Machine Learning and Climate

Spring 2022 | Columbia University

Instructor: Alp Kucukelbir (https://proditus.com/teaching/mlclimate2022)
Course Assistant(s): Nicolas Beltran (nb2838@columbia.edu)
Day and Time: Tuesdays 4:10 – 6:00 p.m.
Location: Zoom + Chandler 401

Last edited: January 17, 2022

1 Schedule (Tentative)

Please see the next page.

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: January 17, 2022. Subject to change.

Week	Topic	Readings
1	**Introduction and prediction |** Methane emission prediction. ▷ Regression, classification.	○ Wang et al. [2020] ○ Gelman et al. [2020]
2	**The power of visualization |** Visualization for planning in forest management. ▷ Clustering, diffusion maps.	○ Tomaselli et al. [2020] ○ Van Der Maaten et al. [2009]
3	**Working with images |** Tracking emissions from satellite imagery. ▷ Image processing, convolutional neural networks.	○ Couture et al. [2020] ○ Vedaldi [2019]
4	**Earth system sciences |** Ocean and weather system modeling. ▷ Linear inverse modeling, hybrid models.	○ Zanna [2012] ○ Bolton and Zanna [2019] ○ Reichstein et al. [2019]
5	**Urban planning |** Energy usage prediction, design for large systems. ▷ Gaussian process regression.	○ Kolter and Ferreira [2011] ○ [Stan Development Team, 2021, Ch. 10]
6	**Energy systems |** Power flow optimization ▷ Non-convex optimization with constraints	○ Donti and Kolter [2021] ○ Donti et al. [2021]
7	**Policy |** How to effect change? Measuring impact of carbon taxes. ▷ Structural causal modeling.	○ Athey [2017] ○ Abrell et al. [2019] ○ Bareinboim and Pearl [2016]
8	**Time series effects |** Earth system applications. ▷ Granger causality, structural causal discovery.	○ Runge et al. [2019]
9	**Geoengineering |** Stratospheric aerosol injection; control for geoengineering. ▷ Reinforcement learning.	○ de Witt and Hornigold [2019] ○ Lei [2021]
10	**Manufacturing |** Data-driven manufacturing; optimization during production. ▷ Bayesian optimization.	○ Attia et al. [2020] ○ Shahriari et al. [2015]
11	**Sustainable building |** Material science for construction; optimizing concrete formulations. ▷ Variational auto encoders.	○ Ge et al. [2019] ○ Doersch [2016]
12	**Aligning ML and Climate |** ML's own carbon footprint; how to quantify and mitigate.	○ Strubell et al. [2020] ○ Lacoste et al. [2019] ○ Schwartz et al. [2020] ○ Henderson et al. [2020]
13	**Final presentations |** Projects.	
14	**Final presentations |** (cont.)	

References

Jiayang Wang, Selvaprabu Nadarajah, Jingfan Wang, and Arvind P Ravikumar. A machine learning approach to methane emissions mitigation in the oil and gas industry. 2020. https://www.climatechange.ai/papers/neurips2020/20/paper.pdf.

Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles C Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák. Bayesian workflow. *arXiv preprint arXiv:2011.01808*, 2020.

Lorenzo Tomaselli, Coty Jen, and Ann B Lee. Wildfire smoke and air quality: How machine learning can guide forest management. *arXiv preprint arXiv:2010.04651*, 2020.

Laurens Van Der Maaten, Eric Postma, and Jaap Van den Herik. Dimensionality reduction: a comparative review. *TiCC TR 2009–005*, 2009.

Heather Couture, Joseph O'Connor, Grace Mitchell, Isabella Söldner-Rembold, Durand D'souza, Krishna Karra, Keto Zhang, Ali Rouzbeh Kargar, Thomas Kassel, Brian Goldman, et al. Towards tracking the emissions of every power plant on the planet. 2020. https://www.climatechange.ai/papers/neurips2020/11/paper.pdf.

Andrea Vedaldi. A convolutional neural network primer, 2019. https://www.robots.ox.ac.uk/~vedaldi/assets/teach/2018/c18-notes.pdf.

Laure Zanna. Forecast skill and predictability of observed atlantic sea surface temperatures. *Journal of Climate*, 25(14):5047–5056, 2012. https://laurezanna.github.io/files/Zanna2012-JClim.pdf.

Thomas Bolton and Laure Zanna. Applications of deep learning to ocean data inference and subgrid parameterization. *Journal of Advances in Modeling Earth Systems*, 11(1):376–399, 2019. https://doi.org/10.1029/2018MS001472.

Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, et al. Deep learning and process understanding for data-driven earth system science. *Nature*, 566(7743):195–204, 2019.

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. https://mc-stan.org/docs/2_28/stan-users-guide/gaussian-processes.html.

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.

Susan Athey. Beyond prediction: Using big data for policy problems. *Science*, 355(6324):483–485, 2017.

Jan Abrell, Mirjam Kosch, and Sebastian Rausch. How effective was the uk carbon tax?-a machine learning approach to policy evaluation. *A Machine Learning Approach to Policy Evaluation (April 15, 2019). CER-ETH–Center of Economic Research at ETH Zurich Working Paper*, 19:317, 2019.

Elias Bareinboim and Judea Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 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.

Christian Schroeder de Witt and Thomas Hornigold. Stratospheric aerosol injection as a deep reinforcement learning problem. *arXiv preprint arXiv:1905.07366*, 2019.

Chen Lei. Deep reinforcement learning. In *Deep Learning and Practice with MindSpore*, pages 217–243. Springer, 2021.

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.

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. https://ieeexplore.ieee.org/iel7/5/7360840/07352306.pdf.

Xiou Ge, Richard T Goodwin, Jeremy R Gregory, Randolph E Kirchain, Joana Maria, and Lav R Varshney. Accelerated discovery of sustainable building materials. *arXiv preprint arXiv:1905.08222*, 2019.

Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.

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. https://ojs.aaai.org/index.php/AAAI/article/download/7123/6977.

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

Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. Green ai. *Communications of the ACM*, 63(12):54–63, 2020. https://dl.acm.org/doi/pdf/10.1145/3381831.

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. https://www.jmlr.org/papers/volume21/20-312/20-312.pdf.