


Machine Learning / Deep Learning

Supervised Learning

- “Traditioneller” Ansatz
- Erfordert ein „Ground Truth“
 - Wir von Menschen / Experten erstellt
 - Aufwändig und teuer
- Modell optimiert auf diese Ground Truth
 - Iterative anpassung der Parameter
 - Minimierung des Fehlers
 - → „Training“ oder „Lernen“



Pawan Gupta. A baby step towards applied machine learning for mobile apps
<https://medium.com/@pawankgupta.se/baby-steps-towards-applied-machine-learning-for-mobile-apps-fb8309ae8380>

DEEP LEARNING BRAUCHT DATEN

Lernen anhand Beispiele


- Deep Learning benötigt viele diverse Daten um die relevanten Merkmale eines Objekts zu erkennen.
 - Unterschiedliche Hunderassen, Automarken, Vogelarten, Kleidung, etc.

Overfitting

- „Sieht“ ein Modell nur rote Autos, geht es davon aus, dass Autos rot sind.
- Sind die Trainingsdaten zu wenig divers, kann es nicht auf „ungesehene“ Daten generalisieren („es erkennt keine blauen Autos“)

Probleme

- Bias: Annotator's Bias, Vorurteile, Stereotype, etc. bilden sich im Annotierungs- / Datensampling-prozess ab
- Persönliche Daten: Identifizierende Informationen in Text, Bild, Ton
 - → Daten Anonymisieren, Synthetische Daten
- Ethische Implikationen: Welche Daten dürfen verwendet werden (z.B. Hautfarbe, Religion, sexuelle Orientierung, Konsumverhalten, etc.)
 - → Ethik Richtlinien, gesetzliche Richtlinien



TRANSFORMERS – BERT & GPT


Self-Supervised Learning

- Ground Truth = \$\$\$
 - → Ground Truth
- Self-Supervised Learning
 - Definition von Hilfsaufgaben
 - Z.B. Wörter oder Sätze vorhersagen, richtige Reihenfolge bestimmen, etc.
- **Vorteil**
 - Ground Truth kann automatisch erstellt werden
 - Modelle können auf „unendlich“ vielen Daten trainiert werden
- **Nachteil**
 - No Human in the Loop


Word Prediction (e.g., center, neighbor)

A quick **brown** fox jumps over the lazy dog

Sentence Prediction (e.g., center, neighbor, next)



Auto-regressive Language Modelling



Masked Language Modelling



Sentence Order Prediction

Sentence 1	Sentence 2	Correct order
I completed high school	Then I did my undergrad	yes
Then I did my undergrad	I completed high school	no


Large Language Models (LLM)

Bigger is Better...


- Early Transformer Models
 - BERT, GPT → 100 Millionen Parameter
- Aktuelle LLMs
 - GPT-3 → 170 Milliarden Parameter
- Large Models brauchen Large Data!
 - 2018 BERT → Wikipedia
 - 2020 GPT-3.5 →

Datasets	Quantity (Tokens)	Weight in Training Mix	Epochs elapsed when training for 300 BN tokens
Common Crawl (filtered)	410 BN	60%	0.44
WebText2	19 BN	22%	2.90
Books1	12 BN	8%	1.90
Books2	55 BN	8%	0.43
Wikipedia	3 BN	3%	3.40

GPT-3: All you need to know about the AI language model <https://www.sigmoid.com/blogs/gpt-3-all-you-need-to-know-about-the-ai-language-model/>



The emergence of Large Language Models (LLMs) <https://thelowdown.momentum.asia/the-emergence-of-large-language-models-llms/>



Transformer models: an introduction and catalog—2023 Edition <https://amatriain.net/blog/transformer-models-an-introduction-and-catalog-2d1e9039f376/>

Large Language Models (LLM)

Bigger is Better...

Daten-Hunger

- Beispiel: “The Pile”
 - 825 GB
 - Kombination aus 22 kleineren, qualitativ hochwertigen Datensätzen





Figure 1: Treemap of Pile components by effective size.


Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., ... & Leahy, C. (2020). The pile: An 800gb dataset of diverse text for language modeling. arXiv preprint arXiv:2101.00027.

CLIP - Connecting text and images


1. Contrastive pre-training



2. Create dataset classifier from label text




3. Use for zero-shot prediction



Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... & Sutskever, I. (2021, July). Learning transferable visual models from natural language supervision. In International conference on machine learning (pp. 8748-8763). PMLR.

Stable Diffusion




Stable Diffusion

- Midjourney, Dalle, etc.

Jay Alammar. The Illustrated Stable Diffusion <https://jalammar.github.io/illustrated-stable-diffusion/>

WO GEHT DIE REISE HIN?

„Herkömmliche“ Technik Entwicklung



Artificial Intelligence 2023



Vielen Dank!

ALEXANDER SCHINDLER

Thematic Coordinator Researchfield Datascience
Data Science & Artificial Intelligence
Center for Digital Safety & Security

AIT Austrian Institute of Technology GmbH

Giefinggasse 4 | 1210 Vienna | Austria

T +43 50550-2902 | M +43 664 8251454

alexander.schindler@ait.ac.at | www.ait.ac.at

