Несортированные ссылки по Deep Learning neural network

А.Е. Шевель

Created: 1 Dec 2021

Corrected: 31 Dec 2021

31 Jan 2022

11 Feb 2022

https://www.nature.com/articles/s41586-021-04301-9??utm_medium=affiliate&utm_source=commission_junction&utm_campaign=CONR_PF018_ECOM_GL_PHSS_ALWYS_PRODUCT&utm_content=productdatafeed&utm_term=PID100095187&CJEVENT=1fdc124e90d711ec807d01680a18050d - Magnetic control of tokamak plasmas through deep reinforcement learning (16 Feb)

https://ai.stackexchange.com/questions/11405/what-are-the-purposes-of-autoencoders - What are the purposes of autoencoders?

https://iopscience.iop.org/journal/2632-2153 - Machine Learning: Science and Technology is a multidisciplinary open access journal that bridges the application of machine learning across the sciences with advances in machine learning methods and theory as motivated by physical insights.

https://sparks.cern/welcome-cern-sparks - CERN launched an annual two-day multidisciplinary science innovation forum and public event. Sparks! the serendipity forum at CERN, brings together renowned scientists from diverse fields around the world, along with decision makers, representatives of industry, philanthropists, ethicists and the public to bring a novel, multi-faceted approach to addressing some of the big questions of our time.

The goal? To foster a new community and develop a platform to spark innovation in issues related to science, technology, engineering and mathematics that are relevant to society, and further CERN's mission towards peace in science.

Sparks! is starting with a cycle of three pilot events from 2021 to 2023, each focusing on a single theme to test the concept. Experience gained from this cycle will lead to a yearly event with multiple themes addressed each year. The event will become a flagship for CERN's new Science Gateway which is scheduled to open its doors to the public in 2023. The theme of the first edition is future intelligence.

https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning - Supervised vs. Unsupervised Learning: What’s the Difference?

https://indico.cern.ch/event/1074354/contributions/4518161/attachments/2322808/3955769/DeepLearning_IceCream_2021.pdf - presentation at CERN

https://en.wikipedia.org/wiki/Generative_adversarial_network GAN https://arxiv.org/pdf/2108.07258.pdf - On the Opportunities and Risks of

Foundation Models (120 pages)

https://towardsdatascience.com/

https://www.wired.com/2016/12/2016-year-deep-learning-took-internet/ - 2016: The Year That Deep Learning Took Over the Internet

Artificial intelligence is remaking tech giants from the inside out, and it's rapidly spreading to the rest of the world.

https://arxiv.org/pdf/2102.09924.pdf - Gradient descent optimization algorithms are the standard ingredients that are used to train artificial neural networks (ANNs). Even though a huge number of numerical simu- lations indicate that gradient descent optimization methods do indeed convergence in the training of ANNs, until today there is no rigorous theoretical analysis which proves (or dis- proves) this conjecture. In particular, even in the case of the most basic variant of gradient descent optimization algorithms, the plain vanilla gradient descent method, it remains an open problem to prove or disprove the conjecture that gradient descent converges in the training of ANNs.

https://www.baeldung.com/cs/epoch-neural-networks - in this tutorial, we'll learn about the meaning of an epoch in neural networks. Then we'll investigate the relationship between neural network training convergence and the number of epochs. Finally, we'll try to understand how we use early stopping to get better generalizing models.

https://www.forbes.com/sites/forbestechcouncil/2021/05/13/the-collective-power-of-swarm-intelligence-in-ai-and-robotics/?sh=1baa6062252f - The Collective Power Of Swarm Intelligence In AI And Robotics

Swarm intelligence is the collective behavior of decentralized, self-organized systems (natural or artificial) that can maneuver quickly in a coordinated fashion. In nature, this closed-loop, collaborative behavior is unique within each species. Ants lay down pheromones directing each other to resources, bees use vibrations, fish feel tremors in the water and birds detect motions spreading through the flock.

https://www.nature.com/articles/s41586-021-03583-3 - Swarm Learning for decentralized and confidential clinical machine learning

... to facilitate the integration of any medical data from any data owner worldwide without violating privacy laws, we introduce Swarm Learning—a decentralized machine-learning approach that unites edge computing, blockchain-based peer-to-peer networking and coordination while maintaining confidentiality without the need for a central coordinator, thereby going beyond federated learning. To illustrate the feasibility of using Swarm Learning to develop disease classifiers using distributed data, we chose four use cases of heterogeneous diseases (COVID-19, tuberculosis, leukaemia and lung pathologies). With more than 16,400 blood transcriptomes derived from 127 clinical studies with non-uniform distributions of cases and controls and substantial study biases, as well as more than 95,000 chest X-ray images, we show that Swarm Learning classifiers outperform those developed at individual sites. In addition, Swarm Learning completely fulfils local confidentiality regulations by design. We believe that this approach will notably accelerate the introduction of precision medicine.

https://towardsdatascience.com/tagged/swarm-intelligence - Swarm Intelligence (many aspects blog)

https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&cad=rja&uact=8&ved=2ahUKEwjcxdftgIH1AhWNrosKHQsUBt4QFnoECBoQAQ&url=https%3A%2F%2Fwww.mdpi.com%2F2504-2289%2F5%2F3%2F36%2Fpdf&usg=AOvVaw29NRKkWZFE7InYs7Lql7Zh -

Nowadays, the high-dimensionality of data causes a variety of problems in machine learning. It is necessary to reduce the feature number by selecting only the most relevant of them. Different approaches called Feature Selection are used for this task. In this paper, we propose a Feature Selection method that uses Swarm Intelligence techniques. Swarm Intelligence algorithms perform optimization by searching for optimal points in the search space. We show the usability of these techniques for solving Feature Selection and compare the performance of five major swarm algorithms: Particle Swarm Optimization, Artificial Bee Colony, InvasiveWeed Optimization, Bat Algorithm, and Grey Wolf Optimizer. The accuracy of a decision tree classifier was used to evaluate the algorithms. It turned out that the dimension of the data can be reduced about two times without a loss in accuracy. Moreover, the accuracy increased when abandoning redundant features. Based on our experiments GWO turned out to be the best. It has the highest ranking on different datasets, and its average iteration number to find the best solution is 30.8. ABC obtained the lowest ranking on high-dimensional datasets.

https://obamawhitehouse.archives.gov/the-press-office/2013/04/02/fact-sheet-brain-initiative - BRAIN Initiative (April 02, 2013)

https://www.braininitiative.org/ - The BRAIN Initiative® seeks to deepen understanding of the inner workings of the human mind and to improve how we treat, prevent, and cure disorders of the brain.

https://www.humanbrainproject.eu/en/ - The Human Brain Project (HBP) is one of the three FET (Future and Emerging Technology) Flagship projects. Started in 2013, it is one of the largest research projects in the world . More than 500 scientists and engineers at over than 140 universities, teaching hospitals, and research centres across Europe come together to address one of the most challenging research targets – the human brain.

https://research.ibm.com/artificial-intelligence - The unreasonable effectiveness of deep learning in artificial intelligence [PNAS December 1, 2020]

https://cogcomp.seas.upenn.edu/ - Our research focuses on the computational foundations of intelligent behavior. We develop theories and systems pertaining to intelligent behavior using a unified methodology -- at the heart of which is the idea that learning has a central role in intelligence.

Our work spans several aspects of this problem -- from theoretical questions in machine learning, knowledge representation and reasoning to experimental paradigms and large scale system development -- and draws on methods from theoretical computer science, probability and statistics, artificial intelligence, linguistics and experimental computer science.

https://www.pnas.org/content/117/48/30033 -

https://www.youtube.com/watch?v=rt4806DzfUY

Google colab: про Юпитер для программистов

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

https://towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464 - The mostly complete chart of Neural Networks, explained

https://www.foreignaffairs.com/articles/united-states/2021-12-10/soon-hackers-wont-be-human

Soon hackers will not be human being.

https://neptune.ai/blog/6-gan-architectures - 6 GAN Architectures You Really Should Know

Generative Adversarial Networks (GANs) were first introduced in 2014 by Ian Goodfellow et. al. and since then this topic itself opened up a new area of research.

Within a few years, the research community came up with plenty of papers on this topic some of which have very interesting names :). You have CycleGAN, followed by BiCycleGAN, followed by ReCycleGAN and so on.

With the invention of GANs, Generative Models had started showing promising results in generating realistic images. GANs has shown tremendous success in Computer Vision. In recent times, it started showing promising results in Audio, Text as well.

Some of the most popular GAN formulations are:

- Transforming an image from one domain to another (CycleGAN),
- Generating an image from a textual description (text-to-image),
- Generating very high-resolution images (ProgressiveGAN) and many more.

https://spectrum.ieee.org/neuromorphic-computing-with-lohi2

Intel's Neuromorphic Chip Gets A Major Upgrade Loihi 2 packs 1 million neurons in a chip half the size of its predecessor

Many AIs may depend on things called neural networks, but there's very little about them that works in the way human and animal brains do. Intel has been experimenting with computers that think more like a brain does for several years now, racking up some impressive if quirky results with their Loihi neuromorphic chip. Now Loihi is getting its first upgrade, and it's a pretty big one. Using a manufacturing process called Intel 4 that's not yet available for commercial chips, the company packed in up to eight-times as many artificial neurons into a chip that's half the area of Loihi. That, and a host of changes motivated by the past few years of experiments, make the Loihi 2 faster and more flexible, says Mike Davies director of Intel's neuromorphic computing lab.

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

Deep learning (also known as deep structured learning) is part of a broader family of machine learning methods based on artificial neural networks with representation learning. Learning can be supervised, semi-supervised or unsupervised.

Deep-learning architectures such as deep neural networks, deep belief networks, deep reinforcement learning, recurrent neural networks and convolutional neural networks have been applied to fields including computer vision, speech recognition, natural language processing, machine translation, bioinformatics, drug design, medical image analysis, material inspection and board game programs, where they have produced results comparable to and in some cases surpassing human expert performance.

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

Artificial neural networks (ANNs), usually simply called neural networks (NNs), are computing systems inspired by the biological neural networks that constitute animal brains.

An ANN is based on a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal to other neurons. An artificial neuron receives a signal then processes it and can signal neurons connected to it. The "signal" at a connection is a real number, and the output of each neuron is computed by some non-linear function of the sum of its inputs. The connections are called edges. Neurons and edges typically have a weight that adjusts as learning proceeds. The weight increases or decreases the strength of the signal at a connection. Neurons may have a threshold such that a signal is sent only if the aggregate signal crosses that threshold. Typically, neurons are aggregated into layers. Different layers may perform different transformations on their inputs. Signals travel from the first layer (the input layer), to the last layer (the output layer), possibly after traversing the layers multiple times.

https://iml-wg.github.io/HEPML-LivingReview/

A Living Review of Machine Learning for Particle Physics

Modern machine learning techniques, including deep learning, is rapidly being applied, adapted, and developed for high energy physics. The goal of this document is to provide a nearly comprehensive list of citations for those developing and applying these approaches to experimental, phenomenological, or theoretical analyses. As a living document, it will be updated as often as possible to incorporate the latest developments. A list of proper (unchanging) reviews can be found within. Papers are grouped into a small set of topics to be as useful as possible. Suggestions are most welcome.

The purpose of this note is to collect references for modern machine learning as applied to particle physics. A minimal number of categories is chosen in order to be as useful as possible. Note that papers may be referenced in more than one category. The fact that a paper is listed in this document does not endorse or validate its content - that is for the community (and for peer-review) to decide. Furthermore, the classification here is a best attempt and may have flaws - please let us know if (a) we have missed a paper you think should be included, (b) a paper has been misclassified, or (c) a citation for a paper is not correct or if the journal information is now available. In order to be as useful as possible, this document will continue to evolve so please check back before you write your next paper. If you find this review helpful, please consider citing it using \cite{hepmllivingreview} in HEPML.bib.

https://towardsdatascience.com/classical-neural-network-what-really-are-nodes-and-layers-ec51c6122e09 - Classical Neural Network: What really are Nodes and Layers?

https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53 - A Comprehensive Guide to Convolutional Neural Networks — the ELI5 way

https://medium.com/@ngocson2vn/a-gentle-explanation-of-backpropagation-in-convolutional-neural-network-cnn-1a70abff508b - A gentle explanation of Backpropagation in Convolutional Neural Network (CNN)

https://qiskit.org/textbook/ch-machine-learning/machine-learning-qiskit-pytorch.html

https://community.qiskit.org/textbook/ch-machine-learning/machine-learning-qiskit-pytorch.html

Hybrid quantum-classical Neural Networks with PyTorch and Qiskit

Machine learning (ML) has established itself as a successful interdisciplinary field which seeks to mathematically extract generalizable information from data. Throwing in quantum computing gives rise to interesting areas of research which seek to leverage the principles of quantum mechanics to augment machine learning or vice-versa. Whether you're aiming to enhance classical ML algorithms by outsourcing difficult calculations to a quantum computer or optimise quantum algorithms using classical ML architectures - both fall under the diverse umbrella of quantum machine learning (QML).

https://towardsdatascience.com/a-gentle-introduction-to-neural-networks-series-part-1-2b90b87795bc A Gentle Introduction To Neural Networks Series — Part 1

https://www.appec.org/news/machine-learning-optimized-design-of-experiments

For over a century now, physicists have designed instruments to detect elementary particles, and radiation in general, exploiting cutting-edge technologies, and in some cases developing entirely new ones. As the complexity of the apparatuses and of the required tasks grew, so did our inventiveness. This has brought a stream of new developments, which culminated in the past two decades with the construction and operation of giant detectors like ATLAS and CMS, which are mind boggling instruments.

https://cerncourier.com/a/designing-an-ai-physicist/

Designing an AI physicist 2 September 2021

Jesse Thaler argues that particle physicists must go beyond deep learning and design AI capable of deep thinking.

If we don't exploit the full power of AI, we will not maximise the discovery potential of the LHC and other experiments

Ultimately, we need to merge the insights gained from artificial intelligence and physics intelligence. If we don't exploit the full power of AI, we will not maximise the discovery potential of the LHC and other experiments. But if we don't build trustable AI, we will lack scientific rigour. Machines may never think like human physicists, and human physicists will certainly never match the computational ability of AI, but together we have enormous potential to learn about the fundamental structure of the universe.

Jesse Thaler is a professor at MIT and director of the US National Science Foundation's Institute for Artificial Intelligence and Fundamental Interactions.

https://cerncourier.com/a/deep-learning-tailors-supersymmetry-searches/

Deep learning tailors supersymmetry searches

23 February 2021

A report from the CMS experiment

https://cerncourier.com/a/hunting-anomalies-with-an-ai-trigger/

Hunting anomalies with an AI trigger

31 August 2021

Jennifer Ngadiuba and Maurizio Pierini describe how ‘unsupervised’ machine learning could keep watch for signs of new physics at the LHC that have not yet been dreamt up by physicists.

https://cerncourier.com/a/learning-to-detect-new-top-quark-interactions/

Learning to detect new top-quark interactions 4 October 2021

A report from the CMS experiment

Ever since its discovery in 1995 at the Tevatron, the top quark has been considered to be a highly effective probe of new physics. A key reason is that the last fundamental fermion predicted by the Standard Model (SM) has a remarkably high mass, just a sliver under the Higgs vacuum expectation value divided by the square root of two, implying a Yukawa coupling close to unity. This has far-reaching implications: the top quark impacts the electroweak sector significantly through loop corrections, and may couple preferentially to new massive states. But while the top quark may represent a window into new physics, we cannot know a priori whether new massive particles could ever be produced at the LHC, and direct searches have so far been inconclusive. Model-inde-pendent measurements carried out within the framework of effective field theory (EFT) are therefore becoming increasingly important as a means to make the most of the wealth of precision measurements at the LHC. This approach makes it possible to systematically correlate sparse deviations observed in different measurements, in order to pinpoint any anomalies in top-quark couplings that might arise from unknown massive particles.

https://www.tensorflow.org - Whether you're an expert or a beginner, TensorFlow is an end-to-end platform that makes it easy for you to build and deploy ML models.

TensorFlow is an end-to-end open source platform for machine learning

TensorFlow makes it easy for beginners and experts to create machine learning models. See the sections below to get started.

https://blog.tensorflow.org/2021/04/reconstructing-thousands-of-particles-in-one-go-at-cern-lhc.html tensorflow blog

TensorFlow is a free and open-source software library for machine learning and artificial intelligence. It can be used across a range of tasks but has a particular focus on training and inference of deep neural networks.[4][5]

Tensorflow is a symbolic math library based on dataflow and differentiable programming. It is used for both research and production at Google.[6][7][8]

TensorFlow was developed by the Google Brain team for internal Google use. It was released under the Apache License 2.0 in 2015.[1][9]

TensorFlow is Google Brain's second-generation system. Version 1.0.0 was released on February 11, 2017.[14] While the reference implementation runs on single devices, TensorFlow can run on multiple CPUs and GPUs (with optional CUDA and SYCL extensions for general-purpose computing on graphics processing units).[15] TensorFlow is available on 64-bit Linux, macOS, Windows, and mobile computing platforms including Android and iOS.

Its flexible architecture allows for the easy deployment of computation across a variety of platforms (CPUs, GPUs, TPUs), and from desktops to clusters of servers to mobile and edge devices.

TensorFlow computations are expressed as stateful dataflow graphs. The name TensorFlow derives from the operations that such neural networks perform on multidimensional data arrays, which are referred to as tensors. During the Google I/O Conference in June 2016, Jeff Dean stated that 1,500 repositories on GitHub mentioned TensorFlow, of which only 5 were from Google.[16]

In December 2017, developers from Google, Cisco, RedHat, CoreOS, and CaiCloud introduced Kubeflow at a conference. Kubeflow allows operation and deployment of TensorFlow on Kubernetes.

In March 2018, Google announced TensorFlow.js version 1.0 for machine learning in JavaScript.[17]

In Jan 2019, Google announced TensorFlow 2.0.[18] It became officially available in Sep 2019.[19]

In May 2019, Google announced TensorFlow Graphics for deep learning in computer graphics.[20]

http://oproject.org/pages/CERN%20GPU%20Condor.html In this tutorial I will teach how to setup your environment, write the scripts and submits your jobs to use the GPUs in HTCondor batch system at CERN. The objective is to train machine learning models using or not jupyter notebooks. The official tutorial can be found at https://batchdocs.web.cern.ch/index.html I am taking only the required for GPUs proccesing.

https://pytorch.org - An open source machine learning framework that accelerates the path from research prototyping to production deployment.

PyTorch is an open source machine learning library based on the Torch library,[3][4][5] used for applications such as computer vision and natural language processing,[6] primarily developed by Facebook's AI Research lab (FAIR).[7][8][9] It is free and open-source software released under the Modified BSD license. Although the Python interface is more polished and the primary focus of development, PyTorch also has a C++ interface.[10]

A number of pieces of deep learning software are built on top of PyTorch, including Tesla Autopilot,[11] Uber's Pyro,[12] Hugging Face's Transformers,[13] PyTorch Lightning,[14][15] and Catalyst.[16][17]

PyTorch provides two high-level features:[18]

Tensor computing (like NumPy) with strong acceleration via graphics processing units (GPU)

Deep neural networks built on a type-based automatic differentiation system

https://realpython.com/pytorch-vs-tensorflow/ - PyTorch vs TensorFlow: What’s the difference? Both are open source Python libraries that use graphs to perform numerical computation on data. Both are used extensively in academic research and commercial code. Both are extended by a variety of APIs, cloud computing platforms, and model repositories.

If they're so similar, then which one is best for your project?

In this tutorial, you'll learn:

- What the differences are between PyTorch and TensorFlow
- What tools and resources are available for each
- How to choose the best option for your specific use case

https://en.wikipedia.org/wiki/Google_Brain - Google Brain is a deep learning artificial intelligence research team under the umbrella of Google AI, a research division at Google dedicated to artificial intelligence. Formed in 2011, Google Brain combines open-ended machine learning research with information systems and large-scale computing resources.[1] The team has created tools such as TensorFlow, which allow for neural networks to be used by the public, with multiple internal AI research projects.[2] The team aims to create research opportunities in machine learning and natural language processing.[2]

https://openai.com - Developers can now join the waitlist to access OpenAI Codex, our AI system that translates natural language into code.

https://www.simplilearn.com/keras-vs-tensorflow-vs-pytorch-article#what_is_keras – Keras is an effective high-level neural network Application Programming Interface (API) written in Python. This open-source neural network library is designed to provide fast experimentation with deep neural networks, and it can run on top of CNTK, TensorFlow, and Theano.

https://machinelearningmastery.com/tutorial-first-neural-network-python-keras/ Your First Deep Learning Project in Python with Keras Step-By-Step

https://www.analyticssteps.com/blogs/how-decide-which-activation-function-and-loss-function-use How to decide which Activation Function and Loss Function to use?

https://mxnet.apache.org - A truly open source deep learning framework suited for flexible research prototyping and production.

https://en.wikipedia.org/wiki/Neural_cryptography - Artificial neural networks are well known for their ability to selectively explore the solution space of a given problem. This feature finds a natural niche of application in the field of cryptanalysis. At the same time, neural networks offer a new approach to attack ciphering algorithms based on the principle that any function could be reproduced by a neural network, which is a powerful proven computational tool that can be used to find the inverse-function of any cryptographic algorithm.

The ideas of mutual learning, self learning, and stochastic behavior of neural networks and similar algorithms can be used for different aspects of cryptography, like public-key cryptography, solving the key distribution problem using neural network mutual synchronization, hashing or generation of pseudo-random numbers.

Another idea is the ability of a neural network to separate space in non-linear pieces using "bias". It gives different probabilities of activating the neural network or not. This is very useful in the case of Cryptanalysis.

Two names are used to design the same domain of research: Neuro-Cryptography and Neural Cryptography.

https://www.cnews.ru/news/top/2021-11-02_google_sozdaet_nastoyashchij 02 ноября 2021

В компании Google ведутся разработки проекта Pathways, новой архитектуры нейронных сетей, которые смогут выполнять сразу множество разных задач и осваивать новые.

https://www.youtube.com/watch?v=5tvmMX8r_OM -- MIT introduction https://www.youtube.com/watch?v=f9oDe4Yq4E0&ab_channel=3Blue1BrowntranslatedbySciberia обучение распознаванию цифр

https://www.youtube.com/c/foo52ru   foo52ru ТехноШаман

https://www.youtube.com/watch?v=HWN_DVRB1G4&t=17&ab_channel=AndreySozykin пример с google colab

https://www.nature.com/articles/s41534-021-00443-w - Article Open Access Published: 15 July 2021

Quantum-inspired machine learning on high-energy physics data

https://qiskit.org/documentation/machine-learning/getting_started.html - Qiskit Machine Learning depends on the main Qiskit package which has its own Qiskit Getting Started detailing the installation options for Qiskit and its supported environments/platforms. You should refer to that first. Then the information here can be followed which focuses on the additional installation specific to Qiskit Machine Learning.

Qiskit Machine Learning has some functions that have been made optional where the dependent code and/or support program(s) are not (or cannot be) installed by default. Those are PyTorch and Sparse. See Optional installs for more information.

https://qiskit.org/textbook/ch-machine-learning/machine-learning-qiskit-pytorch.html#how - The background presented here on classical neural networks is included to establish relevant ideas and shared terminology; however, it is still extremely high-level. If you'd like to dive one step deeper into classical neural networks, see the well made video series by youtuber 3Blue1Brown. Alternatively, if you are already familiar with classical networks, you can skip to the next section.

https://qiskit.org/textbook/ch-states/representing-qubit-states.html - Representing Qubit States

You now know something about bits, and about how our familiar digital computers work. All the complex variables, objects and data structures used in modern software are basically all just big piles of bits. Those of us who work on quantum computing call these classical variables. The computers that use them, like the one you are using to read this article, we call classical computers.

In quantum computers, our basic variable is the qubit: a quantum variant of the bit. These have exactly the same restrictions as normal bits do: they can store only a single binary piece of information, and can only ever give us an output of 0 or 1. However, they can also be manipulated in ways that can only be described by quantum mechanics. This gives us new gates to play with, allowing us to find new ways to design algorithms.

http://quantum.cern/sites/default/files/2021-10/Style-based%20quantum%20generative%20adversarial%20networks%20for%20Monte%20Carlo%20events.pdf - We propose and assess an alternative quantum generator architecture in the context of generative adversarial learning for Monte Carlo event generation, used to simulate particle physics processes at the Large Hadron Collider (LHC). We validate this methodology by implementing the quantum network on artificial data generated from known underlying distributions. The network is then applied to Monte Carlo-generated datasets of specific LHC scattering processes. The new quantum generator architecture leads to an improvement in state-of-the-art implementations while maintaining shallow-depth networks. Moreover, the quantum generator successfully learns the underlying distribution functions even if trained with small training sample sets; this is particularly interesting for data augmentation applications. We deploy this novel methodology on two different quantum hardware architectures, trapped-ion and superconducting technologies, to test its hardware-independent viability.

https://docs.google.com/presentation/d/1bwJDRC777rAf00Drthi9yT2c9b0MabWO5ZlksfvFzx8/edit#slide=id.gf171287819_0_165 - State of AI Report 2021 – ONLINE

https://yandex.ru/lab/yalm - Балабоба демонстрирует, как с помощью нейросетей семейства YaLM можно продолжать тексты на любую тему, сохраняя связность и заданный стиль. Здесь эти нейросети используются для развлечения, но разрабатывались они для серьёзных задач — об этом можно почитать тут.

У Балабобы нет своего мнения или знания. Он умеет только подражать — писать тексты так, чтобы они были максимально похожи на реальные тексты из интернета.

https://deepai.org/ -- The best place to use and discover the latest artificial intelligence.

https://openai.com/ - OpenAI is an AI research and deployment company. Our mission is to ensure that artificial general intelligence benefits all of humanity.

https://deepmind.com/ - Our teams research and build safe AI systems. We're committed to solving intelligence, to advance science and benefit humanity.

https://www.nature.com/articles/s41586-021-03819-2 - Highly accurate protein structure prediction with AlphaFold

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort1,2,3,4, the structures of around 100,000 unique proteins have been determined5, but this represents a small fraction of the billions of known protein sequences6,7. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bio

https://www.nature.com/articles/s41467-021-27606-9#Sec8

Nature 16 Dec 2021

Face detection in untrained deep neural networks

Abstract

Face-selective neurons are observed in the primate visual pathway and are considered as the basis of face detection in the brain. However, it has been debated as to whether this neuronal selectivity can arise innately or whether it requires training from visual experience. Here, using a hierarchical deep neural network model of the ventral visual stream, we suggest a mechanism in which face-selectivity arises in the complete absence of training. We found that units selective to faces emerge robustly in randomly initialized networks and that these units reproduce many characteristics observed in monkeys. This innate selectivity also enables the untrained network to perform face-detection tasks. Intriguingly, we observed that units selective to various non-face objects can also arise innately in untrained networks. Our results imply that the random feedforward connections in early, untrained deep neural networks may be sufficient for initializing primitive visual selectivity.

This network consists of feature extraction and classification networks. The feature extraction network consists of five convolutional layers with rectified linear unit (ReLU) activation and a pooling layer, while the classification network has three fully connected layers.

https://huggingface.co/welcome

https://www.science.org/content/article/breakthrough-2021 - Protein structures for all

AI-powered predictions show proteins finding their shapes

today, the field's central repository, the Protein Data Bank, contains some 185,000 experimentally solved structures. But mapping structures can take years—and cost hundreds of thousands of dollars per protein.

Code for AlphaFold2 and RoseTTAFold is now publicly available, helping other scientists jump into the game. Those programs "developed by Google sister company DeepMind, trains itself on databases of experimentally solved structures".

https://deepmind.com/ - Our teams research and build safe AI systems. We're committed to solving intelligence, to advance science and benefit humanity.

When we started DeepMind in 2010, there was far less interest in the field of AI than there is today. To accelerate the field, we took an interdisciplinary approach, bringing together new ideas and advances in machine learning, neuroscience, engineering, mathematics, simulation and computing infrastructure, along with new ways of organising scientific endeavour.

We achieved early success in computer games, which researchers often use to test AI. One of our programs learned to play 49 different Atari games from scratch, just from seeing the pixels and score on the screen. Our AlphaGo program was also the first to beat a professional Go player, a feat described as a decade ahead of its time.

We joined forces with Google in 2014 to accelerate our work, while continuing to set our own research agenda. Our programs have learned to diagnose eye diseases as effectively as the world's top doctors, to save 30% of the energy used to keep data centres cool, and to predict the complex 3D shapes of proteins - which could one day transform how drugs are invented.

https://deepmind.com/blog/article/FermiNet - FermiNet: Quantum Physics and Chemistry from First Principles

the Fermionic Neural Network or FermiNet is well-suited to modeling the quantum state of large collections of electrons.