

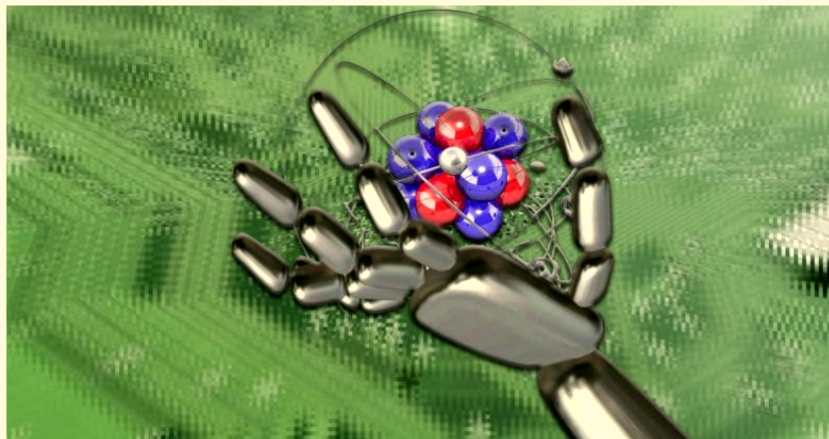
Title: Machine Learning Lecture

Speakers: Mohamed Hibat Allah

Collection: Machine Learning 2023/24

Date: April 02, 2024 - 11:30 AM

URL: <https://pirsa.org/24040046>



Machine Learning for many-body Physics

April 2 - May 2, 2023



Machine Learning for many-body Physics

Course Outline and your TA



Machine Learning for Many-Body Physics April 2 - May 2, 2024 Course Outline

Objective: This course is designed to introduce machine learning techniques for studying classical and quantum many-body problems encountered in quantum matter, quantum information, and related fields of physics. Lectures will emphasize relationships between statistical physics and machine learning. Tutorials and homework assignments will focus on developing programming skills for machine learning using Python.

Lecturer and Academic Staff: Mohamed Hibat-Allah

Email: mhibatallah@perimeterinstitute.ca

Office: 272

How to address me: Mohamed

How to get in touch with me: Just come with your question to my office. If I am not there, just send me an email and we can work out a schedule.

Teaching assistant: Megan Schuyler Moss

Email: msmoss@uwaterloo.ca

How to address me: Schuyler

How to get in touch with me: by email

Guest lecturer: Damian Pope



Schuyler Moss (TA)

PhD Student at UWaterloo and the Perimeter

Quantum Intelligence Lab (PIQUIL)

Schedule

- General course information** 0/3 ^

- Course outline**
- PDF
- References**
- TEXT
- Schedule**
- PDF


- Lecture notes** 0/1 v

- Tutorials** 0/1 v

- Homework assignments** 0/1 v

	Monday	Tuesday	Wednesday	Thursday	Friday
Week 1 (April 1 - 5)		Lecture 1 (Mohamed) 11:30 am - 12:30pm		Lecture 2 (Mohamed) 11:30 am - 12:30pm	
		Tutorial 1 - Python toolkits 2:00 - 3:30pm		Office hour: 12:30pm-1:30pm	
Week 2 (April 8 - 12)		Lecture 3 (Mohamed) 9:00 - 10:00am			
		Lecture 4 (Mohamed) 11:30 am - 12:30pm			
	Solar eclipse day	Tutorial 2 - Linear regression 2:00 - 3:30 pm			
Week 3 (April 15 - 19)	Homework 1 will be posted	Lecture 6 (Mohamed) 11:30 am - 12:30pm		Lecture 7 (Mohamed) 11:30 am - 12:30pm	
	Lecture 5 (Mohamed) 2:00 - 3:00pm	Tutorial 3 - Supervised Learning with Feedforward NNs 2:00 - 3:30pm		Office hour: 12:30pm-1:30pm	Lecture 8 (Mohamed) 3:45pm - 4:45pm
Week 4 (April 22 - 26)		Lecture 10 (Mohamed) 11:30 am - 12:30pm		Lecture 11 (Mohamed) 11:30 am - 12:30pm	Homework 1 and Presentation topic deadlines. Homework 2 will be posted (deadline two weeks after).
	Lecture 9 (Mohamed) 2:00 - 3:00pm	Tutorial 4 - Supervised Learning with Convolutional NNs 2:00 - 3:30pm		Office hour: 12:30pm-1:30pm	Tutorial 5 - Unsupervised learning 2:00 - 3:30pm

Lecture notes

<input type="radio"/>	General course information	0/3	∨
<input type="radio"/>	Lecture notes	0/1	∧
<input checked="" type="radio"/>	Lecture Notes (Last update April 1st)		
	 PDF		
<input type="radio"/>	Tutorials	0/1	∨
<input type="radio"/>	Homework assignments	0/1	∨

1 LECTURE 1

4

1 Lecture 1: Motivation, Goals, Definitions of Supervised, Unsupervised, and Reinforcement Learning

1.1 Background

Interest in machine learning has grown substantially in the last few years. In this course, we focus on many-body physics, but first, let us look at some exciting examples from other areas.

- Image classification (a neural network won a competition in 2012). Can play with this at <https://github.com/tensorflow/models>
- In 2016, AlphaGo won a Go competition against the world champion.
- Technology for self-driving cars uses reinforcement learning.
- Google Language Translation uses machine learning.

Quizzes

General course information 0/3 ∨


Lecture notes 0/1 ∨

Tutorials 0/1 ∨

Homework assignments 0/1 ∨

Quizzes 0/1 ^

Quiz 1: Machine Learning for Many-body physics

 QUIZ · 7 QUESTIONS

QUESTION 1 OF 7

Machine Learning (ML) is the equivalent as artificial intelligence (AI)

Choose only ONE best answer.


A Yes

B No

CONFIRM

Presentations on May 2nd

<input type="radio"/>	General course information	0/3	∨
<input type="radio"/>	Lecture notes	0/1	∨
<input type="radio"/>	Tutorials	0/1	∨
<input type="radio"/>	Homework assignments	0/1	∧
<input type="radio"/>	Student presentations on May 2nd		

 TEXT

- **5 minutes presentation** about a topic of interest at the **intersection of Machine learning and Physics** followed by questions.
- You can find **references** on the **PSI portal** or you can reach out to Mohamed.
- **Thursday, May 2nd at 11:30am – 1:00pm on the Time room.**
- **Deadline** to send a topic for the presentation to Mohamed is on **Friday, April 26th.**

Machine learning popularity

Interest over time [?](#)

Google Trends

● machine learning
Search term



2004

2012

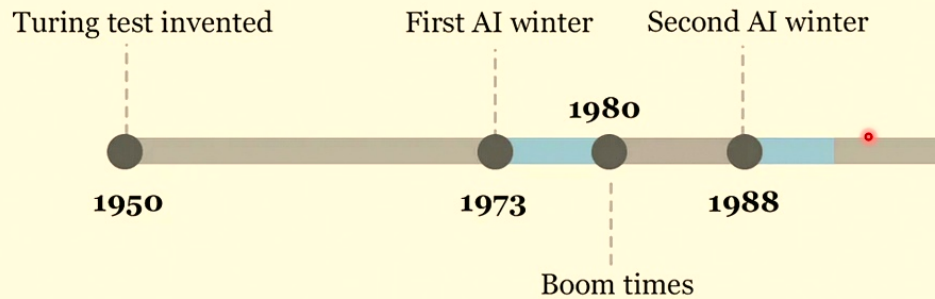
2018

2024

Machine Learning for many-body Physics

8

AI Winter



Credit: towardsdatascience.com

https://en.wikipedia.org/wiki/AI_winter

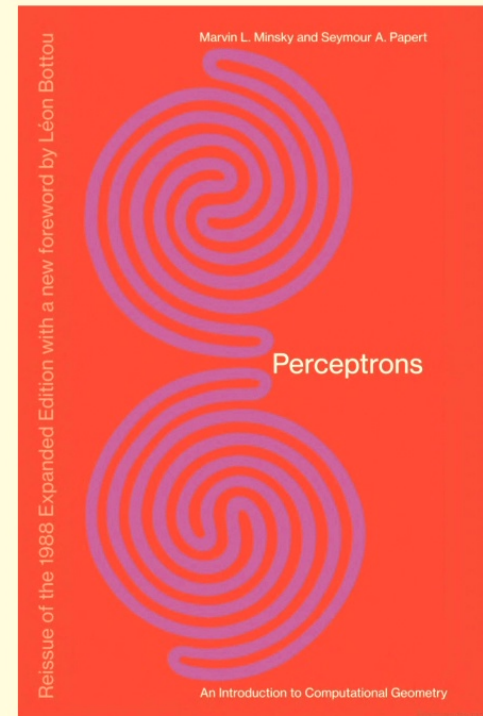


Image Classification (2012)

ImageNet Classification with Deep Convolutional Neural Networks

Alex Krizhevsky
 University of Toronto
 kriz@cs.utoronto.ca

Ilya Sutskever
 University of Toronto
 ilya@cs.utoronto.ca

Geoffrey E. Hinton
 University of Toronto
 hinton@cs.utoronto.ca

Abstract

We trained a large, deep convolutional neural network to classify the 1.2 million high-resolution images in the ImageNet LSVRC-2010 contest into the 1000 different classes. On the test data, we achieved top-1 and top-5 error rates of 37.5% and 17.0% which is considerably better than the previous state-of-the-art. The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax. To make training faster, we used non-saturating neurons and a very efficient GPU implementation of the convolution operation. To reduce overfitting in the fully-connected layers we employed a recently-developed regularization method called "dropout" that proved to be very effective. We also entered a variant of this model in the ILSVRC-2012 competition and achieved a winning top-5 test error rate of 15.3%, compared to 26.2% achieved by the second-best entry.



AlphaGo (2016)



Altmetric: 3193 Citations: 569 [More detail >>](#)

Article

Mastering the game of Go with deep neural networks and tree search

David Silver , Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel & Demis Hassabis 

Nature **529**, 484–489 (28 January 2016)
doi:10.1038/nature16961
[Download Citation](#)

Received: 11 November 2015
Accepted: 05 January 2016
Published: 27 January 2016

AlphaGo seals 4-1 victory over Go grandmaster Lee Sedol

DeepMind's artificial intelligence astonishes fans to defeat human opponent and offers evidence computer software has mastered a major challenge

Steven Borowiec

Tue 15 Mar 2016 10:12 GMT



<https://www.theguardian.com>

You can learn about that on the AlphaGo (documentary movie) 2017

Self-driving cars (2016)

End to End Learning for Self-Driving Cars

Mariusz Bojarski NVIDIA Corporation Holmdel, NJ 07735	Davide Del Testa NVIDIA Corporation Holmdel, NJ 07735	Daniel Dworakowski NVIDIA Corporation Holmdel, NJ 07735	Bernhard Firner NVIDIA Corporation Holmdel, NJ 07735
Beat Flepp NVIDIA Corporation Holmdel, NJ 07735	Prasoon Goyal NVIDIA Corporation Holmdel, NJ 07735	Lawrence D. Jackel NVIDIA Corporation Holmdel, NJ 07735	Mathew Monfort NVIDIA Corporation Holmdel, NJ 07735
Urs Muller NVIDIA Corporation Holmdel, NJ 07735	Jiakai Zhang NVIDIA Corporation Holmdel, NJ 07735	Xin Zhang NVIDIA Corporation Holmdel, NJ 07735	Jake Zhao NVIDIA Corporation Holmdel, NJ 07735
Karol Zieba NVIDIA Corporation Holmdel, NJ 07735			

Abstract

We trained a convolutional neural network (CNN) to map raw pixels from a single front-facing camera directly to steering commands. This end-to-end approach proved surprisingly powerful. With minimum training data from humans the system learns to drive in traffic on local roads with or without lane markings and on highways. It also operates in areas with unclear visual guidance such as in parking lots and on unpaved roads.

The system automatically learns internal representations of the necessary processing steps such as detecting useful road features with only the human steering angle as the training signal. We never explicitly trained it to detect, for example, the outline of roads.

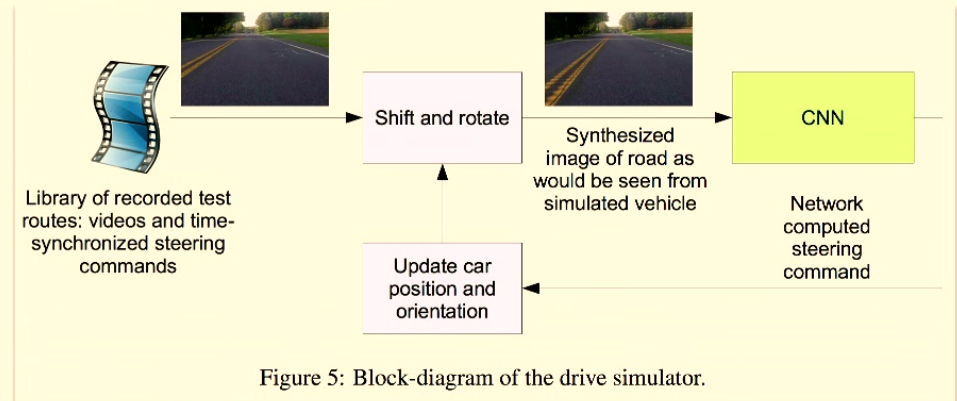


Figure 5: Block-diagram of the drive simulator.

r:1604.07316v1 [cs.CV] 25 Apr 2016

Medical Diagnosis (2017)

CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning


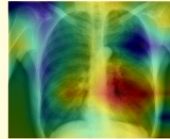
Pranav Rajpurkar^{*1} Jeremy Irvin^{*1} Kaylie Zhu¹ Brandon Yang¹ Hershel Mehta¹
Tony Duan¹ Daisy Ding¹ Aarti Bagul¹ Robyn L. Ball² Curtis Langlotz³ Katie Shpanskaya³
Matthew P. Lungren³ Andrew Y. Ng¹

Abstract

We develop an algorithm that can detect pneumonia from chest X-rays at a level exceeding practicing radiologists. Our algorithm, CheXNet, is a 121-layer convolutional neural network trained on ChestX-ray14, currently the largest publicly available chest X-ray dataset, containing over 100,000 frontal-view X-ray images with 14 diseases. Four practicing academic radiologists annotate a test set, on which we compare the performance of CheXNet to that of radiologists. We find that CheXNet exceeds average radiologist performance on the F1 metric. We extend CheXNet to detect all 14 diseases in ChestX-ray14 and achieve state of the art results on all 14 diseases.

1. Introduction

More than 1 million adults are hospitalized with pneumonia and around 50,000 die from the disease every year. Most (80%) die from bacterial pneumonia.


Input Chest X-Ray Image
CheXNet 121-layer CNN
Output Pneumonia Positive (85%)


:1711.05225v3 [cs.CV] 25 Dec 2017

Language translation (2016)

Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi
yonghui, schuster, zhifengc, qvl, mnorouzi@google.com

Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, Jeffrey Dean

Abstract

Neural Machine Translation (NMT) is an end-to-end learning approach for automated translation, with the potential to overcome many of the weaknesses of conventional phrase-based translation systems. Unfortunately, NMT systems are known to be computationally expensive both in training and in translation inference – sometimes prohibitively so in the case of very large data sets and large models. Several authors have also charged that NMT systems lack robustness, particularly when input sentences contain rare words. These issues have hindered NMT's use in practical deployment and services, where both accuracy and speed are essential. In this work, we present GNMT, Google's Neural Machine Translation system, which attempts to address many of these issues. Our model consists of a deep LSTM network with 8 encoder and 8 decoder layers using residual connections as well as attention connections from the decoder network to the encoder. To improve parallelism and therefore decrease training time, our attention mechanism connects the bottom layer of the decoder to the top layer of the encoder. To accelerate the final translation speed, we employ low-precision arithmetic during inference computations. To improve handling of rare words, we divide words into a limited set of common sub-word units ("wordpieces") for both input and output. This method provides a good balance between the flexibility of "character"-delimited models and the efficiency of "word"-delimited models, naturally handles translation of rare words, and ultimately improves the overall accuracy of the system. Our beam search technique employs a length-normalization procedure and uses a coverage penalty, which encourages generation of an output sentence that is most likely to cover all the words in the source sentence. To directly optimize the translation BLEU scores, we consider refining the models by using reinforcement learning, but we found that the improvement in the BLEU scores did not reflect in the human evaluation. On the WMT'14 English-to-French and English-to-German benchmarks, GNMT achieves competitive results to state-of-the-art. Using a human side-by-side evaluation on a set of isolated simple sentences, it reduces translation errors by an average of 60% compared to Google's phrase-based production system.

arXiv:1609.08144v2 [cs.CL] 8 Oct 2016

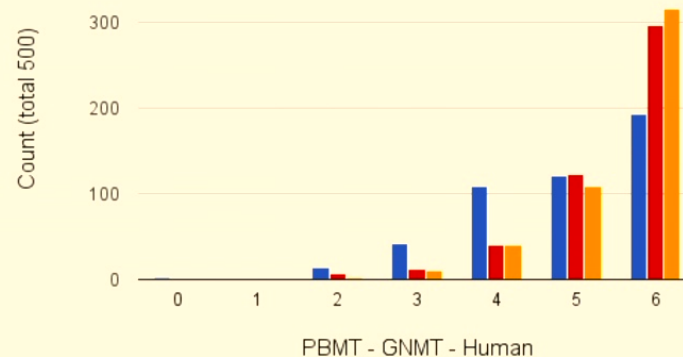


Figure 6: Histogram of side-by-side scores on 500 sampled sentences from Wikipedia and news websites for a typical language pair, here English → Spanish (PBMT blue, GNMT red, Human orange). It can be seen that there is a wide distribution in scores, even for the human translation when rated by other humans, which shows how ambiguous the task is. It is clear that GNMT is much more accurate than PBMT.

Chat-GPT (Launched in 2022)



You

Explain machine learning in one paragraph to a five years old kid



ChatGPT

Machine learning is like teaching a computer how to do something by showing it lots of examples. Just like when you learn to ride a bicycle by practicing, the computer learns from trying different things and getting better over time. So, if you want the computer to recognize cats, you show it many pictures of cats and tell it, "This is a cat!" Then, after seeing many pictures, it learns to recognize cats by itself. It's like having a really smart friend who learns new things by looking at lots of pictures and figuring out what they are.

GPT-4

GPT-4 Technical Report

OpenAI*

Abstract

We report the development of GPT-4, a large-scale, multimodal model which can accept image and text inputs and produce text outputs. While less capable than humans in many real-world scenarios, GPT-4 exhibits human-level performance on various professional and academic benchmarks, including passing a simulated bar exam with a score around the top 10% of test takers. GPT-4 is a Transformer-based model pre-trained to predict the next token in a document. The post-training alignment process results in improved performance on measures of factuality and adherence to desired behavior. A core component of this project was developing infrastructure and optimization methods that behave predictably across a wide range of scales. This allowed us to accurately predict some aspects of GPT-4's performance based on models trained with no more than 1/1,000th the compute of GPT-4.



GPT-3.5 – Handwritten Notes to Latex

Outline:

- ① Recurrent Neural Networks (RNNs):
- ② $\mu_x = \text{RNN}$ for variational Monte Carlo:

Q1 Recurrent Neural Networks (RNNs)

- * Originally built for language processing
- ↳ Ex: Machine translation, speech recognition, Music generation, DNA sequence analysis.

↳ Example:

“I live in France, hence I speak ...”

↳ Most likely the next word is French, so how can we process it in a neural network, so that it predicts the word “French”

↳ There comes Recurrent neural networks to the rescue.

Prompt:

Can you convert the following handwritten notes to latex?

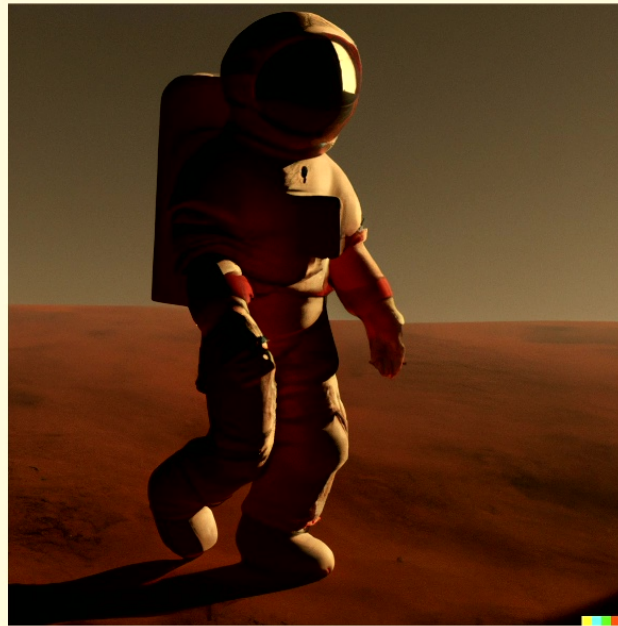
Outline

1. Recurrent Neural Networks (RNNs):
2. $\mu_x = \text{RNN}$ for variational Monte Carlo.

Recurrent Neural Networks (RNNs)

- Originally built for language processing.
 - Ex: Machine translation, speech recognition, music generation, DNA sequence analysis.
- Example:
 - “I live in France, hence I speak...”
 - “Next likely the next word is French, so how can we process it in a neural network so that it predicts the word ‘French’?”
 - Enter convo Recurrent neural networks to the rescue.

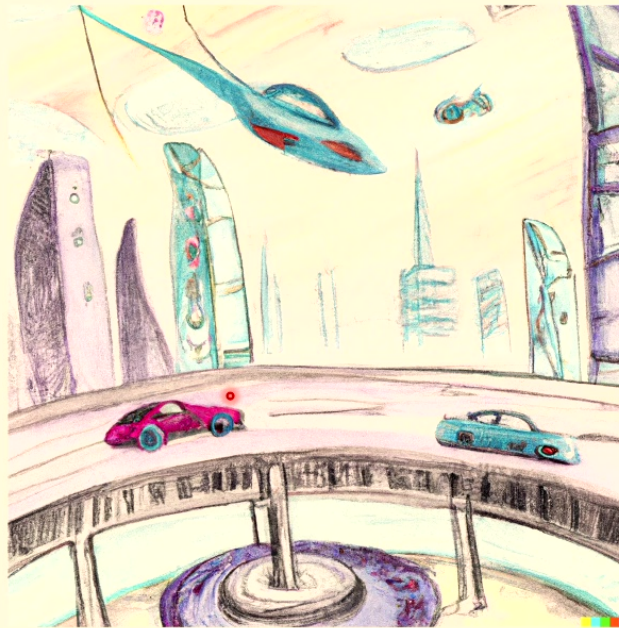
Dall-E



<https://labs.openai.com>

A 3d realistic render of an astronaut walking on Mars

Dall-E



A pencil and watercolor drawing of a bright city in the future with flying cars

Dall-E



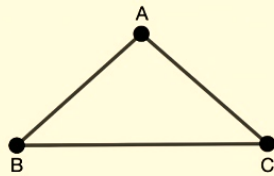
A handpalm with a tree growing on top of it

Artificial Intelligence Virtual Artist



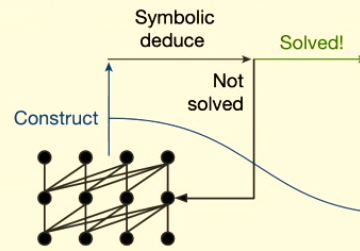
AlphaGeometry: Solving Olympiad Geometry problems

a A simple problem



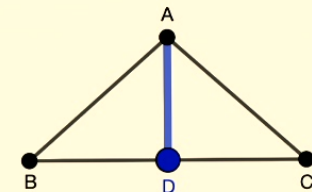
"Let ABC be any triangle with $AB = AC$. Prove that $\angle ABC = \angle BCA$."

b AlphaGeometry



c Language model

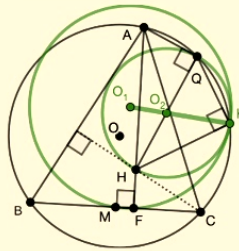
d Solution



Construct D: midpoint BC,
 $AB=AC, BD = DC, AD=AD \Rightarrow \angle ABD=\angle DCA$ [1]
 [1], $B C D$ collinear $\Rightarrow \angle ABC=\angle BCA$

e IMO 2015 P3

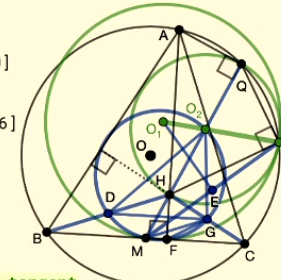
"Let ABC be an acute triangle. Let (O) be its circumcircle, H its orthocenter, and F the foot of the altitude from A. Let M be the midpoint of BC. Let Q be the point on (O) such that $QH \perp QA$ and let K be the point on (O) such that $KH \perp KQ$. Prove that the circumcircles (O_1) and (O_2) of triangles FKM and KQH are tangent to each other."



Alpha-Geometry

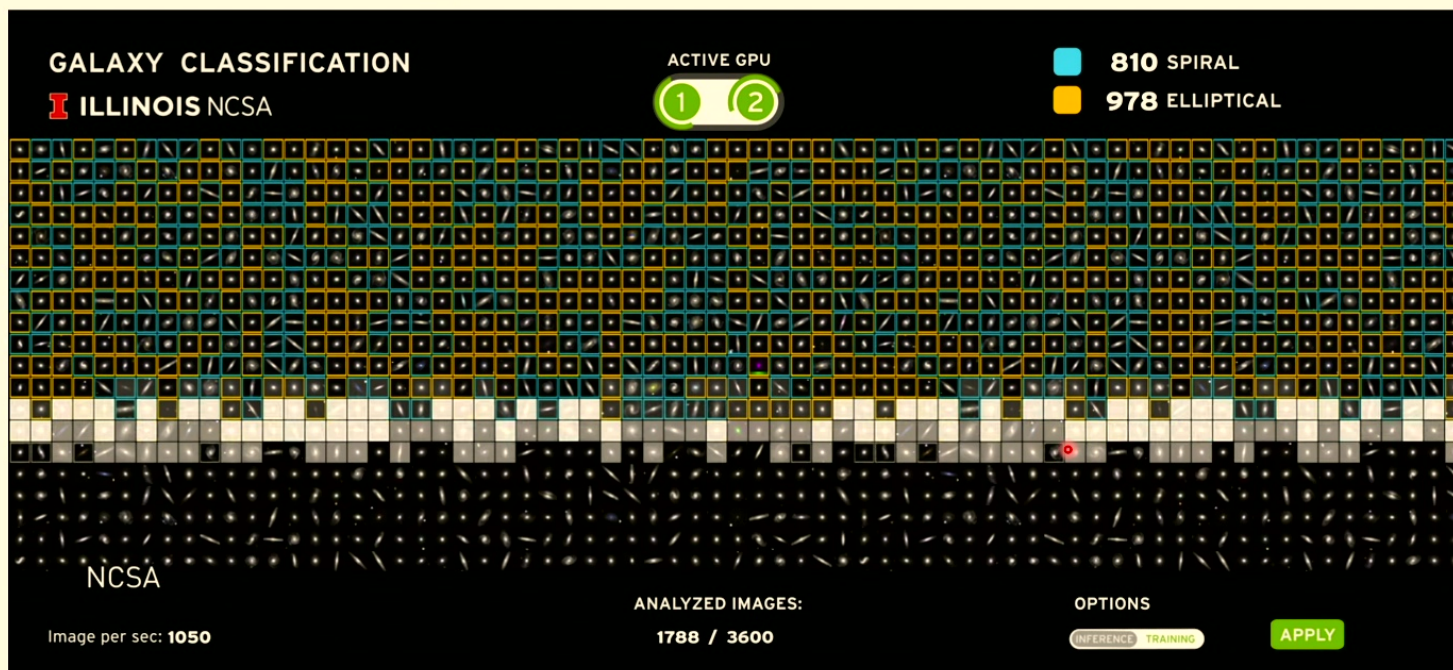
f Solution

Construct D: midpoint BH [a]
 [a], O_2 midpoint HQ $\Rightarrow BQ \parallel O_2D$ [20]
 Construct G: midpoint HC [b] ...
 $\angle GMD = \angle GO_2D \Rightarrow M O_2 G D$ cyclic [26]
 [a], [b] $\Rightarrow BC \parallel DG$ [30]
 Construct E: midpoint MK [c]
 ..., [c] $\Rightarrow \angle KFC = \angle KO_1E$ [104]
 $\angle FKO_1 = \angle FKO_2 \Rightarrow KO_1 \parallel KO_2$ [109]
 [109] $\Rightarrow O, O_2, K$ collinear $\Rightarrow (O_1), (O_2)$ tangent



Solving olympiad geometry without human demonstrations, Nature, 2023

Galaxy Classification



Credit: ILLINOIS NCSA, NVIDIA

Gravitational Waves Detection

Convolutional neural networks: A magic bullet for gravitational-wave detection?

Timothy D. Gebhard,^{1,2,*} Niki Kilbertus^{1,3,†}, Ian Harry,^{4,5} and Bernhard Schölkopf¹

¹Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany

²Max Planck ETH Center for Learning Systems, Universitätsstrasse 6, 8092 Zürich, Switzerland

³Engineering Department, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom

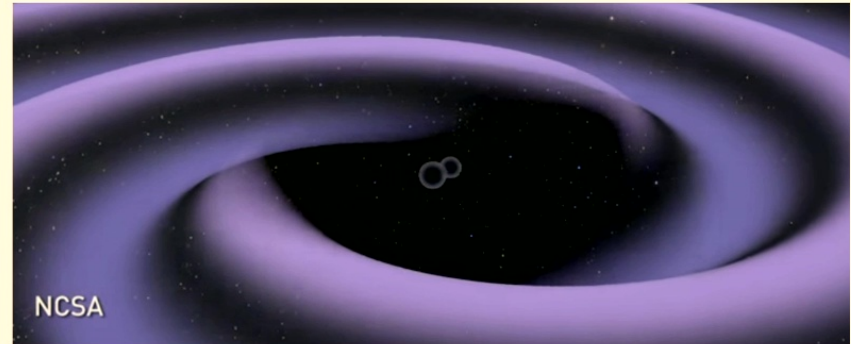
⁴Institute for Cosmology and Gravitation, University of Portsmouth, 1-8 Burnaby Road, Portsmouth, PO1 3FZ, United Kingdom

⁵Max Planck Institute for Gravitational Physics, Am Mühlenberg 1, 14476 Potsdam, Germany

Ⓞ (Received 25 April 2019; published 26 September 2019)

In the last few years, machine learning techniques, in particular convolutional neural networks, have been investigated as a method to replace or complement traditional matched filtering techniques that are used to detect the gravitational-wave signature of merging black holes. However, to date, these methods have not yet been successfully applied to the analysis of long stretches of data recorded by the Advanced LIGO and Virgo gravitational-wave observatories. In this work, we critically examine the use of convolutional neural networks as a tool to search for merging black holes. We identify the strengths and limitations of this approach, highlight some common pitfalls in translating between machine learning and gravitational-wave astronomy, and discuss the interdisciplinary challenges. In particular, we explain in detail why convolutional neural networks alone cannot be used to claim a statistically significant gravitational-wave detection. However, we demonstrate how they can still be used to rapidly flag the times of potential signals in the data for a more detailed follow-up. Our convolutional neural network architecture as well as the proposed performance metrics are better suited for this task than a standard binary classifications scheme. A detailed evaluation of our approach on Advanced LIGO data demonstrates the potential of such systems as trigger generators. Finally, we sound a note of caution by constructing *adversarial examples*, which showcase interesting “failure modes” of our model, where inputs with no visible resemblance to real gravitational-wave signals are identified as such by the network with high confidence.

DOI: 10.1103/PhysRevD.100.063015



Credit: ILLINOIS NCSA, NVIDIA

Discovery of Physical Laws

ScienceAdvances

AI Feynman: A physics-inspired method for symbolic regression

SILVIU-MARIAN UDRESCU AND MAX TEGMARK [Authors Info & Affiliations](#)

SCIENCE ADVANCES · 15 Apr 2020 · Vol 6, Issue 16 · DOI: 10.1126/sciadv.abb2631

48,019 196



Abstract

A core challenge for both physics and artificial intelligence (AI) is symbolic regression: finding a symbolic expression that matches data from an unknown function. Although this problem is likely to be NP-hard in principle, functions of practical interest often exhibit symmetries, separability, compositionality, and other simplifying properties. In this spirit, we develop a recursive multidimensional symbolic regression algorithm that combines neural network fitting with a suite of physics-inspired techniques. We apply it to 100 equations from the Feynman Lectures on Physics, and it discovers all of them, while previous publicly available software cracks only 71; for a more difficult physics-based test set, we improve the state-of-the-art success rate from 15 to 90%.



Example: the neural network was able to find the energy of a particle in special relativity.

Predicting research trends

RESEARCH ARTICLE | PHYSICAL SCIENCES




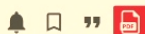
Predicting research trends with semantic and neural networks with an application in quantum physics

Mario Krenn  and Anton Zeilinger  [Authors Info & Affiliations](#)

Contributed by Anton Zeilinger, October 24, 2019 (sent for review August 19, 2019; reviewed by Ebrahim Karimi and Terry Rudolph)

January 14, 2020 | 117 (4) 1910-1916 | <https://doi.org/10.1073/pnas.1914370116>

 13,277 | 37



Significance

The corpus of scientific literature grows at an ever increasing speed. While this poses a severe challenge for human researchers, computer algorithms with access to a large body of knowledge could help make important contributions to science. Here, we demonstrate the development of a semantic network for quantum physics, denoted SEMNET, using 750,000 scientific papers and knowledge from books and Wikipedia. We use it in conjunction with an artificial neural network for predicting future research trends. Individual scientists can use SEMNET for suggesting and inspiring personalized, out-of-the-box ideas. Computer-inspired scientific ideas will play a significant role in accelerating scientific progress, and we hope that our work directly contributes to that important goal.

Particle Physics

Machine learning at the energy and intensity frontiers of particle physics

Alexander Radovic^{1*}, Mike Williams^{2*}, David Rousseau³, Michael Kagan⁴, Daniele Bonacorsi^{5,6}, Alexander Himmel⁷, Adam Aurisano⁸, Kazuhiro Terao⁴ & Taritree Wongjirad⁹

Our knowledge of the fundamental particles of nature and their interactions is summarized by the standard model of particle physics. Advancing our understanding in this field has required experiments that operate at ever higher energies and intensities, which produce extremely large and information-rich data samples. The use of machine-learning techniques is revolutionizing how we interpret these data samples, greatly increasing the discovery potential of present and future experiments. Here we summarize the challenges and opportunities that come with the use of machine learning at the frontiers of particle physics.

What is Machine Learning (ML)?

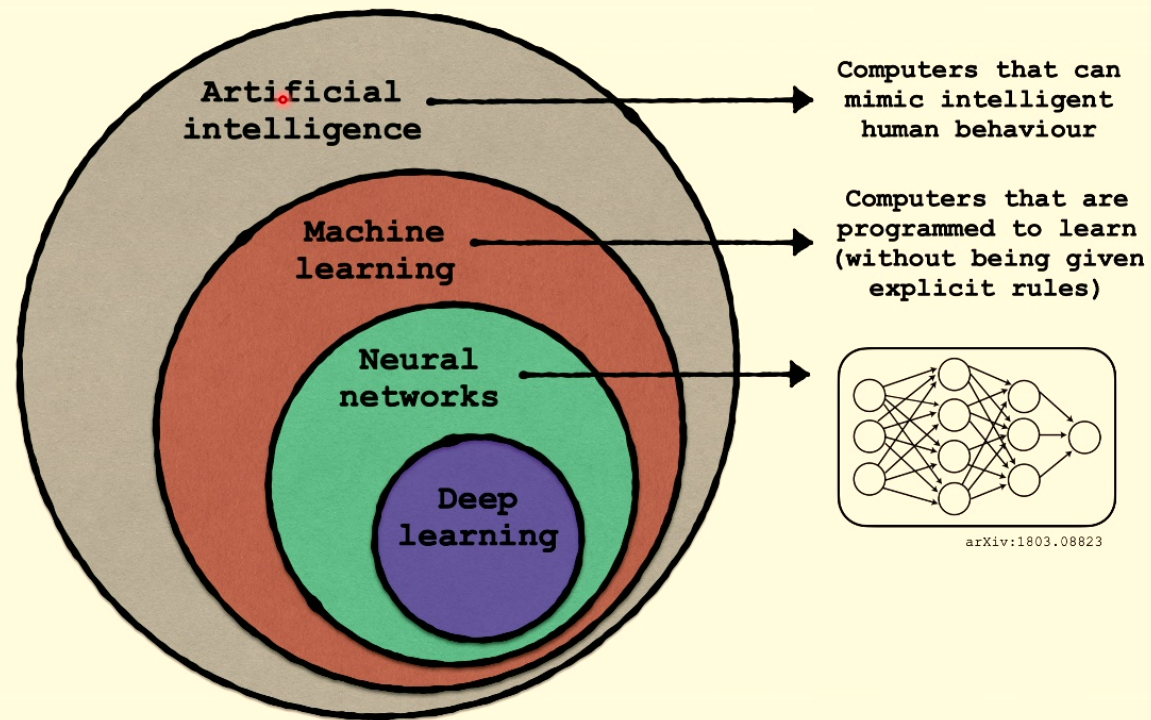
“Machine learning is a field of computer science that uses statistical techniques to give computer systems the ability to “learn” (i.e., progressively improve performance on a specific task) with data, without being explicitly programmed.”

<https://en.wikipedia.org>

“[Machine learning] is about finding out regularities in data and making use of them for fun and profit.”

L.-G. Liu, S.-H. Li and L. Wang, <http://wangleiphy.github.io>

Important remark to make



Let's get started



Machine Learning for many-body Physics

Machine Learning for
many-body physics
Lecture 1

Goal of MB Physics: Predict and explain macroscopic phenomena (features)
from microscopic quantities (Data)

ML me
↳ Sto

ML methods tend to be more fundamentally Data-Driven

↳ Start from physical measurements (Data) and then learn about
macroscopic properties (features), such as phases of Matter

$s_1 \quad s_2 \quad \dots \quad s_N$

↓↑

possibilities = 2^N

$N = O(100)$

① Supervised Learning (SL)

Data $\leftarrow D = \{(\vec{x}, \vec{y})\}$

Data point $\leftarrow \vec{x} = (x_1, x_2, \dots, x_d)$

Label $\leftarrow \vec{y} = (y_1, y_2, \dots, y_d)$

Task: fit some function $f(\vec{x})$ to \vec{y}

* Regression, \vec{y} are continuous.

* Classification, \vec{y} are Discrete.

② Unsupervised Learning (UL)

$D = \{\vec{x}\}$ (Unlabelled datapoints)

* Regression. \vec{y} are continuous.

* Classification. \vec{y} are Discrete.

② Unsupervised Learning (UL)

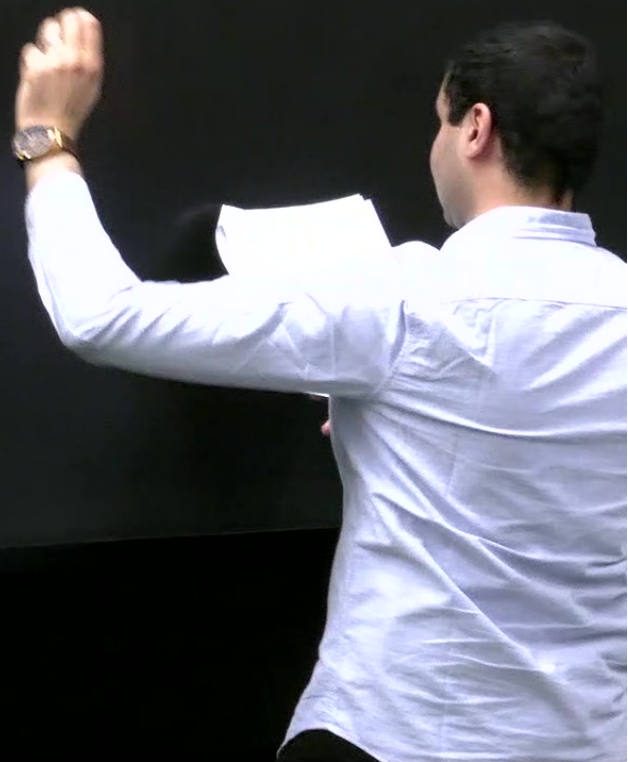
\vec{y}

$D = \{ \vec{x} \}$ (Unlabelled datapoints)

Task: Get meaningful features from the dataset.

③ Reinforcement Learning (RL)

Task:



③ Reinforcement Learning (RL)

Task: given an environment, we want to
to take an action so that a "reward" is maximized.

crete.

(U)
unlabelled datapoints)
meaningful features from
dataset.



Tentative Outline

→ SL (Lectures 1-7)

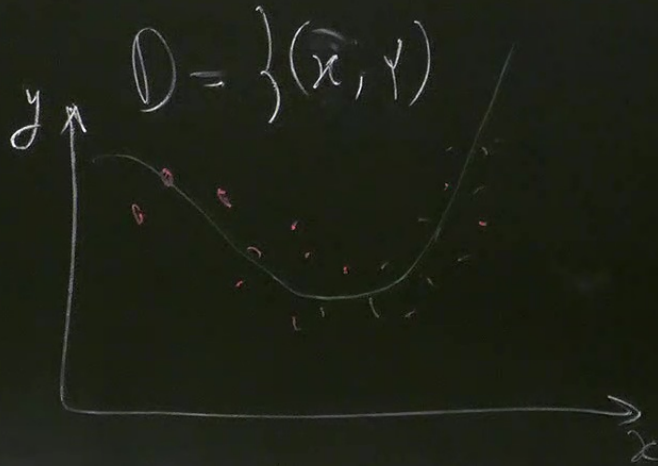
→ UL (Lectures 8-11)

→ Quantum Machine Learning. (Lecture 12)

→ Your presentations (Lecture 13)

SL examples:

→ 1D regression:

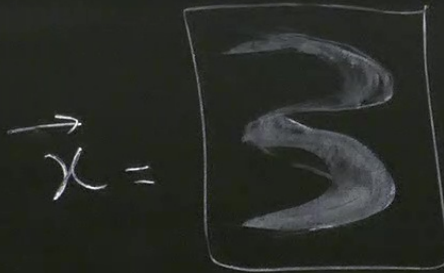


→ Handwritten digits classifier

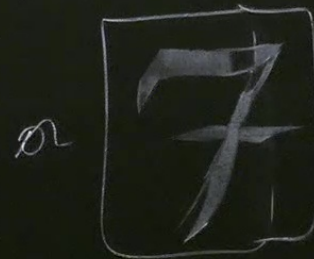
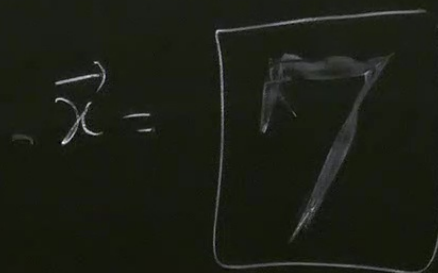
$$D = \{(\vec{x}, y)\}$$

Image of a
handwritten
digit

Corresponding
number.



, $y = 3$



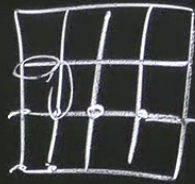
→ $y = 7$

Goal: Train a computer to determine the labels of new handwritten digits.

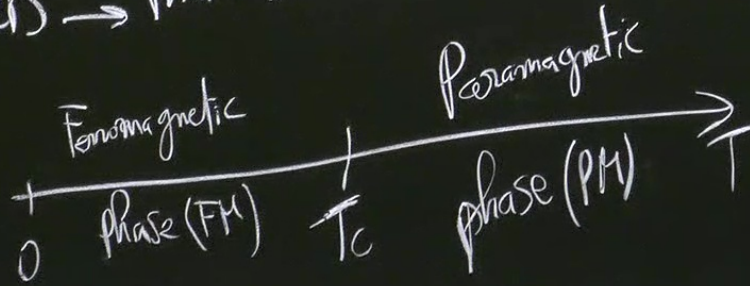
→ Ising model

$$H = - \sum_{\langle ij \rangle} J_{ij} S_i S_j$$

$$S_i = \uparrow \text{ or } \downarrow \quad (-1 \text{ or } 1)$$



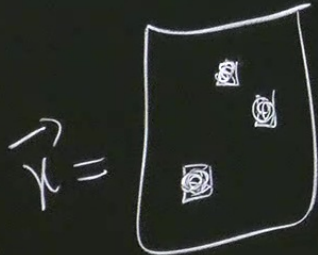
↳ I_{2D} → Phase transition at $T_c = 2.269$



$$D = \left\{ (\vec{x}, y) \right\}$$

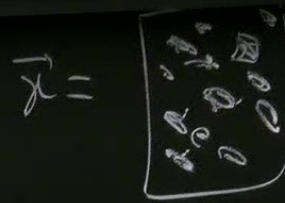
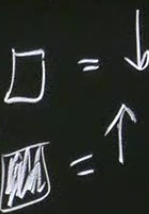
(s_1, s_2, \dots, s_N)

FM or PM
(0) (1)



$\vec{x} =$

$y = 0$ (FM)



$\vec{x} =$

$y = 1$ (PM)

$\rightarrow D = \{(\vec{x}, y)\}$
 Spin conf. \rightarrow Energy of each spin conf.

$$H_{\text{regression}} = - \sum_{i=1}^N \sum_{j=1}^N J_{ij} S_i S_j$$

\downarrow
 Unknown

Next lecture (SL)

\hookrightarrow Linear regression

\hookrightarrow Logistic regression