

**Title:** Lecture - Machine Learning, PHYS 777

**Speakers:** Mohamed Hibat Allah

**Collection/Series:** Machine Learning (Elective), PHYS 777, February 24 - March 28, 2025

**Subject:** Condensed Matter, Other

**Date:** February 25, 2025 - 9:00 AM

**URL:** <https://pirsa.org/25020015>



# Machine Learning for many-body Physics

Feb 25 - Mar 28, 2025



Machine Learning for many-body Physics

# Course Outline



**Machine Learning for Many-Body Physics**  
Feb 25 - March 28, 2025  
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.

**PSI Fellow:** Subhayan Sahu

Email: [ssahu@perimeterinstitute.ca](mailto:ssahu@perimeterinstitute.ca)

How to address me: Subhayan

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.

**Lecturer:** Mohamed Hibat-Allah

Email: [mhibatallah@uwaterloo.ca](mailto:mhibatallah@uwaterloo.ca)

How to address me: Mohamed

How to get in touch with me: just send me an email and we can work out a schedule.

**Teaching assistant:** Cole Coughlin

Email: [ccoughlin@perimeterinstitute.ca](mailto:ccoughlin@perimeterinstitute.ca)

How to address me: Cole

How to get in touch with me: by email



**Subhayan Sahu**



**Cole Coughlin**

# Schedule

- General course information** 0/3 ^

---




- Course outline
  - PDF
- References
  - TEXT
- Schedule**
  - PDF
- Lecture notes** 0/1 ∨
- Tutorials** 0/1 ∨
- Homework assignments** 0/1 ∨

	Monday	Tuesday	Wednesday	Thursday	Friday
<b>Week 1</b> (Feb 24 - 28)		Lecture 1 (Mohamed) 9:00 am - 10:00am	Lecture 2 (Mohamed) 9:00 am - 10:00am		Lecture 3 (Mohamed) 9:00 - 10:00am
<b>Week 2</b> (Mar 3 - 7)	Tutorial 1 - Linear regression 3:45 - 5:00pm	Lecture 4 (Mohamed) 9:00 - 10:00am	Lecture 5 (Mohamed) 9:00 - 10:00am	Tutorial 2 - Supervised Learning with Feedforward NNs 2:00 - 3:30 pm	Lecture 6 (Mohamed) 9:00 - 10:00am
<b>Week 3</b> (Mar 10 - 14)	Homework 1 will be posted	Lecture 7 (Mohamed) 9:00 - 10:00am		Lecture 8 (Mohamed) 9:00 - 10:00am	Tutorial 3 - Supervised Learning with Convolutional NNs 2:00 - 3:30pm
<b>Week 4</b> (Mar 17 - 21)	Homework 2 will be posted	Lecture 9 (Mohamed) 9:00 - 10:00am		Tutorial 4 - Unsupervised learning 2:00 - 3:30pm	Lecture 10 (Mohamed) 9:00 - 10:00am  Homework 1 and Presentation topic deadlines.
<b>Week 5</b> (Mar 24 - 28)	Tutorial 5 - Quantum Machine Learning 3:45 - 5:15pm	Lecture 11 (Mohamed) 9:00 - 10:00am	Lecture 12 (Mohamed) 9:00 - 10:00am	Tutorial 6 - Neural Quantum States 2:00 - 3:30pm	Lecture 13 (Mohamed) 9:00 - 10:00am



# References


## Machine Learning for Many-Body Physics Useful References

<input type="radio"/>	<b>General course information</b>	0/3	^
<input type="radio"/>	Course outline		
	 PDF		
<input checked="" type="radio"/>	<b>References</b>		
	 TEXT		
<input type="radio"/>	Schedule		
	 PDF		
<input type="radio"/>	<b>Lecture notes</b>	0/1	v
<input type="radio"/>	<b>Tutorials</b>	0/1	v
<input type="radio"/>	<b>Homework assignments</b>	0/1	v

### Books and review articles:

- Nielsen, "[Neural Networks and Deep Learning](#)"
- Goodfellow, Bengio and Courville, "[Deep Learning](#)", MIT Press (2016)
- Liu, Li and Wang, "[Lecture Note on Deep Learning and Quantum Many-Body Computation](#)" (2018)
- Mehta, Bukov, Wang, Day, Richardson, Fisher and Schwab, "A high-bias, low-variance introduction to Machine Learning for physicists", Physics Reports **810** 1-124 (2019), [arXiv:1803.08823](#)
- Carleo, Cirac, Cranmer, Daudet, Schuld, Tishby, Vogt-Maranto and Zdeborová, "Machine learning and the physical sciences", Rev. Mod. Phys. **91**, 045002 (2019), [arXiv:1903.10563](#)
- Carrasquilla and Torlai, "Neural networks in quantum many-body physics: a hands-on tutorial", [arXiv:2101.11099](#)
- Schuld and Petruccione, "Machine Learning with Quantum Computers", Second edition (2021)
- Dawid et al., "Modern applications of machine learning in quantum sciences", <https://arxiv.org/ftp/arxiv/papers/2204/2204.04198.pdf> (2022)
- **Monte Carlo methods:** Newman and Barkema, "Monte Carlo Methods in Statistical Physics" (1999)

# Lecture notes

- General course information** 0/2 ∨
- Lecture notes** 0/1 ∧
- Lecture Notes (last update Feb 24th)**  
 PDF
- Tutorials** 0/2 ∨
- 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

# 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>*

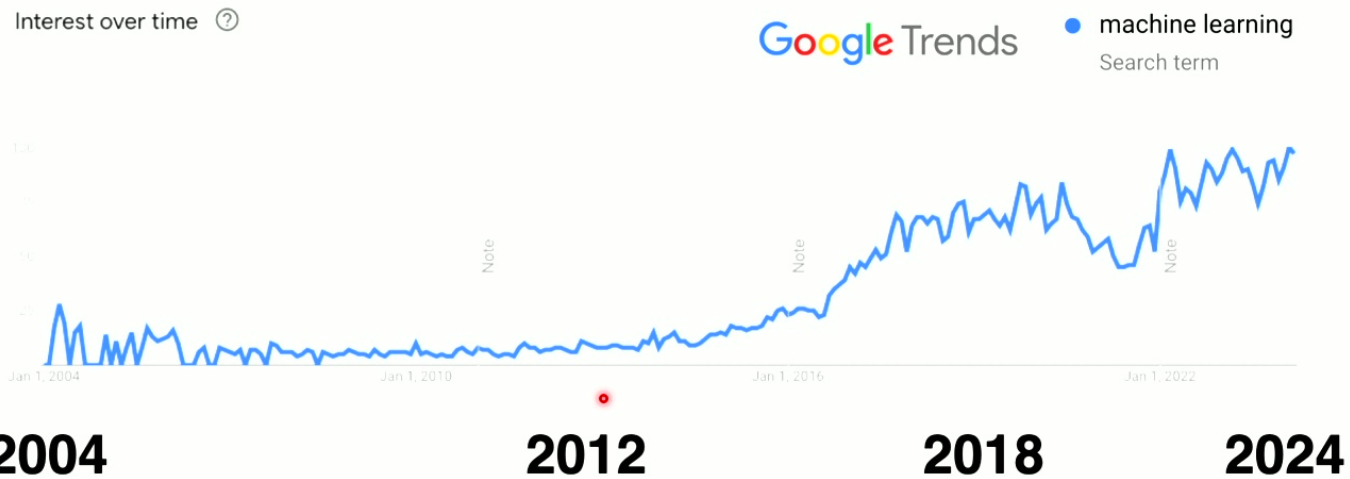
“[Machine learning] is about finding out regularities in data and making use of them for fun and profit.”

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

# Important remark to make



# Machine learning popularity



# 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 ILSVRC-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)

**nature**  
International journal of science

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) | Received: 11 November 2015  
doi:10.1038/nature16961 | Accepted: 05 January 2016  
Download Citation | 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 10:16 GMT



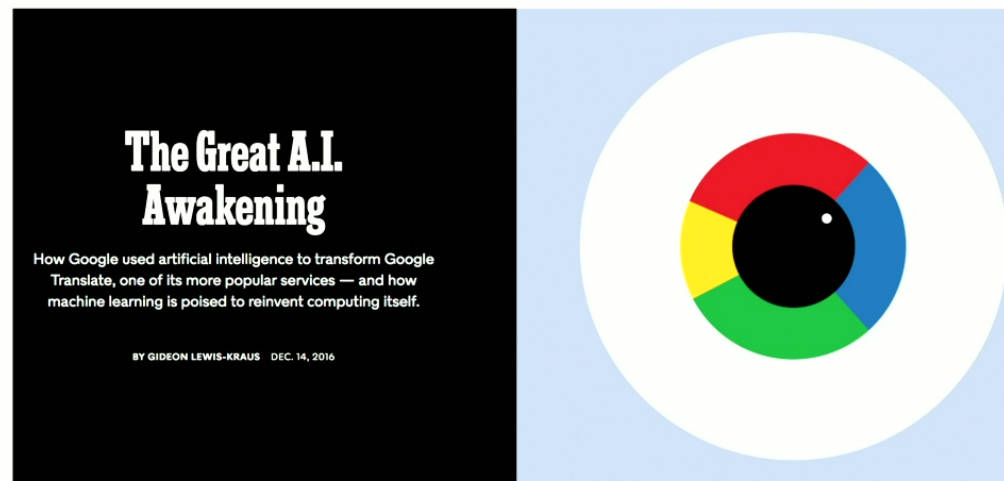
<https://www.theguardian.com>

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



# Language translation (2016)

---



**The New York Times**

# Large Language Models

---

**M** You  
Explain quantum physics in 30 words exactly?

**G** ChatGPT  
Quantum physics describes the behavior of subatomic particles, emphasizing wave-particle duality, superposition, and entanglement. It challenges classical notions, revealing a probabilistic nature where particles exist in multiple states simultaneously.



# GPT4o – Handwritten Notes to Latex

Outline:

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

2.1 Recurrent Neural Networks (RNNs)

\* Originally built for language processing

↳ Ex Machine translation, speech recognition, Music generation, DNA sequence analysis.

Sequential

→ 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.  $|\Psi\rangle = \text{RNN}$  for variational Monte Carlo.

## 2.1 Recurrent Neural Networks (RNNs)

Originally built for language processing.

Examples:

- Machine translation
- Speech recognition
- Music generation
- DNA sequence analysis

Note: These tasks involve sequential data.

## Example

"I live in France, hence I speak ..."


It is highly likely that the next word is French. So, how can we process this in a neural network to predict the word French?

Here comes Recurrent Neural Networks to the rescue.

Machine Learning for many-body Physics

# Dall-E

---



<https://labs.openai.com>

**A 3d realistic render of an astronaut walking on Mars**

# AlphaProof and AlphaGeometry

---

RESEARCH

AI achieves silver-medal standard  
solving International Mathematical  
Olympiad problems

25 JULY 2024

AlphaProof and AlphaGeometry teams

Google DeepMind

CS479/679

17

# 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>

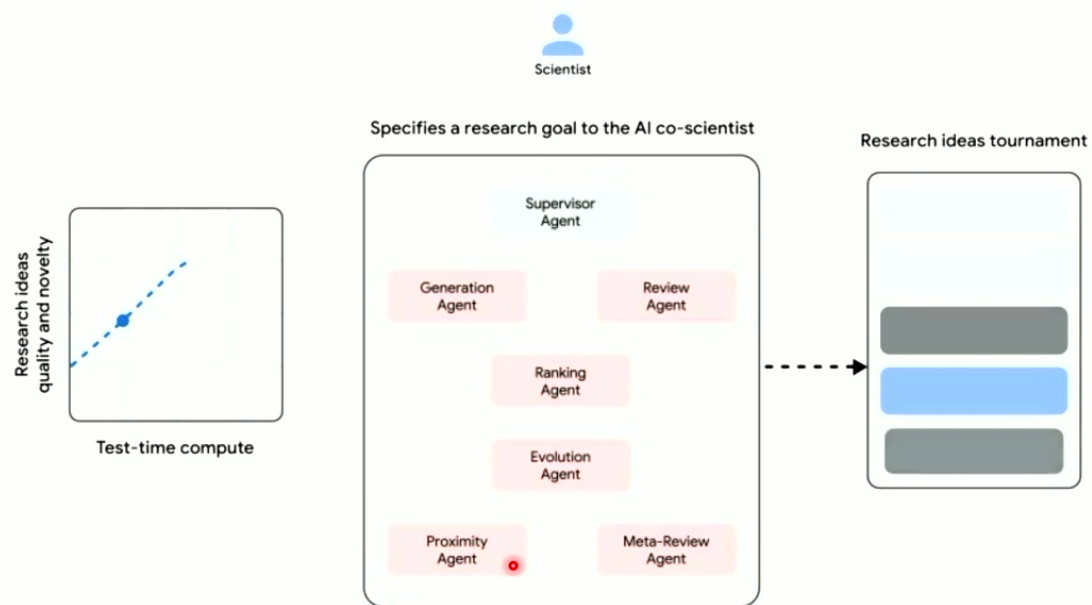
13,277 | 37



### 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.

# AI co-scientist




<https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>



# < ML | Physics >

---



# Nobel prize in Physics 2024

---



Ill. Niklas Elmehed © Nobel Prize Outreach

**John J. Hopfield**

Prize share: 1/2



Ill. Niklas Elmehed © Nobel Prize Outreach

**Geoffrey Hinton**

Prize share: 1/2

Lecture 1. Machine Learning  
for many-body physics

→ Computational methods are often useful in many-body physics.

↳ Solving Schrödinger equation

$$\hat{H}|\psi\rangle = E|\psi\rangle$$

↳ Finding Boltzmann distribution

$$P(\vec{\sigma}) = \frac{\exp(-\beta E(\vec{\sigma}))}{Z}$$

↳ Examples: → Monte Carlo

→ Tensor Networks

→ Exact diagonalization, ...



Important points about many-body physics:

↳ Relies on many-particles interacting.

↳ Needed to describe phases of matter (excluding single quasi-particle theories).

ML methods tend to be more fundamentally data-driven!

Example: → Start from physical measurements (data)  
and learn macroscopic properties (features)  
such as phases of matter.



# ① Supervised Learning (SL)

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

↳ Datapoints  $\vec{x} = (x_1, x_2, \dots, x_{d_x})$

↳ Labels  $\vec{y} = (y_1, y_2, \dots, y_{d_y})$

Task. "fit" some function  $f(\vec{x}) \rightarrow \vec{y}$

↳ Regression: when  $\vec{y}$  continuous.

↳ Classification: when  $\vec{y}$  are discrete.





→ Exact (unlabeled) points, etc.

## ② Unsupervised Learning (UL)

$$D = \{ \vec{x} \} \quad (\text{unlabelled datapoints})$$

Task Extract meaningful features from the dataset



### ③ Reinforcement Learning (RL)

↳ Given environment, take actions that result in "reward" being maximized.



## Tentative outline:

- \* SL (Lectures #1-7)
- \* UL (Lectures #8, 9 and 11)
- \* Quantum ML (Lecture #10)
- \* Neural Quantum states (#12)

\* Recent advances in (ML | many-body phys) + ML ethics

SL examples:

↳ Regression:

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

$$(dx = 1 = dy)$$



→ Goal: determine a curve to describe the data.



phase 1  
phase 2

Ex 2


Classification of Handwritten digits


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

Image of a handwritten digit

The corresponding number

$\vec{x} =$   ;  $y = 3$

$\vec{x} =$  

or 

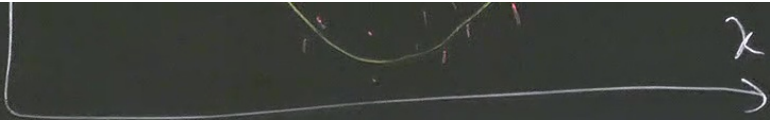
;  $y = 7$



E.g. 3: Ising model.

$$H = -J \sum_{\langle i, j \rangle} s_i s_j$$

$$s_i = \begin{matrix} \uparrow \\ \downarrow \end{matrix} \text{ or } \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ or } \begin{pmatrix} +1 \\ -1 \end{pmatrix}$$



Recall in 2D, the Ising model has a phase transition at  $T_c/J \approx 2.269$

→ For  $T < T_c$  Ferromagnetic phase (FM)



→ For  $T > T_c$  Paramagnetic phase (PM)




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

$(S_{11}, S_{21}, S_{31}, \dots, S_{N1})$

$y=0$  (FM)  
 $y=1$  (PM)



$y=0$  (FM)





Eq. 4 regression (Learn the Ising Hamiltonian couplings)

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

(Tutorial 1).

spin config

Energy of  
a spin config

$$H = \sum_{i=1}^N \sum_{j=1}^N J_{ij} S_i S_j$$

regression

Goal: Use  $D$  to determine the matrix of coupling  $J_{ij}$ .



Eq. 4 regression (Learn the Ising Hamiltonian couplings)

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

(Tutorial 1).

spin config

$$\vec{x} = (S_1, S_2, \dots, S_N)$$

Energy of a spin config

Goal:



## Next lecture:

- ↳ Linear regression
- ↳ Gradient Descent (GD)
- ↳ Logistic regression (classification)

Linear regression (LR):

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

$f(\vec{x})$  to  $y(\vec{x})$

$d_y = 1$

$$f(\vec{x}) = \sum_{j=1}^{d_x} w_j x_j = \vec{w}^T \cdot \vec{x}$$

$$\vec{w} = (w_1, \dots, w_{d_x})$$



$M = \#$  of datapoints in  $\mathcal{D}$ .

$x_j^{(i)}$  =  $j^{\text{th}}$  element of datapoint  $x^{(i)}$ ,

$1 \leq i \leq M; 1 \leq j \leq d_x$ .

$X \in \mathbb{R}^{M \times d_x}$

$X_{ij} = x_j^{(i)}$

$\mathcal{L} = \text{Error}(f(\vec{x}), \vec{y})$