Machine Learning and Climate

Reading Schedule (Tentative)

Template for each session:

Week	Topic	Readings
X	**Session Title |** Climate topic(s). ▷ ML topics that pertain to climate topic(s) above.	○ Climate-related reading ○ ML-related reading ○ (Additional readings, if necessary)

Last updated: November 1, 2023. Subject to change.

Please see the next page.

Week	Topic	Readings
1	**Introduction** | The scale of the problem. ▷ Context and background.	○ Kaack et al. [2022]
2	**Biodiversity** | Forecasting biodiversity using regression. ▷ Predictive vs. causal models.	○ Sirén et al. [2022] ○ Reference: [Pearl et al., 2016, Ch. 1,2]
3	**Climate event attribution** | What causes heatwaves? ▷ Counterfactual analysis using data.	○ Hannart et al. [2016] ○ Reference: [Pearl et al., 2016, Ch. 3]
4	**Time series effects** | Earth system applications. ▷ Granger causality, structural causal discovery.	○ Runge et al. [2019]
5	**Monitoring emissions** | Tracking emissions from photographic imagery. ▷ Image processing, convolutional neural networks.	○ Wang et al. [2020] ○ Vedaldi [2019]
6	**Energy systems** | Power flow optimization ▷ Non-convex optimization with constraints.	○ Donti and Kolter [2021] ○ Donti et al. [2021]
7	**Policy** | Tackling climate change doubt; analyzing text-based patterns ▷ Topic modeling.	○ Boussalis and Coan [2016] ○ Grimmer and Stewart [2013]
8	**Urban planning** | Energy usage prediction, design for large systems. ▷ Gaussian process regression.	○ Kolter and Ferreira [2011] ○ [Stan Development Team, 2021, Ch. 10]
9	**Accelerating science** | Material science discovery, optimizing wind farm layouts. ▷ Bayesian optimization.	○ Hellan et al. [2023] ○ Shahriari et al. [2015]
10	**Manufacturing** | Data-driven manufacturing; optimization during production. ▷ Bayesian optimization.	○ Attia et al. [2020] ○ Shahriari et al. [2015]
11	**Computational material science** | New materials for climate change mitigation applications. ▷ Diffusion models	○ Düreth et al. [2023] ○ Luo [2022]
12	**Aligning ML and Climate** | ML's own carbon footprint; how to quantify and mitigate.	○ Strubell et al. [2020] ○ Lacoste et al. [2019] ○ Henderson et al. [2020] ○ Luccioni and Hernandez-Garcia [2023]
13	**Final presentations** | Projects.	

References

Lynn H Kaack, Priya L Donti, Emma Strubell, George Kamiya, Felix Creutzig, and David Rolnick. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, pages 1–10, 2022.

Alexej PK Sirén, Chris S Sutherland, Ambarish V Karmalkar, Matthew J Duveneck, and Toni Lyn Morelli. Forecasting species distributions: Correlation does not equal causation. *Diversity and Distributions*, 28(4):756–769, 2022.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.

Alexis Hannart, J Pearl, FEL Otto, P Naveau, and M Ghil. Causal counterfactual theory for the attribution of weather and climate-related events. *Bulletin of the American Meteorological Society*, 97(1):99–110, 2016.

Jakob Runge, Sebastian Bathiany, Erik Bollt, Gustau Camps-Valls, Dim Coumou, Ethan Deyle, Clark Glymour, Marlene Kretschmer, Miguel D Mahecha, Jordi Muñoz-Marí, et al. Inferring causation from time series in earth system sciences. *Nature communications*, 10(1):1–13, 2019.

Jingfan Wang, Lyne P Tchapmi, Arvind P Ravikumar, Mike McGuire, Clay S Bell, Daniel Zimmerle, Silvio Savarese, and Adam R Brandt. Machine vision for natural gas methane emissions detection using an infrared camera. *Applied Energy*, 257:113998, 2020.

Andrea Vedaldi. A convolutional neural network primer, 2019.

Priya L Donti and J Zico Kolter. Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, 46:719–747, 2021.

Priya L Donti, David Rolnick, and J Zico Kolter. Dc3: A learning method for optimization with hard constraints. *arXiv preprint arXiv:2104.12225*, 2021.

Constantine Boussalis and Travis G Coan. Text-mining the signals of climate change doubt. *Global Environmental Change*, 36:89–100, 2016.

Justin Grimmer and Brandon M Stewart. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political analysis*, 21(3):267–297, 2013.

J Kolter and Joseph Ferreira. A large-scale study on predicting and contextualizing building energy usage. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, 2011.

Stan Development Team. *Stan User's Guide*, 2021.

Sigrid Passano Hellan, Christopher G Lucas, and Nigel H Goddard. Bayesian optimisation against climate change: Applications and benchmarks. *arXiv preprint arXiv:2306.04343*, 2023.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104 (1):148–175, 2015.

Peter M Attia, Aditya Grover, Norman Jin, Kristen A Severson, Todor M Markov, Yang-Hung Liao, Michael H Chen, Bryan Cheong, Nicholas Perkins, Zi Yang, et al. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578(7795): 397–402, 2020.

Christian Düreth, Paul Seibert, Dennis Rücker, Stephanie Handford, Markus Kästner, and Maik Gude. Conditional diffusion-based microstructure reconstruction. *Materials Today Communications*, 35:105608, 2023.

Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*, 2022.

Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13693–13696, 2020.

Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*, 2019.

Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248):1–43, 2020.

Alexandra Sasha Luccioni and Alex Hernandez-Garcia. Counting carbon: A survey of factors influencing the emissions of machine learning. *arXiv preprint arXiv:2302.08476*, 2023.