



Climate Change AI

Expanded Introduction to Machine Learning

This document contains extended reading materials for different topics in machine learning (ML). It is intended as a list of resources for participants of the Climate Change AI Summer School 2024 who are interested in knowing more about particular topics in machine learning at an introductory level. It is not intended as a self-contained introduction to machine learning.

1. Introduction to AI

This topic covers an overview of Artificial Intelligence (AI), tracing its historical development from early foundations, the AI winters, the resurgence, and the current state of AI. Also includes some of the branches and intersections of AI with other fields and their impact and challenges.

Readings:

- [Video] Introduction to Artificial Intelligence
<https://www.youtube.com/watch?v=16Dir4QqCUg>
- [Video] AI History
https://www.youtube.com/watch?v=z8fEXuH0mu0&list=PLoROMvodv4rOca_Ovz1DvdtWuz8BfSWL2&index=3
- [Video] Artificial Intelligence Today
https://www.youtube.com/watch?v=C0lhR4D5KYc&list=PLoROMvodv4rOca_Ovz1DvdtWuz8BfSWL2&index=4
- [Book] Artificial Intelligence: A Modern Approach <https://aima.cs.berkeley.edu/>

2. Infrastructure for AI

This topic covers the infrastructure useful for developing AI models. It includes the importance of Integrated Development Environments (IDEs) and virtual environments for development, with a special focus on Conda as a management tool. Google Colab is a cloud-based platform that offers free GPU resources, making it an excellent choice for learning and experimenting with simple AI/ML models. Version control with Git is another crucial topic, as it allows for efficient tracking of changes, collaboration among team members, and maintaining a history of project evolution. Additionally, popular deep learning frameworks such as TensorFlow and PyTorch are specifically designed for building DL models effectively.

Readings:

- [Blog] The Python Tutorial <https://docs.python.org/3/tutorial/>
- [Video series] PyTorch Beginner Series
https://www.youtube.com/watch?v=IC0_FrIX-sw&list=PL_IsbAsL_o2CTlGHgMxNrKhZP97BaG9ZN
- [Video series] TensorFlow in Google Colaboratory
https://www.youtube.com/watch?v=inN8seMm7UI&list=PLQY2H8rRoyvyK5aEDAI3wUUqC_F0oEroL&index=2
- [Blog] Github Essentials
<https://docs.github.com/en/get-started/start-your-journey/hello-world>
- Documentations and official tutorials for [Python](#), [PyTorch](#), [Tensorflow](#), [Conda](#), [Github](#) are great resources to refer

3. Neural Networks

This topic delves into the core concepts and mechanics behind neural networks. Neural networks are computational models designed to recognize patterns and predict based on data. This topic can cover exploring the structure and function of neural networks, explaining how neurons, layers, and activation functions work together to process information. The role of gradients and backpropagation in training neural networks, highlighting how these methods optimize model performance. It can also cover designing loss functions and metrics for evaluating models.

Readings:

- [Book] Deep Learning <https://www.deeplearningbook.org/>
- [Video] Introduction to Neural Networks
<https://www.youtube.com/watch?v=MfljxPh6Pys&list=PLoROMvodv4rMiGQp3WXShMGgzqpfVfbU&index=12>
- [Video Series] Neural Networks
https://www.youtube.com/watch?v=aircAruvnKk&list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi




Image Credit: [Colah's blog](#)

4. Paradigms of ML

This section covers the different paradigms of deep learning, exploring the various approaches and techniques that drive modern AI advancements. Supervised learning, one of the most common paradigms, trains models using labeled data. Unsupervised and self-supervised learning emphasize the significance of pre-trained models. Transfer learning highlights how pre-trained models can be adapted to new tasks through fine-tuning. Another area, reinforcement learning, focuses on how agents learn to make decisions in an environment to maximize rewards. Active and online learning explore areas where models continually learn from new data. Semi-supervised learning showcases its efficiency in leveraging both labeled and

unlabeled data. Generative modeling involves generating new data samples from learned distributions.

Readings:

- [Video] Transfer learning
https://www.youtube.com/watch?v=yofjFQddwHE&list=PLkDaE6sCZn6E7jZ9sN_xHwSHOdjUxUW_b&index=19
- [Video] Supervised learning
https://www.youtube.com/watch?v=sca5rQ9x1cA&list=PLkDaE6sCZn6FNC6YRfRQc_FbeQrF8BwGI&index=4
- [Video] Unsupervised learning
https://www.youtube.com/watch?v=gG_wl_uGfIE&list=PLkDaE6sCZn6FNC6YRfRQc_FbeQrF8BwGI&index=6
- [Book chapter] Probabilistic Machine Learning: An Introduction, Chapter 1
<https://probml.github.io/pml-book/book1.html>
- [Blog] Reinforcement Learning: An Introduction With Python Examples by Bex Tuychiev
<https://www.datacamp.com/tutorial/reinforcement-learning-python-introduction>

5. Computer Vision

This topic explores the fundamental techniques and advanced methods that enable machines to interpret and understand visual data. Convolutional Neural Networks (CNNs) learn the spatial hierarchies of features and form the backbone of many computer vision applications. Vision transformers that use self-attention mechanisms for image analysis is another emerging architecture. It also covers the entire image processing pipeline, outlining the steps from raw image acquisition to final analysis. Additionally, delving into commonly performed computer vision tasks such as image segmentation (dividing images into meaningful segments) can be useful for a deeper understanding.

Readings:

- [Book] Computer Vision: Algorithms and Applications
<https://szeliski.org/Book/>

- [Paper] An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale <https://arxiv.org/abs/2010.11929>
- [Video] But what is a convolution? <https://www.youtube.com/watch?v=KuXjwB4LzSA>
- [Blog] Conv Nets <https://colah.github.io/posts/2014-07-Conv-Nets-Modular/>

6. Natural Language Processing

This topic covers the techniques and models that enable computers to understand and generate human language. An introduction to NLP fundamentals include tokenization, stemming, and lemmatization, which are crucial for text preprocessing. Word embeddings play a crucial role in capturing semantic relationships between words. The various NLP models can include from traditional approaches like n-grams to advanced neural architectures such as Transformers and BERT. The key tasks in NLP, such as sentiment analysis, named entity recognition, and machine translation showcase the wide array of applications and challenges involved in building robust language modeling tools.

Readings:

- [Video] What are Transformer Neural Networks? <https://www.youtube.com/watch?v=XSSTuhyAmnI>
- [Video] NLP with Deep Learning <https://www.youtube.com/watch?v=LWMzyfvuehA&list=PLoROMvodv4rMFqRtEuo6SGjY4XbRIVRd4&index=8>
- [Video series] Natural Language Processing (NLP) with Tensorflow <https://www.youtube.com/playlist?list=PLQY2H8rRoyvzDbLUZkbudP-MFQZwNmU4S>
- [Blog] The Illustrated Transformer <https://jalammar.github.io/illustrated-transformer/>
- [Blog] Understanding LSTM Networks <https://colah.github.io/posts/2015-08-Understanding-LSTMs/>

7. Math fundamentals

This topic includes the mathematical foundations essential for machine learning. Linear algebra forms the backbone, providing tools for vector and matrix operations crucial in neural networks. Calculus is important for understanding optimization, enabling the calculation of gradients for backpropagation. Probability and statistics are fundamental for modeling uncertainty and interpreting model predictions. Key concepts in optimization, such as gradient descent, help in training model parameters. These mathematical principles collectively form the foundation upon which deep learning algorithms are built and optimized.

Readings:

- [Book] Introduction to Probability
https://www-sop.inria.fr/members/Giovanni.Neglia/probas/bertsekas_tsitsiklis_probability.pdf
- [Book] Introduction to Linear Algebra
<https://math.mit.edu/~gs/linearalgebra/ila6/indexila6.html>
- [Video series] Essence of linear algebra
https://www.youtube.com/playlist?list=PLZHQObOWTQDPD3MizzM2xVFitgF8hE_ab
- [Video series] Essence of calculus
<https://www.youtube.com/playlist?list=PLZHQObOWTQDMsr9K-rj53DwVRMYO3t5Yr>

8. Interpretability and explainability

In many climate-relevant applications, there is a human operator who looks at the model output and makes the final decision. This may for example be the case where machine learning models are used to inform the operations of power grids. When having such a high responsibility, like making sure the lights stay on, the operator would naturally want to know why the model predicts one way or another, and then decide to intervene or override what the model outputs based on their own experience. Many machine learning models (not all!), however, are not “interpretable” in this way. This means that we do not understand what particular aspect of the input data was important for a particular output. Importantly, this does not mean that we don’t understand the technical side of the model – even the developer of the

model will not be able to explain why it is predicting one thing over another even though they certainly understand what they built.

Unfortunately, neural networks used in deep learning are not interpretable. That is why ML engineers and researchers apply a number of tricks to “explain” what they are doing. Visualization techniques, such as activation maps and feature importance approaches, provide insights into how these models make predictions and identify patterns in data. Interpretability methods, including SHAP values and LIME, help in explaining model decisions, making complex models less opaque. Techniques like attention mechanisms and saliency maps reveal which parts of the input data are most influential in the model's decisions. Together, these tools and methods provide an insight into the inner workings of deep learning models.

Interpretability aims:

- Oversight
 - Regulatory oversight and recourse
 - Real-time settings: Intervening & overriding model outputs
 - Domain-informed model debugging
- Credibility
 - Allowing stakeholders to decide whether to trust
- Scientific discovery
 - Working with and expanding domain knowledge

Readings:

- [Book] Interpretable machine learning (Christoph Molnar)
<https://christophm.github.io/interpretable-ml-book/>
- [Article] Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead (Cynthia Rudin)
<https://www.nature.com/articles/s42256-019-0048-x>
- [Tutorial] PSCC 2024 tutorial: Trustworthy AI for Power System:
<http://www.chatziva.com/pssc2024.html>
- [Paper] Explainable machine learning for public policy: Use cases, gaps, and research directions:
<https://www.cambridge.org/core/journals/data-and-policy/article/explainable-machine-learning-for-public-policy-use-cases-gaps-and-research-directions/B5B66B3C3B16196482984E878D795161>
- [Paper] The Values Encoded in Machine Learning Research:
<https://arxiv.org/abs/2106.15590>

- [Blog] Finding the Words to Say: Hidden State Visualizations for Language Models (Jay Alammar): <https://jalammar.github.io/hidden-states/>
- [Blog] The Illustrated GPT-2 (Visualizing Transformer Language Models) (Jay Alammar): <https://jalammar.github.io/illustrated-gpt2/>
- [Blog] The Illustrated Word2vec (Jay Alammar): <https://jalammar.github.io/illustrated-word2vec/>
- [Blog] Visualizing representations by Colah: <https://colah.github.io/posts/2015-01-Visualizing-Representations/>

9. Causal Inference

Causal inference encompasses methods that are used to determine the cause-and-effect relationship between variables. Unlike correlation, which only indicates that two variables change together, causal inference seeks to establish that a change in one variable causes change in another. The identification of causal relationships is crucial in many fields related to climate change, from energy economics to transportation, where understanding the true impact of interventions, policies, or treatments is essential. For instance, researchers might use causal inference to determine whether a subsidy for solar panels caused a decrease in greenhouse gas emissions in a specific region. The challenge lies in isolating the causal effect from other confounding factors. This requires sophisticated statistical techniques and experimental or quasi-experimental designs.

Readings:

- [Video Lecture] Machine Learning for Healthcare - Causal Inference by David Sontag <https://www.youtube.com/watch?v=gRkUh9Wb-I>
- [Blog] An introduction to Causal inference by Fabian Dablander <https://fabiandablander.com/r/Causal-Inference.html>
- [Lecture series] Introduction to Causal Inference by Maya Petersen & Laura B. Balzer: <https://ctml.berkeley.edu/introduction-causal-inference>

10. Physics-informed machine learning

Physics-informed machine learning aims to integrate the laws of physics into machine learning algorithms. Incorporating known physical principles into neural networks can significantly reduce the amount of data needed for training, improve the accuracy of predictions and help prevent catastrophic failure. Physics-informed machine learning is particularly useful in complex systems where data may be

scarce or expensive to obtain. By embedding physical laws into the learning process, it also ensures that the models adhere to constraints of the real world, leading to more robust and generalizable solutions. This approach is transforming fields such as fluid dynamics, material science, and climate modeling, where traditional machine learning models may struggle to provide reliable results due to the intricate nature of the underlying physical processes, which naive neural networks often fail to fully capture.

Readings:

- [Video Explanation] Scientific Machine Learning: Physics-Informed Neural Networks by Craig Gin <https://www.youtube.com/watch?v=RTPo6KgpvBA>
- [Blog] Introduction to Physics-informed Neural Networks: A hands-on tutorial with PyTorch by Mario Dagrada <https://towardsdatascience.com/solving-differential-equations-with-neural-networks-afdcf7b8bcc4>
- [Book chapter] Introduction to Scientific Machine Learning through Physics-Informed Neural Networks https://book.sciml.ai/notes/03-Introduction_to_Scientific_Machine_Learning_through_Physics-Informed_Neural_Networks/
- [Paper] Kashinath, K., et al (2021). Physics-informed machine learning: case studies for weather and climate modelling: <https://royalsocietypublishing.org/doi/10.1098/rsta.2020.0093>
- [Paper] Aurora: A Foundation Model of the Atmosphere: <https://www.microsoft.com/en-us/research/publication/aurora-a-foundation-model-of-the-atmosphere/>
- [Workshop] Climate Change and Machine Learning: Opportunities, Challenges, and Considerations: <https://icml.cc/virtual/2022/tutorial/18443>