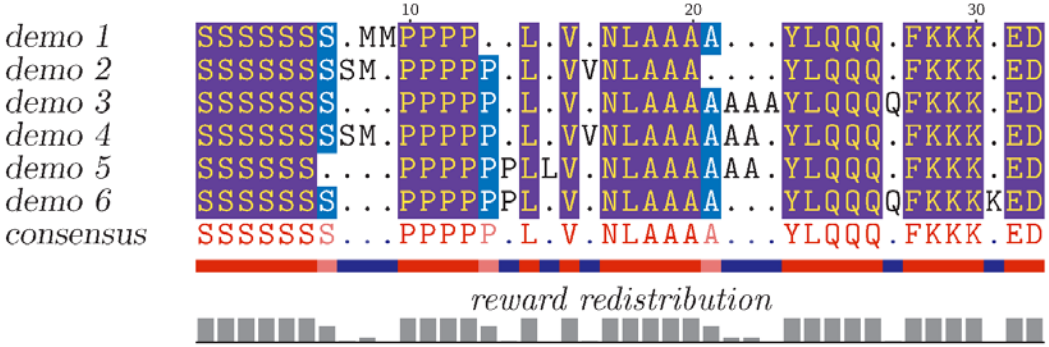

Align-RUDDER: Five Steps

I) Defining Events

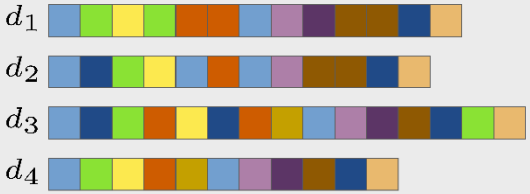

II) Scoring Matrix S

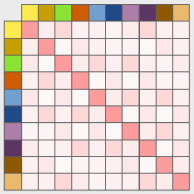

$$s_{ij} = \ln\left(\frac{q_{ij}}{p_i p_j}\right)/\lambda^*$$

III) Multiple Sequence Alignment

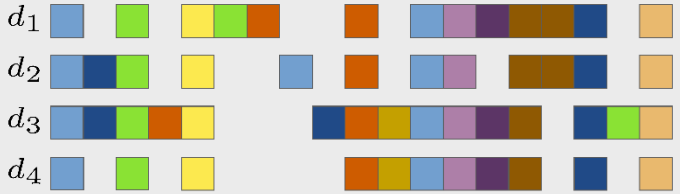

IV a) PSSM and Profile

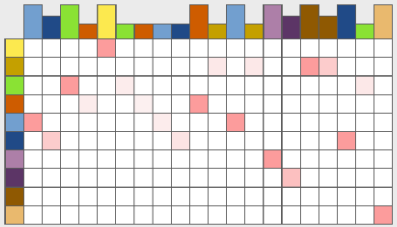

$$s_{i,t} = \ln\left(\frac{q_{i,t}}{p_i}\right)/\lambda_t^*$$

V) Reward Redistribution

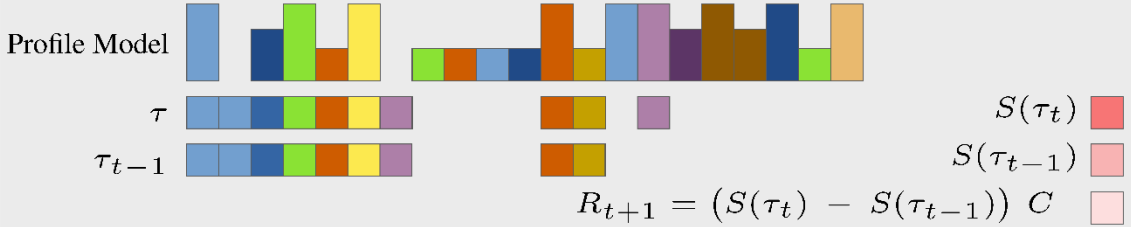

$$R_{t+1} = \left(S(\tau_t) - S(\tau_{t-1})\right) C$$

Align-RUDDER: MineCraft

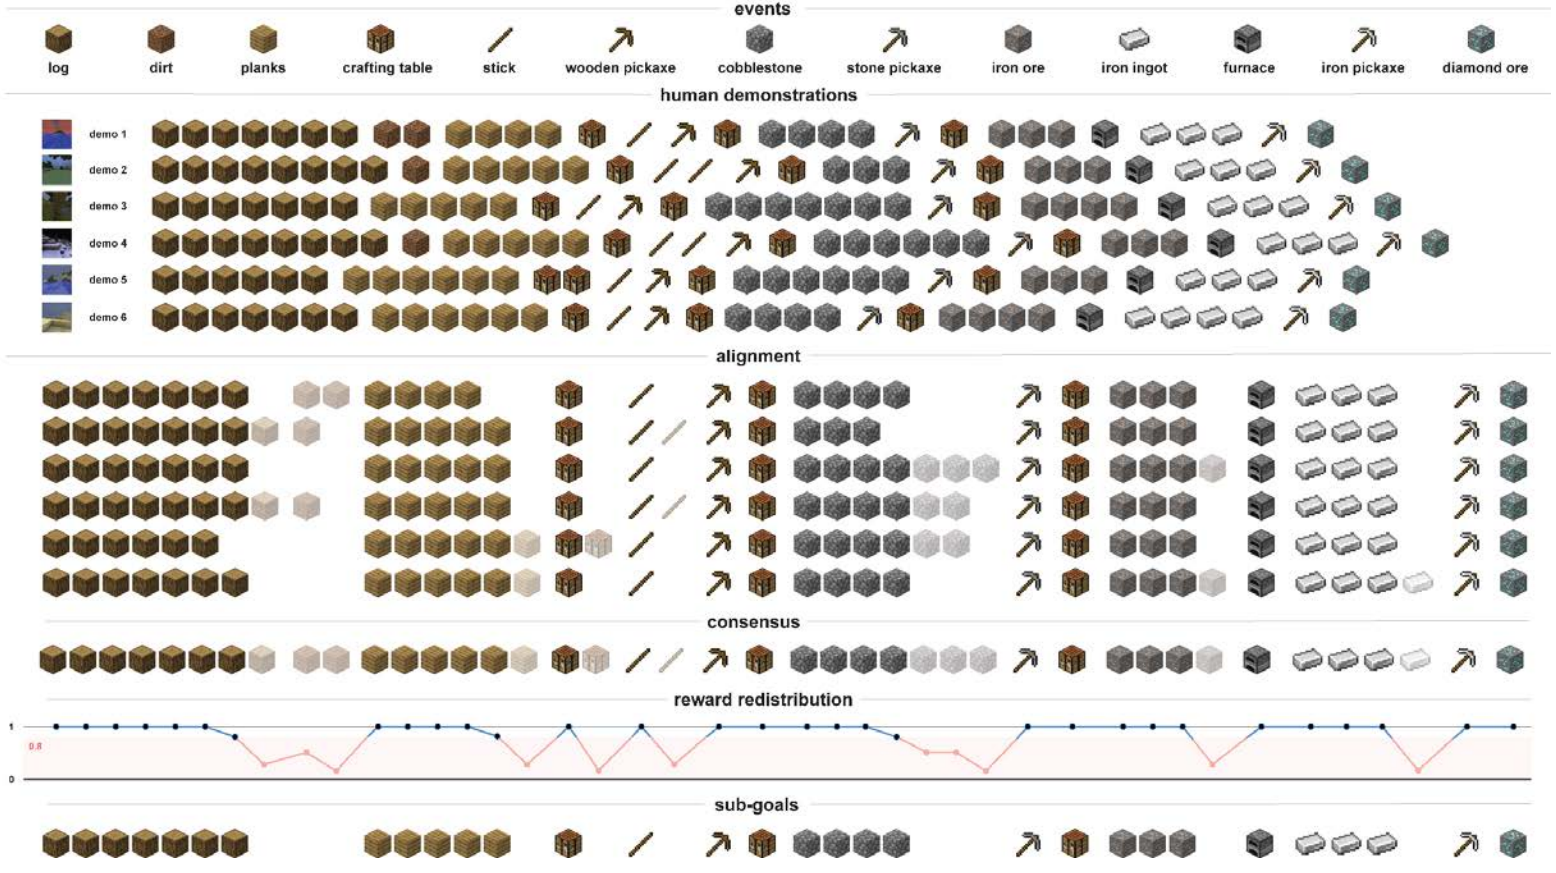

JKU

Align-RUDDER: Learning From Few Demonstrations by Reward Redistribution

A general Reinforcement Learning method showcased on the example of Minecraft

NEURAL INFORMATION PROCESSING SYSTEMS

2020


END

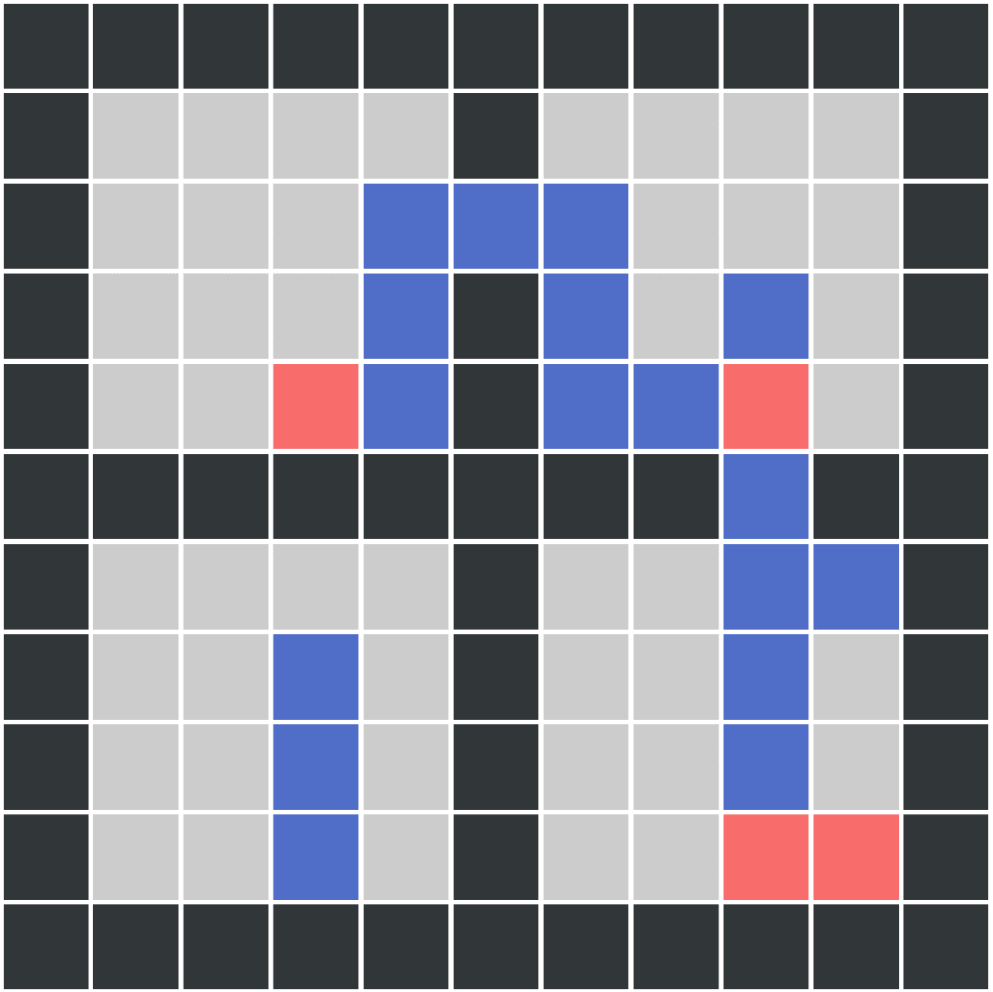

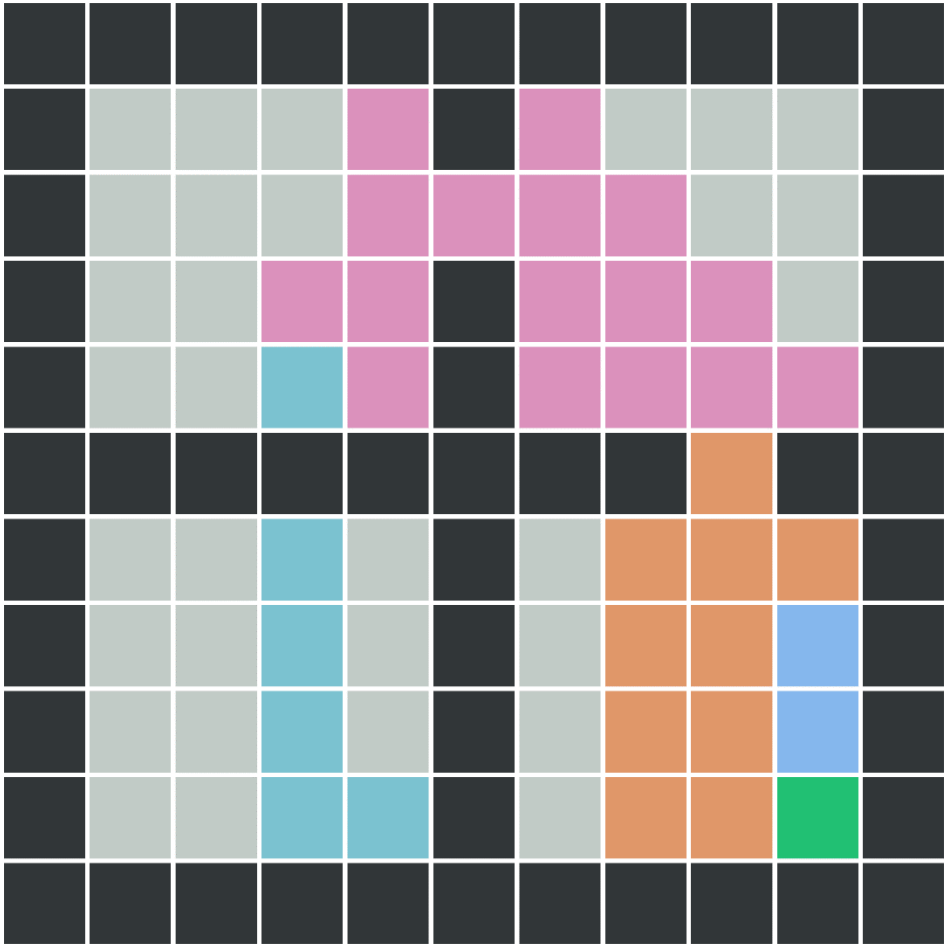

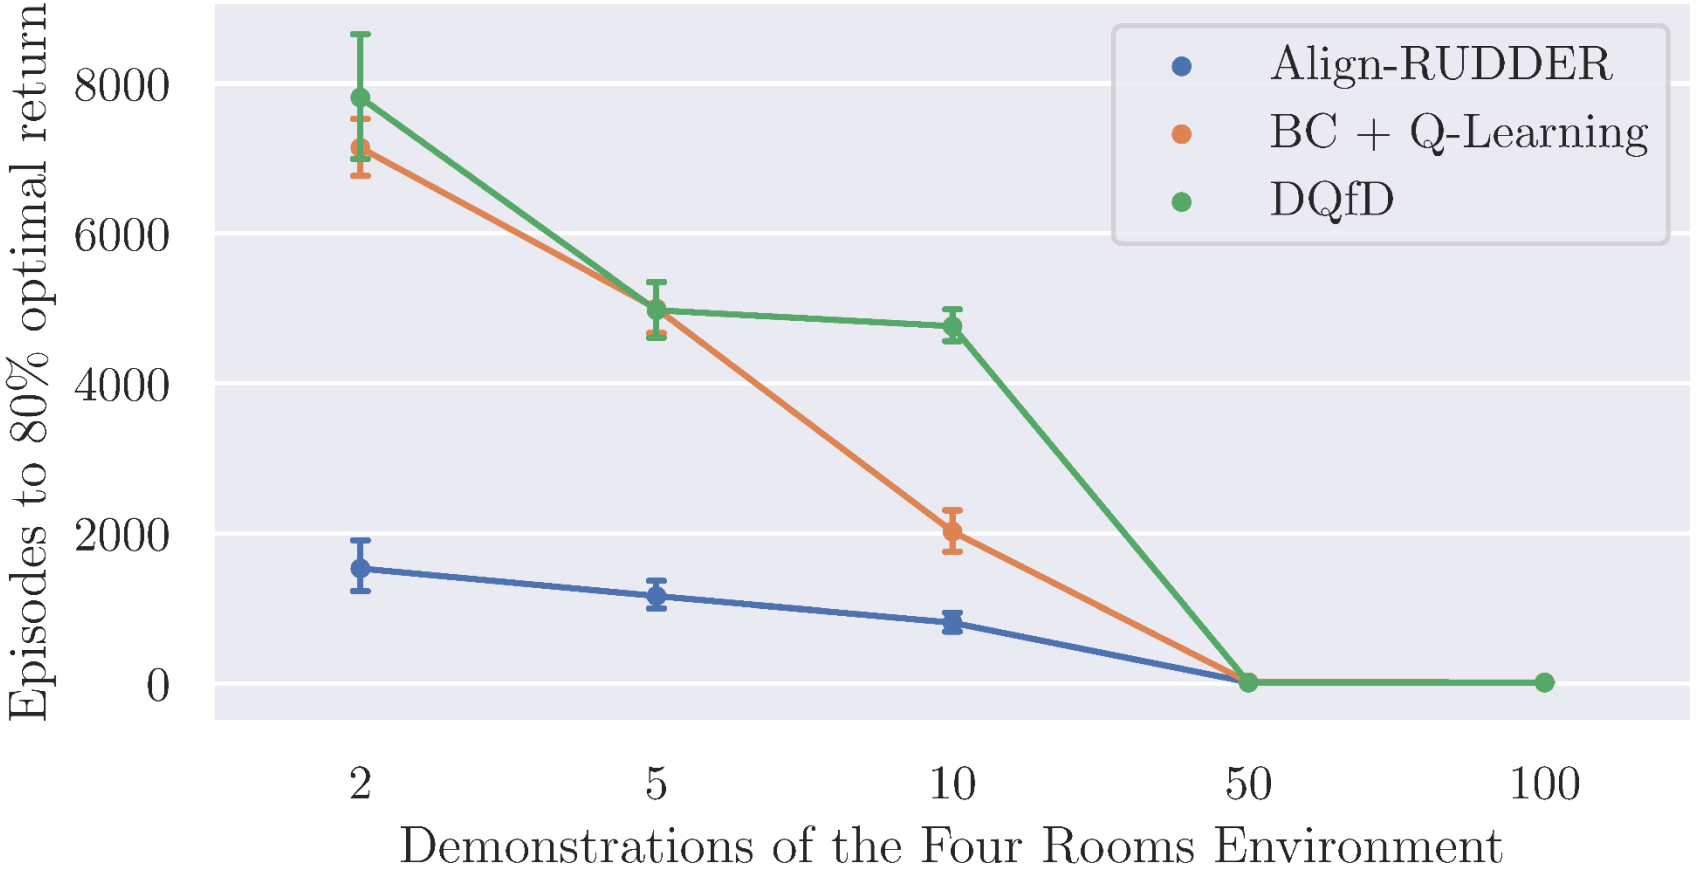
Episodes to 80% optimal return
8000
6000
4000
2000
0
Align-RUDDER
BC + Q-Learning
DQfD
2
5
10
50
100
Demonstrations of the Four Rooms Environment

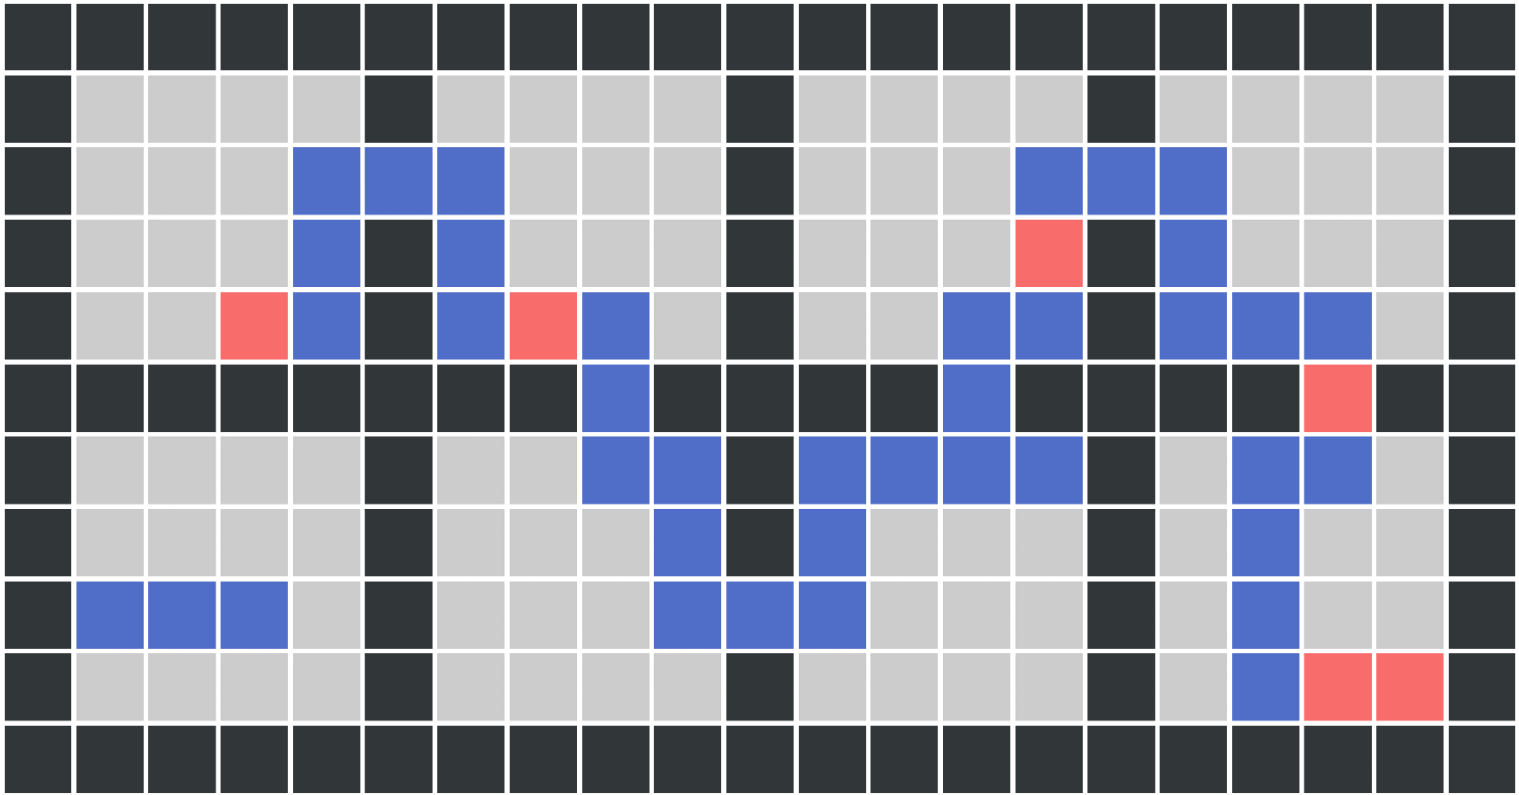

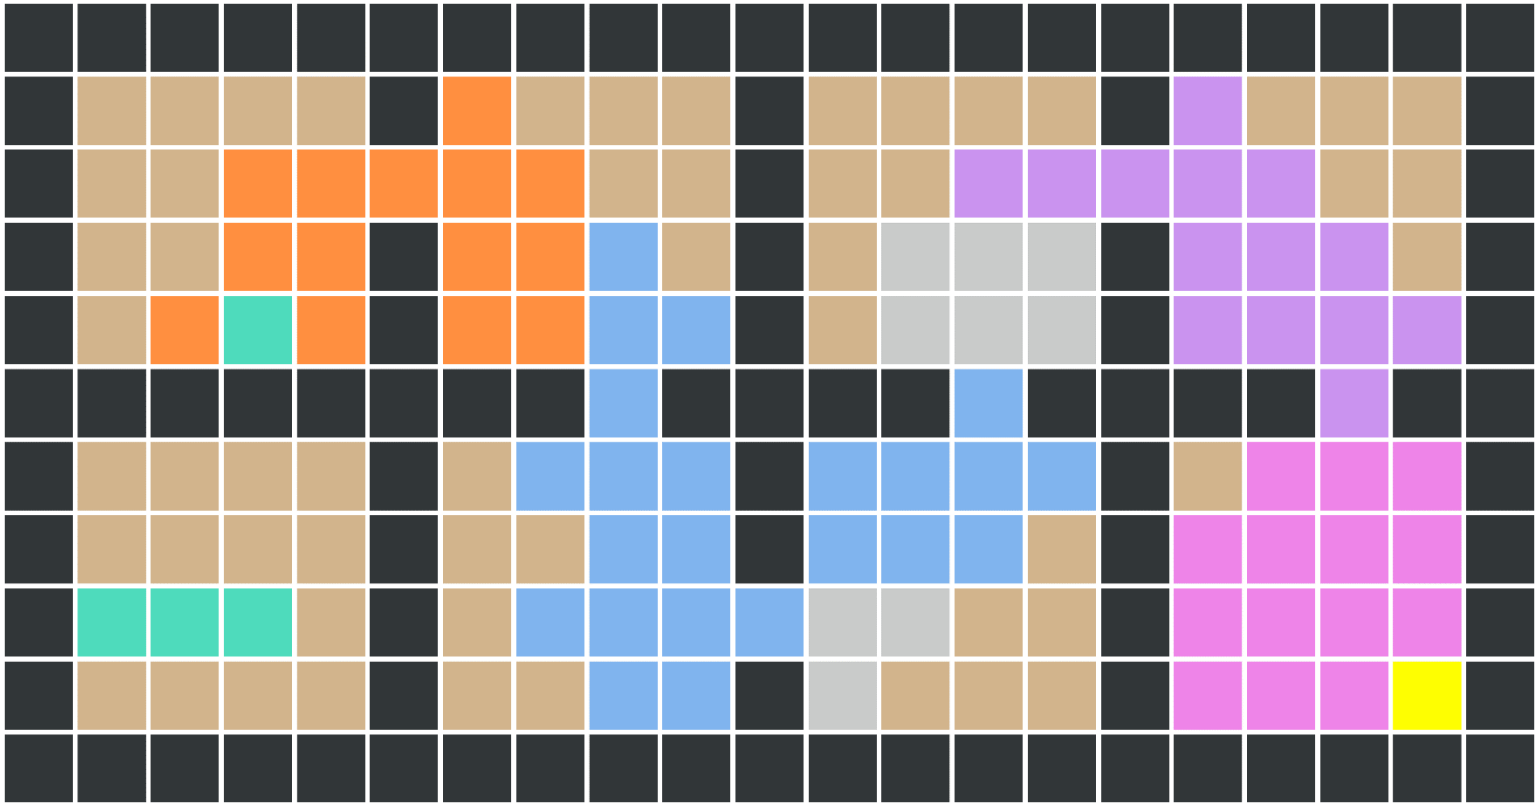

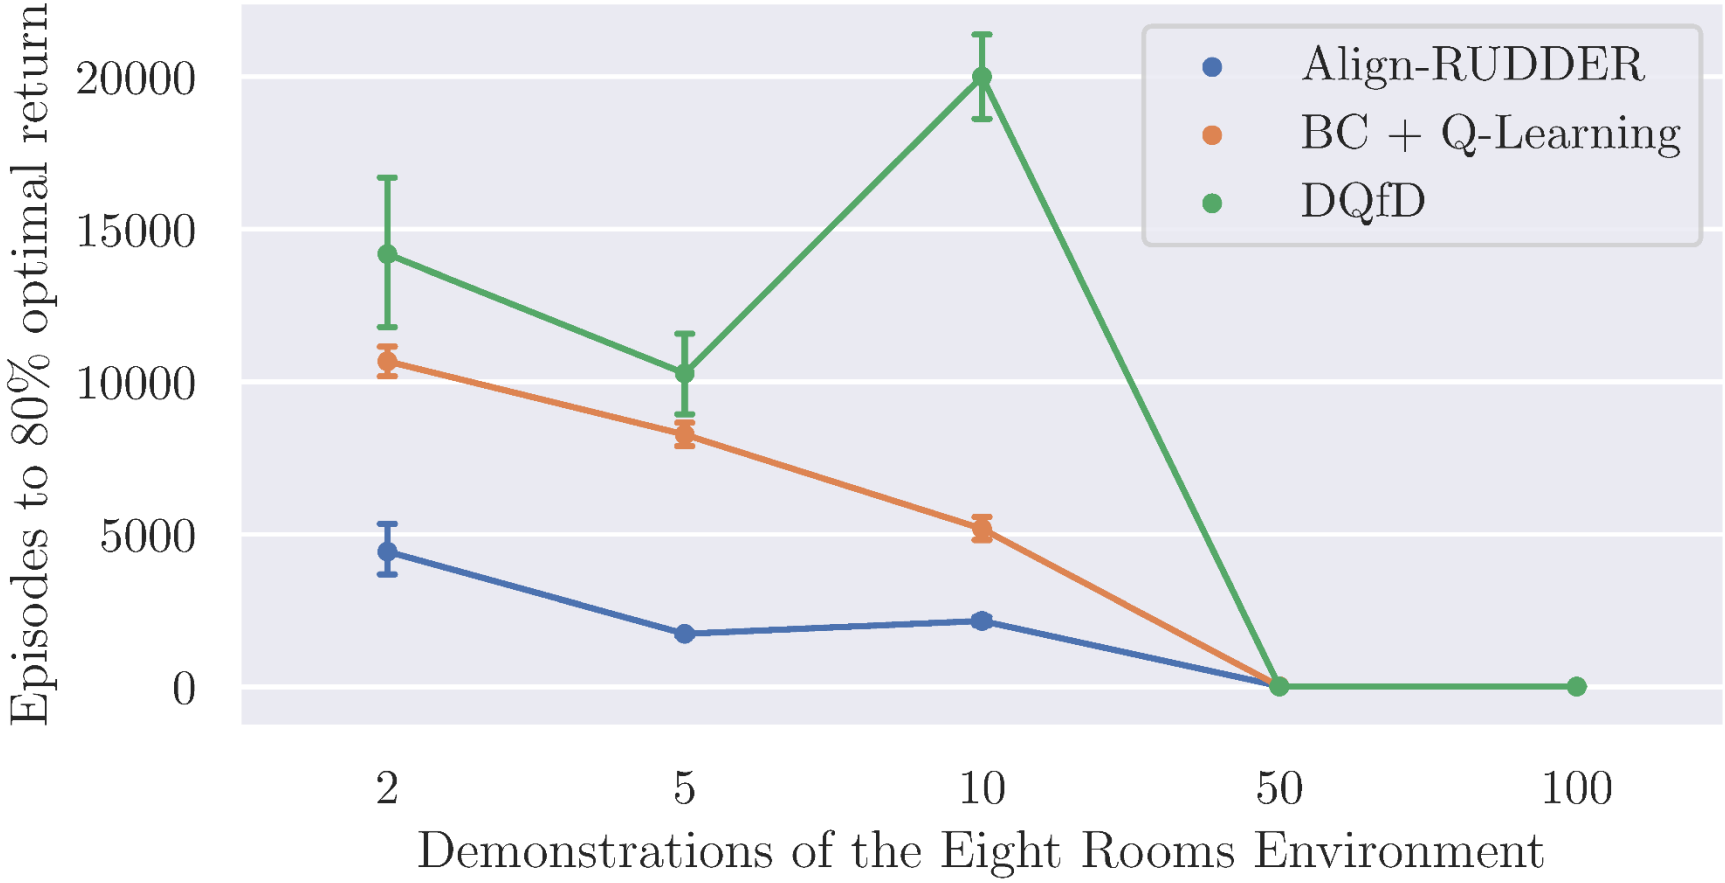
Episodes to 80% optimal return
20000
15000
10000
5000
0
2
5
10
50
100
Demonstrations of the Eight Rooms Environment
Align-RUDDER
BC + Q-Learning
DQfD