Machine Learning as Applied Technology

Lessons learned from productizing machine learning

> Dietmar Millinger twingz development GmbH aiaustria.com

https://www.youtube.com/watch?v=mzZWPcgcRD0

1000

N

LEE SEDOL 00:01:00

.

2

ALPHAGO

Autonomous vehicles

Autonomous vehicles drive the development of many new machine learning technologies.

The most advanced cases are autonomous trucks in the US. This is happening **today** (e.g. UPS).

Waymo has already 20 million miles of experience with self-driving taxis on public roads.



https://www.theverge.com/2019/8/15/20805994/ups-self-driving-trucks-autonomous-delivery-tusimple https://www.youtube.com/watch?v=yjztvddhZmI

Tesla Hydranet architecture

- Neural network architecture for full self-driving (FSD) function in Tesla cars
- Sensor fusion of video streams from eight cameras into one unified 3D vector space
- Based on transformer architecture
- Solves multiple isolated problems in one architecture



3-Dimensional "Vector Space"





GPT-3, DALL-E, ...

Autoregressive **language model** trained on 6 million text articles.

GPT-3's full version has a capacity of **175 billion** machine learning **parameters**. It shows signs of a **deeper language understanding**

DALL-E, a variant of GPT-3 was trained on a combination of **text and images**.

It delivers impressive results for text to image tasks.

Latest image generation model shows very impressive capabilities (**stable diffusion**).

Human: Hey, how are you doing?
AI: I'm good! What would you like to chat about?
Human: Have you done any coding with Python?
AI: Yep! I've done a lot of Python coding as it is my favorite language.
Human: Nice, it's my favorite too! What have you done with it?
AI: Oh, a lot. I have done everything from creating Python modules to creating microwebsites with it.
Human: What web framework do you use?
AI: I used to use Django, but it was too difficult for me. I have found Flask is a better option as it's lighter and simpler.
Human: Wow, I agree 100%. Have you used WebSocket at all?
AI: Yes, I actually have written a chatbot using it!
Human:



https://www.theverge.com/21346343/gpt-3-explainer-openai-examples-errors-agi-potential https://towardsdatascience.com/have-you-seen-this-ai-avocado-chair-b8ee36b8aea https://arxiv.org/pdf/2102.12092.pdf https://www.twilio.com/blog/ultimate-guide-openai-gpt-3-language-model

Google SayCan Robot

Concept for clever robots

- use of autoregressive language model to continue the currently observed status
- collect many different possible continuation stories
- filter best option in relation to abilities (affordances) of the robot





(a) Large Language Models (LLMs)



(b) SayCan

https://arxiv.org/pdf/2204.01691.pdf

Alpha Fold

The **protein folding** problem requires a solution that can estimate the 3D structure of a protein given it's amino acid sequence as input.

Breakthrough by deepmind solution in 2020.

The Alpha Fold 2 model is an end-to-end approach from amino acid sequence input to 3D shape and 3D distances between amino acid residues.

The model contains a transformer core which uses parts of the output also as input to perform evolutional refinement in iterative steps.

Now available as open source and as a protein database with 200 million entries in 2022.





Median Free-Modelling Accuracy

https://www.deepmind.com/research/highlighted-research/alphafold/timeline-of-a-breakthrough https://www.nature.com/articles/s41586-021-03819-2

Transformers

Based on a refinement of the attention idea.

Learns to relate information from different parts of the structured input.

Attention can be understood as a **key value store** where the keys, values and queries are learned from the data in order to fulfill a given supervised task.

Ashish Yaswari Gogle Brain anzawani Gogle Brain anzawani Gogle Brain and Gogle Brain	Ashish Vaswani" Google Brain avaswani@google.com Llion Jones* Google Research llion@google.com	Noam Shazeer* Google Brain	Niki Parm	
Lion Jons* Google Research Diversity of Toronto. Lukasz Kaiser* Google Brain Lukaszkaiser@google.com Mine Nomer* † University of Toronto. Mine Nomer* † University of Toronto. Mine Schweiter Mine Polsenkhin* † Ilis polseukhin@gnail.com Mine Schweiter Mine Schweiter Mine Polseukhin * Mine Schweiter Mine Schweiter Mine Polseukhin * Mine Schweiter Mine Schweiter Mine Polseukhin@gnail.com Mine Schweiter Mine Schweiter Mine Schweiter <th>Llion Jones* Google Research llion@google.com</th> <th>noamegoogre.com</th> <th>Google Rese nikip@google</th> <th>ar* Jakob Uszkoreit arch Google Research e.com usz@google.com</th>	Llion Jones* Google Research llion@google.com	noamegoogre.com	Google Rese nikip@google	ar* Jakob Uszkoreit arch Google Research e.com usz@google.com
Inte Polosukhin ^{* ‡} Itilia polosukhin@gmilt.com		Aidan N. Gome: University of Toro aidan@cs.toronto	z*† onto o.edu luka	Łukasz Kaiser* Google Brain aszkaiser@google.com
Abstract The dominant sequence transduction models are based on complex recurrent of convolutional neural networks that include an encoder and a decoder. The base includes the sequence of the seconder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformé predrainse, We propose a new simple network architecture, the Transformé predrainse of the seconder and decoder through an attention predrain translation mechanisms, dispensing with recurrence and convolutions of the second second second second second second second second predrain translation mechanisms, dispensing with the second second second predrain translation mechanisms, dispensing with the second second second predrain translation mechanisms, dispensing were the section best result. Use one of 41.8 affet instantion the literature. We show that the Transformer generalizes well but best models from the literature. We show that the Transformer generalizes well but standed translation tasks is provided table-of-the transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature. We show that the Transformer generalizes well but set models from the literature show the transformer generalizes well but se		Illia Polosu illia.polosukhi	khin* [‡] n@gmail.com	
The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English- to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfally to English constituency parsing both with large and limited training data.		Abstra	ict	
	The dominant seque convolutional neurar performing models mechanism. We pr based solely on atter entirely. Experime be superior in qualiti- less time to train. It to-German translati ensembles, by over ' our model establishe training for 3.5 day best models from th other tasks by apply large and limited tra	ence transduction mode al networks that include also connect the encor- oropose a new simple n into mechanisms, dispe- nts on two machine try by while being more par Our model achieves 28. to nask, improving ov 2 BLEU. On the WMT 2 es a new single-model st s on eight OPUs, a sma le literature. We show th ying it successfully to E aining data.	els are based on an encoder and leer and decoder etwork architectu sning with recurrs nalation tasks si allelizable and r 4 BLEU on the er the existing b 014 English-to-F Builsh-to-f the-art BL1 all fraction of the at the Transform nglish constituer	complex recurrent or a decoder. The best through an attention are, the Transformer, ence and convolutions how these models to equiring significantly WMT 2014 English- est results, including irench translation task, EU score of 41.8 after training costs of the er generalizes well to acy parsing both with
1 Introduction Recurrent payral activates long chart term memory [13] and gated recurrent [7] payral activ	1 Introduction	long short-term memor	v [13] and gated	recurrent [7] neural network

efficient interence and visualizations. Luxasz and Adan spent countiess long days designing various parts of and implementing tensor2tensor, replacing our catifier codebase, greatly improving results and massively accelerating our research. [†]Work performed while at Google Brain.

[‡]Work performed while at Google Research.

2017

6 Dec

[cs.CL]

arXiv:1706.03762v5

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

Control of tokamak plasma

Plasma shape and configuration control is a key function in nuclear fusion reactors based on the tokamak principle.

A model developed by google DeepMind shows promising results. The model is based on **reinforcement learning** methods and it controls the plasma via control coils.

Training requires a very **detailed simulation model** of the system.

Article Magnetic control of tokamak plasmas through deep reinforcement learning

 https://doi.org/10.1038/s41586-021-04301-9
 Jo

 Received: 14 July 2021
 Tri

 Accepted: 1D December 2021
 Ja

 Published online: 16 February 2022
 Ss

 Donn access
 December 2021

Jonas Degrave¹³, Federico Felici²³⁸³, Jonas Buchl¹³⁴³, Michael Neuner¹⁴, Brendan Tracey¹³⁸⁷, Francesco Carpanese¹²⁴³, Timo Ewalda¹⁴, Roland Hafner¹³, Abbas Abdolmalek¹, Digo de las Casas¹, Craig Donner¹, Leslie Fritz, Cristian Galperti¹, Andrea Huber¹, James Keeling¹, Maria Taimpoukell¹¹, Jackie Kay¹, Antoine Merle¹, Jean-Marc Moret¹, Seb Noury¹, Federico Pesamosca², David Platu¹, Olivier Sauter², Cristian Somanriva³, Stefano Coda³, Basil Duval², Ambrogio Fasol¹¹, Pushmeet Kohli¹¹, Koray Kavukcuoglu¹, Demis Hasasbi¹¹ é Martin Ridemiller¹¹

Check for updates

Nuclear fusion using magnetic confinement, in particular in the tokamak configuration, is a promising path towards sustainable energy. A core challenge is to shape and maintain a high-temperature plasma within the tokamak vessel. This requires high-dimensional, high-frequency, closed-loop control using magnetic actuator coils, further complicated by the diverse requirements across a wide range of plasma configurations. In this work, we introduce a previously undescribed architecture for tokamak magnetic controller design that autonomously learns to command the full set of control coils. This architecture meets control objectives specified at a high level, at the same time satisfying physical and operational constraints. This approach has unprecedented flexibility and generality in problem



414 | Nature | Vol 602 | 17 February 2022

Near future of AI development

- ➤ Multi-modal transformers
- Learning from little data (one shot learner)
- > Explainability, safety and ethical AI, regulation
- ➢ Self structuring systems (AutoML)
- > Neuromorphic computer (silicon Al)
- Quantum-* (machine-learning, training,...)
- > Artificial general intelligence



Anomaly detection in energy flow data

What is given

- connected sensors that deliver residual current measurements (RCM) from industrial power installations
- > data in 1 to 60 second resolution

What is required

- alarm signals for abrupt anomalies in real-time data
- alarm signals for dangerous long term trends

Challenge

- ➤ each installation is different
- many installations exceed normative limits for short times



Table 2. Measured sensor channels.

Name	Unit	Description
DC	mA	Direct current component of residual current
AC	mA	Sum of all alternating current components
50 Hz	mA	Residual current in 50 Hz band
<100 Hz	mA	Residual current in below 100 Hz band
150 Hz	mA	Residual current in 150 Hz band
100 Hz–1 kHz	mA	Residual current in mid frequency band
>1 kHz	mA	Residual current in 1 kHz band
>10 kHz	mA	Residual current in high frequency band

a State Estimation Training Pipeline

The current solution

Autoencoder

 variational autoencoder for control over shape of latent space

Measurement of the reconstruction error

- ➤ estimation of anomaly
- ➤ identification of type of anomaly

Clustering of latent space

 k-Means clustering for estimation of operational states of system



Figure 1. Simplified models of the data processing pipelines: (**a**): Principle of the state estimation training pipeline. (**b**): Principle of the state estimation prediction pipeline.

Autoencoder

Symmetrical neural network with a small latent space vector and identical input and output configuration.

Trained to reproduce input data at output.

Generates a dense vector representation of training data in the latent space.





https://arxiv.org/pdf/2003.05991.pdf

https://hackernoon.com/autoencoders-deep-learning-bits-1-11731e200694

https://towardsdatascience.com/applied-deep-learning-part-3-autoencoders-1c083af4d798

Applications

- ➤ anomaly detection
- ➤ generative models
- reconstruction of input data

Reconstruction error

After training an autoencoder to reconstruct the input x, the construction error can be calculated as:

 $error_k = \sum_{i=1}^m (\mathbf{x}^i - \hat{\mathbf{x}}^i)^2$

The error of reconstruction can be used for:

- anomaly detection: use the reconstruction >error as a signal for anomalies. E.g. high error indicates **unseen** data.
- input reconstruction: use the \succ reconstruction error to find regions with distortions (e.g. noise).



Line plot of reconstruction_error

Data Index

Variational Autoencoder

Variational autoencoders have a latent space vector which consists of **two vectors** of size N. One part is the **mean value** μ of the latent variable and the second part is the **variance** σ of the latent variable.





https://towardsdatascience.com/intuitively-understanding-variational-autoencoders-1bfe67eb5daf

Variational Autoencoder

The **loss function** for the training contains two parts.

- One to reduce the **reconstruction error** of the input values and
- one that forces the distribution of the values in the latent space to be normal distributed (mean 0 and std dev 1).

This results in a tense latent space distribution of the representations of the input values. This makes it possible to sample from the latent space and use the **decoder part for generation**.



Application and outlook

Application in

- industrial installations with high energy consumption or fire risk
- industrial installations with high availability demand
- > photovoltaic systems

Further development

- integration of several anomaly detection models
- clustering of anomalies to filter regular anomalies from irregular ones
- experiment with time-series transformers in order to distinguish between anomalies in time sequence and anomalies in relation between features

Open Access Article

Analysis of Residual Current Flows in Inverter Based Energy Systems Using Machine Learning Approaches

by 😩 Holger Behrends ^{1.*,1,‡} 🖾 📀, 😩 Dietmar Millinger ^{2,‡} 🖾, 🤮 Werner Weihs-Sedivy ^{2,‡} 🖾, 🤮 Anže Javornik ^{3,‡} 🖾, 🥸 Gerold Roolfs ⁴ 🖾 and 😵 Stefan Gelßendörfer ¹ 🖾 💿

¹ German Aerospace Center (DLR), Institute of Networked Energy Systems, 26129 Oldenburg, Germany

² Twingz Development GmbH, 1060 Vienna, Austria

³ Pointar d.o.o., 4220 Škofja Loka, Slovenia

⁴ Doepke Schaltgeräte GmbH, 26506 Norden, Germany

* Author to whom correspondence should be addressed.

- [†] Current address: Carl-von-Ossietzky-Strasse 15, 26129 Oldenburg, Germany.
- [‡] These authors contributed equally to this work.

Academic Editors: Wei-Hsin Chen and Rui Xiong

Energies 2022, 15(2), 582; https://doi.org/10.3390/en15020582

Received: 15 October 2021 / Revised: 20 December 2021 / Accepted: 27 December 2021 / Published: 14 January 2022

(This article belongs to the Collection Artificial Intelligence and Smart Energy)



Abstract

Faults and unintended conditions in grid-connected photovoltaic systems often cause a change of the residual current. This article describes a novel machine learning based approach to detecting anomalies in the residual threshold-based detection methods. For this study, a power-hardware-in-the-loop approach was carried out, in which typical faults have been injected under ideal and realistic operating conditions. The investigation shows that faults in a photovoltaic convertional threshold-based detection methods. For this study, a power-hardware-in-the-loop approach was carried out, in which typical faults have been injected under ideal and realistic operating conditions. The investigation shows that faults in a photovoltaic convertien system cause a unique behaviour of the residual current and fault patterns can be detected and identified by using pattern recognition and variational autoencoder machine learning algorithms. In this context, it was found that the residual current is not only affected by malfunctions of the system, but also by volatile external influences. One of the main challenges here is to separate the regular residual currents caused by the interferences from those caused by faults. Compared to conventional methods, which respond to absolute changes in residual current, the two machine learning models detect faults that do not affect the absolute value of the residual current. View Full-Text

Keywords: renewable energies; photovoltalc; predictive maintenance; reliability; anomaly detection; residual current; machine learning; reconstruction error

V Show Figures





PV-DiStAnS 2.0



Machine Learning in physics

The subjects of the sciences of physics are sometimes way more complex than the subjects of the science of machine learning (personal hypothesis)

So it may be way more complex for a data scientist to enter the center zone than for a physicist to develop from the blue zone to the center zone.

However, a strong interaction between domain expertise and data science knowledge is required for successful relevant machine learning projects.



Why are ML projects different to software projects?

Development of software has become a well understood and managed capability of organizations. Machine learning has not yet reached this status. Many ML projects have to deal with the uncertainty if the expected results can be delivered.



Mindsets I

Team mindset

- > focus on **domain impact** vs data science
- it is very easy to be drawn into very interesting data science and machine learning problems

Project mindset

- > set **realistic goals** for ML projects
- manage customer expectations
- ➤ avoid isolation of ML teams



Mindsets II

Fail-fast principle

- support experimentation with ML
- ➢ error tolerant project setup
- > provide suitable ML infrastructure
- test and monitor extensively
- ➢ implement continuous integration for ML

Focus on end-to-end problem

- focus on complete cycle from data acquisition to integration of results
- > start with a **baseline model**
- implement incremental improvements



Life cycle of ML projects



More things to consider

Data management

- > organize your data
- at least metadata should go into a database

Model management

> prevent use of old or wrong models

Model health monitoring

 detect long term degradation of model quality since data distributions may change over time

Machine learning pipeline

> investment will pay off in iteration cycles



Machine learning generates statistics not causality

ML enforces simple models.

Errors in data may lead to low quality models.

E.g. Husky case from LIME paper:

- training data contained wolves in snowy background and huskies in gras background.
- Network learned to look at the background only



(a) Husky classified as wolf

(b) Explanation

Conclusions

Machine learning and deep learning provide many valuable methods and solutions which can be utilized in real-world, production level applications.

However, adaptations and deviations from normal software project setups have to be considered.

A crucial decision is the composition of the project team as a mixture of domain experts and data science experts.

Unless practical experience has already been gathered with ML in a problem domain, prepare for iterations and even failure of approaches.

Extra care shall be given to data and model management as well as permanent monitoring of model and data health.

Do not put too much trust in ML models since current ML does learn statistics but not necessarily causality.

Questions?

dietmar.millinger@aiaustria.com