AI-driven Grid Optimization Can Reduce Emissions

Kyri Baker[1,2] **Jackie Kay**[1] **Miha Zgubič**[1] **Laura Toni**[1,3] **Steven Bohez**[1]

Eric Perim[1] **Marc Peter Deisenroth**[1] **Sofia Liguori**[1] **Sephora Madjiheurem**[1]

Sims Witherspoon[1] **Sophie Elster**[1] **Kyle Levin**[1] **Luis Piloto**[1]

[1] Google DeepMind [2] University of Colorado Boulder [3] University College London

Abstract

Power systems are an essential backbone of modern society, and notoriously hard to operate optimally, as supply and demand must be balanced in near real-time. Due to the nonlinear dynamics of these networks, many grid operators in the United States and worldwide use linear approximations to clear markets and optimize generator setpoints, introducing inefficiencies like increased losses and unnecessary excess generation. In this paper, we show that the carbon footprint incurred by training a model to learn AC optimal power flow solutions is drastically offset by the gains in operational efficiency (in terms of wasted energy generation) from using these models to optimize grid operations. In particular, we show that it generally takes on the timescale of minutes for these models to offset their initial training footprint.

1 Introduction

The electric power grid is one of the most complex human-made systems in existence. The equations describing the physics of power flow within the grid are nonlinear, and the ground-truth optimization problem that determines generator setpoints is generally NP-hard (Molzahn and Hiskens, 2019). To help make this problem tractable, many independent system operators (ISOs) have used linear models to clear markets and determine generator dispatch (power plant scheduling and output) (Cain et al., 2012). This solution is later adjusted by a post-processing procedure that solves a series of nonlinear equations (AC power flow) to ensure physical relationships are satisfied. However, the use of simplified models for grid operations can unnecessarily increase operational carbon emissions (Winner et al., 2023), increase the cost to generate electricity (Bichler and Knörr, 2023), and decrease grid reliability (Abedi et al., 2020).

Ideally, the AC Optimal Power Flow (AC OPF) problem, a more accurate representation, would be solved instead. However, this problem cannot be reliably solved within the five-minute window often required for many real-time operations (Nair et al., 2022). To address this issue, deep learning models have been proposed to "learn" AC OPF solutions, shifting the computational burden offline, to the training data generation. These models have shown significant improvements in quickly obtaining accurate AC OPF solutions (Baker, 2022; Donti et al., 2021; Fioretto et al., 2020; Piloto et al., 2024). Despite small optimality gaps and high levels of constraint satisfaction, open questions remain regarding how the footprint from model training compares to the benefit of improved operations.

In this paper, we consider the graph neural network (GNN) CANOS (Piloto et al., 2024), which was trained on grids with up to 10,000 buses and is robust to $N-1$ perturbations. We analyze the energy and carbon required for training and inference and corresponding improvements in grid operations. The results indicate that the carbon footprint from training and validating these models, even on large grid sizes, is negligible compared to the potential benefits of operating grid assets more effectively.

2 AI for Optimal Power Flow

The AC OPF problem seeks to minimize generation cost subject to physical constraints. Due to the nonconvex and computationally challenging nature of this problem, grid operators in the United States often use linear approximations of AC OPF (Cain et al., 2012); namely, the DC OPF.

After DC OPF is solved, an AC Power Flow (PF) is solved to recover a physically feasible solution. The AC power flow is not an optimization problem, but encompasses a series of nonlinear equations describing relationships between complex power and voltage. This process solves for variables that the DC OPF does not model, such as reactive power and voltage magnitudes, but keeps all but one generator (the "slack" generator) fixed to the DC solution.

The use of a linear dispatch model introduces an optimality gap. While this gap has been analyzed previously in terms of impact on prices (Overbye et al., 2004), less attention has been given to the impact on power generation. Excess generated power represents not only a direct monetary cost, but also unnecessary carbon emissions.

2.1 Theoretical Benefit from Improved Grid Optimization

We first analyzed the total active power generation difference in DC OPF plus AC PF versus AC OPF. We considered scenarios across network sizes ranging from 30 buses to 10,000 buses (see the Supplemental Information for more details on the considered networks).

In the considered cases, the overall increase in generation for DC OPF plus AC PF ranged from 0.01% to above 8% for some scenarios, with an average of around 2.75%, as compared to AC OPF. Even a 2% decrease in power generation nationwide is significant. In 2022, it is estimated that 4,070 TWh of electrical energy was consumed in the United States, with 2% of this being equivalent to 81.4 TWh. This is roughly equivalent to the electrical energy use of 7.4 million U.S. homes (U.S. Energy Information Administration (EIA), 2020).

2.2 Sources of Energy Consumption

For the status quo scenario, we define the considered energy consumption to be the following:

- Energy from solving DC OPFs plus AC PFs.
- Energy from the power plants, given the dispatch solutions determined by the above.

For the AI-based grid optimizer, we include the following sources of annual energy consumption:

- Energy from dataset generation (which involves solving AC OPFs).
- Energy from model training, evaluation, and hyperparameter tuning. To be conservative, we assume this is an annual process.
- Energy from performing inference (using CANOS to make a prediction) plus AC PFs.
- Energy from the power plants, given the dispatch decisions from CANOS + AC PF.

3 Experiment Setup

Using the models and evaluation procedures for the 10,000 bus system from Piloto et al. (2024), our goal is to calculate the overall energy consumed throughout the entire data generation, training, and inference pipeline. We measured the energy consumption of multiple training and evaluation jobs, leveraging the NVIDIA Management Library for GPU power estimation. The 10,000 bus system model is used as an upper-bound on energy consumption for all states, as this system size is larger than the vast majority of all states' power grids.

Using the Linux `powerTOP` tool (van de Ven, 2022), we measured the energy consumption to perform a typical DC OPF + AC PF using IPOPT (Wächter and Biegler, 2006). We assume an OPF is run every 5 minutes during real-time operation (105,120 across a year).

Dataset generation involves solving AC OPF problems, which have similar energy use to that of solving DC OPFs + AC PFs. Inference involves elementary algebraic operations and function evaluations, and, across the considered operational period (105,120 inferences per year), this totals to merely tens of kWh annually, or less than a single home's annual energy consumption.

To represent a potential ceiling on the annual energy consumption for model and hyperparameter tuning, we take the energy required to train a CANOS model and increase it by two orders of magnitude to account for hyperparameter tuning and possible retraining throughout the year. Depending on the hardware, batch size, number of training steps, and number of parameters, we found that energy consumption ranged from 26 kWh to 55 kWh.

4 Results

Table 1 shows the estimated annual energy consumption allocated to each end-use for the AI-based case and the status quo case, excluding energy generation, rounded to the nearest kWh. We do not show CO_2 emissions here, as those are location-specific and would depend on the generation mixture of the grid powering the process.

Table 1: Consumption breakdown excluding generation

Energy Use Type	**Annual kWh**
Data Generation (AC OPF)	20–780
Training (incl. tuning)	2,600–5,500
Inference	11–12
Assumed Total for Analysis	6,000
DC OPF + AC PF	36–652

As expected, model training incurs a large energy cost compared to DC OPF, although solving OPFs is still not negligible; for challenging cases, conventional solvers may take many iterations to converge, increasing energy use. For the analysis in this paper, we assume a total of 6 MWh for data generation, training, and inference.

4.1 Carbon Reductions

Using active power prediction errors (0.1–1.5%) consistent with the models in Piloto et al. (2024) (including an AC PF after inference, which introduces an additional 1% modification), and applying this error using the basis of a reduction of total generation of 2%, we estimate hypothetical savings under an AI-based optimizer. While grids are typically not operated state-by-state (with some exceptions), we can more granularly apply state-level data to determine the potential impact on individual states. A more realistic estimate would take into account interstate power exchanges.

The 10,000 bus system model is used as a proxy for all states, as this system size is larger (e.g. would require more energy to train) than the vast majority of all states' power systems. The percentage savings (2%), however, assumed from the analysis performed across systems ranging in size from 30 buses to 10,000 buses (see Supplemental Information), is taken to be one standard deviation below the average percentage generation savings. Across all states, 6 MWh for training and evaluation is equivalent to less than one hour of power generation, and on the scale of minutes for most grids (U.S. Energy Information Administration (EIA), 2023b). Compared to other AI models, the chosen GNN is relatively lightweight to train, as it has a focused task (learning AC OPF solutions). The challenge of training these models comes when we try to satisfy AC power flow constraints while pursuing

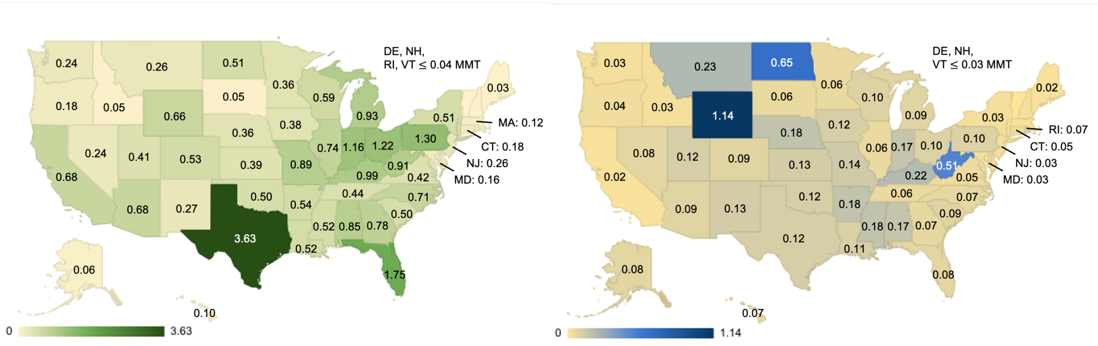

Figure 1: Annual decrease in CO_2 from generation, total (MMT, left) and per capita (MT, right).

optimality of a large-scale nonconvex program. Since this is not achieved with zero error, we must run an AC power flow as a post-processing step, just as we do with DC OPF.

Using the total CO_2 from electric power for each state (U.S. Energy Information Administration (EIA), 2023a) from 2023 with the model accuracy and percent savings assumed above, we plot a country-wide map of the hypothetical carbon reductions per state. Figure 1 (left) shows the potential annual decrease in emissions per state if an AI-based optimizer is used to determine generator dispatch. Numbers are overlaid in the image to indicate the actual decrease in million metric tons (MMT) of CO_2. Despite California being the most populous state, the largest decrease in emissions by far is seen in Texas, with its large amount of in-state generation and extreme levels of energy consumption. California has a milder climate and imports around a quarter of their power (U.S. Energy Information Administration, 2025), reducing their relative benefit from grid efficiency increases.

We also plot the decrease in carbon emissions per capita in Figure 1 (right). Wyoming, which has abundant generation from coal plants and the lowest population of any state has the highest potential reduction. North Dakota and West Virginia are also states with high levels of fossil fuel generation and relatively low populations (fourth and twelfth, respectively), making these states strong contenders for where AI-based grid optimization can have the most impact per capita. While some of the least populous states (Wyoming, North Dakota) benefit from a large per capita reduction, others do not (Alaska, Maine). Some of these differences are due to varying climates and levels of electrification per state. For example, Maine and Alaska primarily use fossil-based sources as primary heating fuels, whereas electricity is often predominant in the southern states U.S. Energy Information Administration (EIA) (2020). Thus, states with higher levels of electrification often benefit more.

4.2 Impact

Using these levels of emission reductions, we can consider how these potential reductions relate to more tangible quantities. If the U.S. converted to AI-based grid optimization, for example, this would be equivalent to 27.98 MMT of CO_2 per year, or: Removing 6.5 million gas-powered mid-size passenger vehicles from the road for a year; 95% of the annual CO_2 emissions of Denmark; 50,000 roundtrip flights from San Francisco to New York; the production of one hamburger for every person on the planet; or the annual CO_2 emissions of 26 gas-fired power plants.

5 Conclusion and Future Work

In this paper, we demonstrated that AI-based grid optimization, despite model training costs, could have the potential to reduce operational grid emissions by improving the quality of dispatch solutions, lowering losses, and reducing the computational power needed to solve conventional optimization problems. Future work could include ISO-specific modeling, time-varying emissions signals, and explicit consideration of slack generator(s) type(s) rather than average emissions factors.

Supplemental Information

DC vs AC generation gap

For ten power networks of varying sizes, we randomly generated 100 load profiles using the perturbation technique described in Fioretto et al. (2020); Lovett et al. (2024). We then calculated the average percent gap in overall active power generation between DC OPF (plus an AC Power Flow) and AC OPF. The results are shown in Figure 2, with the average gap across all considered networks resulting in 2.75%. No clear trend between network size and excess generation from DC OPF was observed, as this is likely topology and line parameter dependent (the generation gap between DC and AC OPF tends to increase as system loading increases, however Baker (2021)).

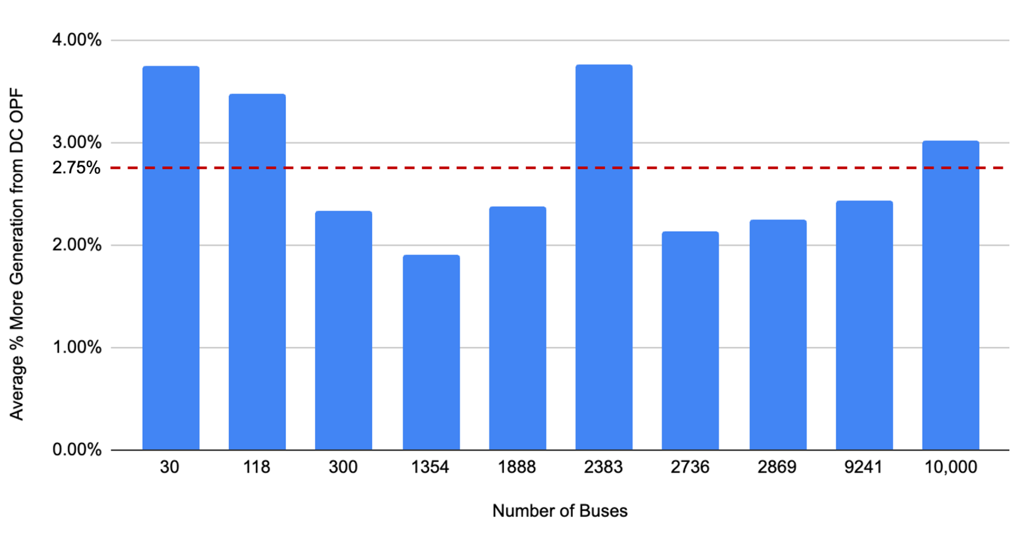

Figure 2: Percent increase in generation from suboptimal grid operations (DC OPF + AC PF) across 100 trials per network, with a mean across all trials of 2.75% and a standard deviation of 0.70%.

References

A. Abedi, L. Gaudard, and F. Romerio. Power flow-based approaches to assess vulnerability, reliability, and contingency of the power systems: The benefits and limitations. *Reliability Engineering & System Safety*, 201:106961, 2020. ISSN 0951-8320. doi: https://doi.org/10.1016/j.ress.2020.106961. URL `https://www.sciencedirect.com/science/article/pii/S0951832019304454`.

K. Baker. Solutions of dc opf are never ac feasible. In *Proceedings of the Twelfth ACM International Conference on Future Energy Systems*, e-Energy '21, page 264–268, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383332. doi: 10.1145/3447555.3464875. URL `https://doi.org/10.1145/3447555.3464875`.

K. Baker. Emulating ac opf solvers for obtaining sub-second feasible, near-optimal solutions, 2022. URL `https://arxiv.org/abs/2012.10031`.

M. Bichler and J. Knörr. Getting prices right on electricity spot markets: On the economic impact of advanced power flow models. *Energy Economics*, 126:106968, 2023. ISSN 0140-9883. doi: https://doi.org/10.1016/j.eneco.2023.106968. URL `https://www.sciencedirect.com/science/article/pii/S0140988323004668`.

M. Cain, R. O'Neill, and A. Castillo. History of optimal power flow and formulations. In *Federal Energy Regulatory Commission (FERC)*, pages 1–36, 2012.

P. L. Donti, D. Rolnick, and J. Z. Kolter. Dc3: A learning method for optimization with hard constraints, 2021. URL `https://arxiv.org/abs/2104.12225`.

F. Fioretto, P. V. Hentenryck, T. W. Mak, C. Tran, F. Baldo, and M. Lombardi. Lagrangian duality for constrained deep learning, 2020. URL `https://arxiv.org/abs/2001.09394`.

S. Lovett, M. Zgubic, S. Liguori, S. Madjiheurem, H. Tomlinson, S. Elster, C. Apps, S. Witherspoon, and L. Piloto. Opfdata: Large-scale datasets for ac optimal power flow with topological perturbations, 2024. URL `https://arxiv.org/abs/2406.07234`.

D. K. Molzahn and I. A. Hiskens. A survey of relaxations and approximations of the power flow equations. *Foundations and Trends® in Electric Energy Systems*, 4(1-2):1–221, 2019. ISSN 2332-6557. doi: 10.1561/3100000012. URL `http://dx.doi.org/10.1561/3100000012`.

A. S. Nair, S. Abhyankar, S. Peles, and P. Ranganathan. Computational and numerical analysis of ac optimal power flow formulations on large-scale power grids. *Electric Power Systems Research*, 202:107594, 2022. ISSN 0378-7796. doi: https://doi.org/10.1016/j.epsr.2021.107594. URL `https://www.sciencedirect.com/science/article/pii/S0378779621005757`.

T. Overbye, X. Cheng, and Y. Sun. A comparison of the ac and dc power flow models for lmp calculations. In *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*, pages 9 pp.–, 2004. doi: 10.1109/HICSS.2004.1265164.

L. Piloto, S. Liguori, S. Madjiheurem, M. Zgubic, S. Lovett, H. Tomlinson, S. Elster, C. Apps, and S. Witherspoon. Canos: A fast and scalable neural ac-opf solver robust to n-1 perturbations, 2024. URL `https://arxiv.org/abs/2403.17660`.

U.S. Energy Information Administration. California state energy profile, 2025. URL `https://www.eia.gov/state/print.php?sid=CA`.

U.S. Energy Information Administration (EIA). Residential energy consumption survey (recs), 2020. URL `https://www.eia.gov/consumption/residential/data/2020/`.

U.S. Energy Information Administration (EIA). Table co2.6. electric power sector co2 emissions estimates from energy consumption, 2023, 2023a. URL `https://www.eia.gov/state/seds/sep_sum/html/pdf/sum_co2_eu.pdf`.

U.S. Energy Information Administration (EIA). Electric power annual 2023, 2023b. URL `https://www.eia.gov/electricity/annual/pdf/epa.pdf`.

A. van de Ven. Powertop. `https://https://github.com/fenrus75/powertop`, 2022.

A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.

C. Winner, J. Garland, C. Crozier, and K. Baker. Carbon emissions resulting from different power flow models for dispatch. In *2023 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5, 2023. doi: 10.1109/PESGM52003.2023.10252655.