
A Learning-boosted Quasi-Newton Method for AC Optimal Power Flow

Kyri Baker

University of Colorado Boulder
Renewable and Sustainable Energy Institute
Boulder, CO, USA
kyri.baker@colorado.edu

Abstract

Power grid operators typically solve large-scale, nonconvex optimal power flow (OPF) problems throughout the day to determine optimal setpoints for generators while adhering to physical constraints. Despite being at the heart of many OPF solvers, Newton-Raphson can be slow and numerically unstable. To reduce the computational burden associated with calculating the full Jacobian and its inverse, many Quasi-Newton methods attempt to find a solution to the optimality conditions by leveraging an approximate Jacobian matrix. In this paper, a Quasi-Newton method based on machine learning is presented which performs iterative updates for candidate optimal solutions without having to calculate a Jacobian or approximate Jacobian matrix. The proposed learning-based algorithm utilizes a deep neural network with feedback. With proper choice of weights and activation functions, the model becomes a contraction mapping and convergence can be guaranteed. Results shown for networks up to 1,354 buses indicate the proposed method is capable of finding approximate solutions to AC OPF very quickly.

1 Introduction

AC optimal power flow (OPF) is a canonical power systems operation problem that is at the heart of optimizing large-scale power networks. Newton-based methods are at the core of many AC OPF solvers used in grids today. However, traditional Newton-Raphson can suffer from slow and numerically unstable Jacobian matrix inversions. To reduce the computational burden associated with calculating the full Jacobian and its inverse, many Quasi-Newton methods attempt to find a solution to the optimality conditions by leveraging an approximate Jacobian matrix [1, 2]. For example, the full Jacobian can be replaced with a function of the gradient of the objective, or the chord method can be used, which fixes the Jacobian to a constant value from the first iteration. Various ways to construct the Jacobian using data-driven techniques have been explored for power flow in [3, 4] with promising results. Here, we avoid calculating a Jacobian by replacing the Newton-Raphson step with a purely data-driven machine learning (ML) model that learns subsequent iterations. The ML model is trained on Newton-Raphson iterations and learns how to imitate the Newton-Raphson algorithm without having to construct Jacobian matrices or calculate matrix inverses.

Learning for OPF is a rapidly growing area due to the immense power that ML and deep learning in particular can provide for representing extremely complex variable relationships and performing inference (making predictions) extremely quickly. For example, [5, 6] attempt to learn AC OPF solutions; [7] develops a method to learn *feasible* AC OPF solutions; ML was used to reduce AC OPF complexity in [8, 9], and warm-start points for AC OPF were obtained using ML in [10, 11], to name a few applications (a more in-depth overview can be found in [12]). Instead of learning the AC OPF solution directly, this paper aims to develop an iterative model that learns Newton-Raphson descent

directions depending on the current candidate solution. This technique aims to learn a representation such as ‘‘Given my current step, what direction should I move to decrease the objective function?’’ rather than ‘‘Given an initial guess, what is the exact optimal solution of the nonconvex AC OPF?’’ A similar technique was explored in [13] where a reinforcement learning agent was trained to solve general unconstrained optimization problems with great success.

The model is comprised of a fully connected three-layer neural network F_R with feedback, where input \mathbf{x}^k is the candidate optimal solution vector at iteration k . Reminiscent of a simple recurrent neural network, the model iteratively uses feedback from the output layer as inputs until convergence ($\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \leq \epsilon$). The model thus bypasses any construction of a Jacobian matrix or associated inverse. Simulation results demonstrate that although the solution is approximate, a solution can be obtained orders of magnitudes faster than traditional Newton-Raphson, and numerical issues due to nearly-singular Jacobians are avoided.

2 Learning-boosted Quasi-Newton

A general nonconvex optimization problem with n -dimensional optimization variable vector \mathbf{x} , cost function $f(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$, M equality constraints $g_i(x) = 0$, $g_i(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$, and P inequality constraints $h_j(x) \leq 0$, $h_j(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ can be written as

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad (1a)$$

$$\text{s.t.} : g_i(\mathbf{x}) = 0, \quad i = 1, \dots, M \quad (1b)$$

$$h_j(\mathbf{x}) \leq 0, \quad j = 1, \dots, P \quad (1c)$$

Newton-Raphson (sometimes just called Newton’s Method) utilizes first and second-order derivatives of the Karush Kuhn Tucker (KKT) optimality conditions of (1) to find a stationary point of the KKT conditions. For AC OPF, this results in finding a local (or maybe global) minimum of the nonconvex problem (1). The Newton-step, that is, the iterative update of candidate optimal solution \mathbf{x}^{k+1} at iteration k can be written as

$$\mathbf{x}^{k+1} = \mathbf{x}^k - \alpha J^{-1}(\mathbf{x}^k) d(\mathbf{x}^k), \quad (2)$$

where $d(\mathbf{x}^k)$ is a vector of KKT conditions evaluated at the current candidate solution \mathbf{x}^k , $J^{-1}(\mathbf{x}^k)$ is the Jacobian matrix of the KKT conditions evaluated at \mathbf{x}^k , and α is an optional step-size parameter where $0 < \alpha \leq 1$. Equation (2) is repeated until convergence, which is typically when iterations cease to change significantly (e.g., $\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \leq \epsilon$) or when the KKT conditions are nearly satisfied (e.g. $d(\mathbf{x}^k) \leq \epsilon$) for some small ϵ , for example.

2.1 Quasi-Newton Methods

The term ‘‘Quasi-Newton’’ simply refers to methods where the Jacobian matrix in (2) or its inverse is approximated. These methods are typically used for two reasons: First, calculating the Jacobian and its inverse is expensive and time-consuming, which may not be appropriate for large problems or fast-timescale optimization. Second, the Jacobian can be singular or close to singular at the optimal solution [14, 15], making the inverse challenging. Thus, Quasi-Newton methods replace J^{-1} with an approximate Jacobian \tilde{J}^{-1} or sometimes with an approximate Jacobian inverse.

The learning-boosted method replaces $\alpha^k J^{-1}(\mathbf{x}^k) d(\mathbf{x}^k)$ in (2) with a fully connected neural network (NN) model $F_R(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that takes in \mathbf{x}^k as an input and provides \mathbf{x}^{k+1} as an output; e.g.

$$\mathbf{x}^{k+1} = F_R(\mathbf{x}^k). \quad (3)$$

A fully connected three-layer NN is used here. The variable vector $\mathbf{x}^k = [\mathbf{v}^k, \boldsymbol{\theta}^k, \mathbf{P}_g^k, \mathbf{Q}_g^k]^T$, where \mathbf{v}^k contains the complex voltage magnitudes at each bus, $\boldsymbol{\theta}^k$ contains the complex voltage angles, and \mathbf{P}_g and \mathbf{Q}_g are the real and reactive power outputs at each generator, respectively.

2.2 Network architecture

Choosing the number of nodes in the hidden layer can be performed by using a popular heuristic [16]: $N_h \approx \sqrt{N_i \cdot N_o} = \sqrt{(2N_L + n) \cdot n}$, where N_i is the number of nodes in the input layer and N_o is the number of nodes in the output layer. For the considered problem, the number of inputs is the number of optimization variables n and the number of loads $2N_L$ (N_L real demands and N_L reactive demands). Note that even though the loads are inputs to the NN, they are fixed and do not change as k changes. These numbers were later refined through trial-and-error. A rectified linear unit (ReLU) was used as the activation function on the input layer; a hyperbolic tangent (tanh) activation function was used in the hidden layer, and a thresholded linear function was used on the output.

3 Convergence Analysis

Quasi-Newton iterations are fixed-point iterations. That is, they can be written in the form $\mathbf{x}^{k+1} = F(\mathbf{x}^k)$, where $F(\mathbf{x}^k)$ is the right-hand side of (2). Unfortunately, when solving AC OPF it is generally difficult to show $\mathbf{x}^{k+1} = \mathbf{x}^k$ for any k . However, we can design our NN to guarantee convergence.

In [17], conditions on the activation functions and bounds on the weights are given to ensure convergence of a recurrent neural network (RNN). The proposed NN architecture is a basic version of an RNN, as the output of the last layer becomes an input of the input layer. Consider the mapping F_R in (3). In order to show that the proposed NN converges to a unique fixed point for any initial x^0 and that F_R is a contraction, it must first be true that each activation function f_i must be bounded, continuous, differentiable, real-valued, and have bounded derivatives. In the proposed NN, thresholds on both the input (ReLU) and output (Linear) layers were placed to bound their output. The ReLU function on input x , $\max\{0, x\}$, is technically not differentiable at $x = 0$. Thus, a more formal analysis would have to be performed to determine if exactly the same convergence results apply in every situation with a ReLU, although they were achieved in simulation. Similar but slightly worse results were achieved with a tanh function on the input layer, which does satisfy these requirements.

From Theorem 1 and 2 in [17] we have a bound on each weight w_{ij} connecting nodes i and j that must hold for the network to converge to a unique fixed point for any given x^0 :

$$|w_{ij}| < c^* < \frac{1}{N_n \cdot f'_{max}}, \quad (4)$$

where $'$ denotes a first-order derivative, N_n is the total number of network nodes, f_l is activation function $l = 1, \dots, N_n$, and

$$f_{max} = \max_{y \in \mathbb{R}, l=1, \dots, N_n} (|f'_l(y)|). \quad (5)$$

Then F_R is a contraction with $0 < c < 1$, and $c = N_n c^* |f'_{max}|$, and w_{ij} can be designed by (4).

4 Network Architecture

Four networks were considered: The IEEE 30 bus, IEEE 300 bus, PG-lib 500-bus [18], and 1,354-bus PEGASE networks [19]. The number of loads, lines, generators, and total (real) generation capacity for each network is shown in Table 1. Line flow constraints were neglected (although some networks did not have them to begin with). The data generation, training and testing of the network, and simulations were performed on a 2017 MacBook Pro laptop with 16 GB of memory. Keras with the Tensorflow backend was used to train the neural network using the Adam optimizer. The chosen number of nodes in the hidden layer and training dataset sizes are shown in Table 2.

4.1 Dataset generation

The MATPOWER Interior Point Solver (MIPS) [20] was used to generate the data and was used as the baseline for comparison with the NN model. The termination tolerance of the MIPS solver was set to 10^{-9} for data generation and 10^{-4} for testing. The tolerance of the learning-boosted solver was

Table 1: Considered network parameters

Case	# of Loads	# of Lines	# of Gens	Total Real Gen. Capacity
30-bus	20	41	6	335 MW
300-bus	191	411	69	32.68 GW
500-bus	200	597	56	12.19 GW
1,354-bus	621	1991	260	128.74 GW

Table 2: Number of nodes and training samples

Case	Input Nodes	Output Nodes	Hidden Nodes	Training Samples
30-bus	112	72	100	72,111
300-bus	1,120	738	800	91,432
500-bus	1,512	1,112	2,300	111,674
1,354-bus	4,470	3,228	6,000	126,724

set to 10^{-4} , where convergence is reached when $\|\mathbf{x}^{k+1} - \mathbf{x}^k\| \leq \epsilon$. A smaller tolerance was used for data generation to promote smoother convergence and “basins of attraction” within the ML model. For a fair comparison, the same convergence criteria was used for the NN model and MIPS during testing. 1,000 different loading scenarios were randomly generated at each load bus from a uniform distribution of $\pm 40\%$ around the given base loading scenario in MATPOWER. In some cases, the generated load profile resulted in an infeasible solution; these samples were excluded.

Table 2 shows the number of training samples ($\mathbf{x}^k, \mathbf{x}^{k+1}$) generated for each scenario. Note that a different number of iterations to convergence is encountered for each scenario and network. It is recognized that generating a diverse and representative dataset is an important and essential thrust of research within learning-based OPF methods. This is an important direction of future work.

5 Simulation Results

5.1 Prediction results

1,000 testing scenarios were generated for each network. Each simulation was initialized with a flat start ($\mathbf{v}^0 = \mathbf{P}_g^0 = \mathbf{Q}_g^0 = 1$ and $\boldsymbol{\theta}^0 = 0$). 1,000 test scenarios were generated using the same methodology in 4.1. The mean absolute error (MAE) for voltage magnitudes (in pu) and real power generation (in MW) is shown in Table 3. The fourth column in the table shows the mean absolute percentage error (MAPE) for the OPF objective function across all scenarios. Interestingly, the 300-bus case had a much worse MAE than other cases. This could be because many solutions in the training dataset had nearly-singular Jacobians near the optimal solution, and thus iterations in some areas were less well-defined and harder for the NN to learn.

Fig. 1 shows 200 test scenarios for the 30-bus system. The NN model does a good job of approximating the actual optimal generation values, given as black dashed lines, for each generator.

5.2 Computation time

The mean, worst-case (max), and variance of time to convergence for the MIPS solver and the proposed NN is shown for all networks across all test scenarios in Table 4.

Table 3: Mean absolute error across testing dataset

Case	MAE: Voltage Magnitude (pu)	MAE: Active Power (MW)	MAPE: Cost (%)
30-bus	0.004 pu	0.64 MW	0.29%
300-bus	0.009 pu	10.47 MW	0.65%
500-bus	0.099 pu	0.62 MW	0.66%
1,354-bus	0.019 pu	7.55 MW	1.16%

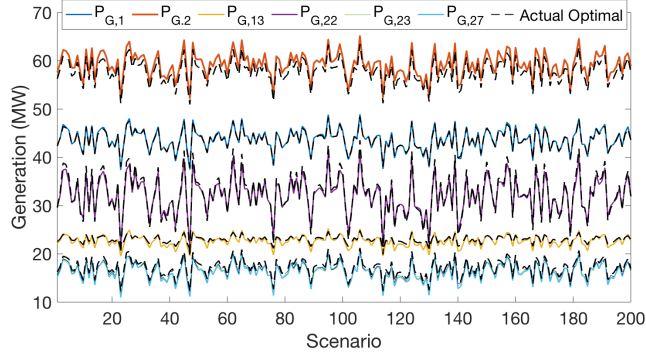


Figure 1: Predicted generation values for 200 test scenarios (colors) and actual optimal values (black dashed line) for the IEEE 30-bus system.

Table 4: Time to convergence for each network.

Case	Mean Time (s)	Max Time (s)	Variance in Time (s)	Mean Speedup
30-bus MIPS	0.04	0.41	4.52e-04	
30-bus NN	0.06	0.42	3.53e-04	-0.66x
300-bus MIPS	1.09	10.87	2.09	
300-bus NN	0.03	0.42	0.001	36.3x
500-bus MIPS	1.46	2.96	0.49	
500-bus NN	0.08	0.45	3.25e-4	18.3x
1,354-bus MIPS	7.64	19.89	15.55	
1,354-bus NN	0.34	0.69	0.0016	22.5x

There is almost no benefit for using a data-driven approach for smaller networks, as these can already be solved in real-time using a solver. However, for larger networks, the learning-boosted Quasi-Newton method obtains solutions extremely quickly compared to the MIPS solver. In addition, the learning-boosted approach has lower variance in computational time, making it more reliable for providing solutions on regular intervals. Interestingly, the 300-bus system takes MIPS longer to solve for some OPF scenarios than the 500-bus. In these cases, the Jacobian was close to singular.

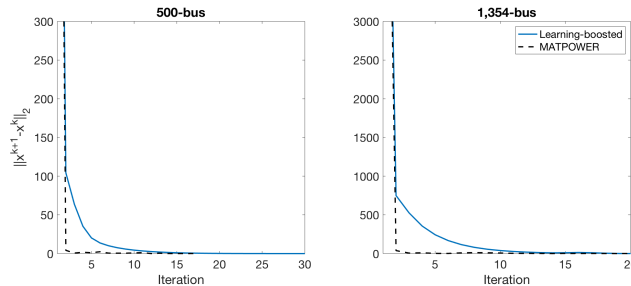


Figure 2: Norm of $\mathbf{x}^{k+1} - \mathbf{x}^k$ for two scenarios in the 500-bus (left) and 1,354-bus (right) systems.

Fig. 2 shows the norm between two iterations, $\|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2$, for two typical test scenarios from the 500 and 1,354-bus systems. While MIPS almost always results in fewer iterations to convergence (because it is using the exact Jacobian), each iteration takes longer to perform. Thus, although each iteration in the learning-boosted method is inexact, each iteration is very fast. In addition, convergence of the NN is guaranteed if the conditions are satisfied in Section 3.

5.3 Assessing feasibility

Variable bounds are ensured by thresholding the output of the output layer on the NN. However, while convergence of the NN is guaranteed, there is no guarantee of convergence to an AC-feasible

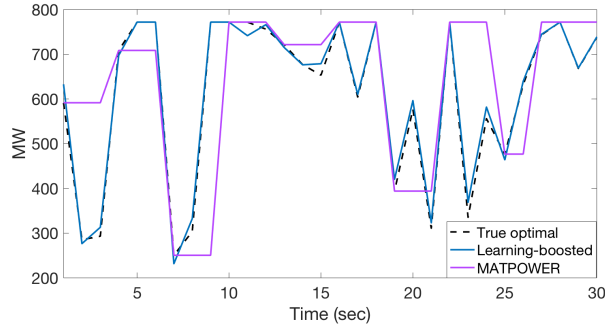


Figure 3: Slack bus generation in the 500-bus system. The learning-boosted solution, although an approximation, is able to track faster changes in demand more accurately than MIPS.

solution (just as with DC OPF). In some cases, like with the 30-bus system, the NN actually outputs solutions with smaller mean constraint violations than MIPS. This is likely due to the fact that the training data was generated for a convergence tolerance of $\epsilon = 10^{-9}$ but during testing, the algorithm is terminated at $\epsilon = 10^{-4}$. Despite the prediction error for the 1,354-bus being relatively low (with a mean optimality gap of 1.16% across the testing set), high errors in the satisfaction of the power flow equations indicate a possible need for an increased training dataset size. Table 5 shows the mean constraint violation for the AC power flow constraints across all networks and test scenarios.

Table 5: Mean constraint violation across all networks

Case	30-bus	300-bus	500-bus	1,354-bus
Mean Constraint Violation (pu)	0.05	0.43	0.32	9.95

5.4 Tracking optimal solutions

Assuming measurements of load at each bus are made available on a one-second basis (which is a reasonable assumption; in fact, SCADA systems should provide measurements with latency of less than one second [21]), the benefit of a real-time optimization approach can be further assessed. Typically, real-time adjustments to generators are done via automatic generation control (AGC); a rule-based, suboptimal affine control policy. Here, a simple experiment was performed where both MATPOWER and the learning-boosted algorithm received load updates every one second. Figure 3 shows the ability of the learning-boosted approach to track optimal generation setpoints with higher accuracy than MATPOWER, despite being an approximation, due to the extremely fast time to obtain new setpoints. This figure illustrates thirty-seconds of the slack bus output for the 500-bus system, which is the generator whose output varies the most in this test case.

As MATPOWER took nearly three seconds to obtain each solution, setpoints were not updated on a one-second scale. Alternatively, the learning-boosted approach provides reasonably accurate solutions in less than a second and keeps up with one-second fluctuations in demand. In reality, generators would adjust their real-time outputs in accordance with the aforementioned AGC heuristics; for illustrative purposes, we show only OPF solutions here. With computationally light methodologies for calculating approximate OPF solutions, optimal solutions could potentially be provided in real-time.

6 Conclusion

A data-driven model for approximating Newton-Raphson for AC OPF was implemented here. Results show that for smaller networks, not much benefit, other than avoiding singular Jacobian matrices and taking inverses of ill-conditioned matrices, can be found. However, for larger networks, the NN model can provide approximate solutions extremely quickly and at very regular intervals compared to a state-of-the-art OPF solver. The optimality gap was small for all considered cases; however, the feasibility gap for the 1,354-bus system was large. This suggests a need for more training data or an increase in model complexity, as the proposed model only contained one hidden layer.

References

- [1] C. Broyden, “Quasi-newton methods and their application to function minimization,” *Mathematics of Computation*, vol. 21, no. 9, pp. 368 – 381, Jul. 1967.
- [2] P. Hennig and M. Kiefel, “Quasi-newton methods: A new direction,” *Journal of Machine Learning Res.*, vol. 14, pp. 843 – 865, Mar. 2013.
- [3] Y. C. Chen, J. Wang, A. D. Domínguez-García, and P. W. Sauer, “Measurement-based estimation of the power flow Jacobian matrix,” *IEEE Transactions on Smart Grid*, vol. 7, no. 5, pp. 2507–2515, 2016.
- [4] X. He, L. Chu, R. C. Qiu, Q. Ai, and W. Huang, “Data-driven estimation of the power flow jacobian matrix in high dimensional space,” Feb 2019, [Online] Available at: <https://arxiv.org/abs/1902.06211>.
- [5] M. Chatzos, F. Fioretto, T. W. K. Mak, and P. V. Hentenryck, “High-fidelity machine learning approximations of large-scale optimal power flow,” *arXiv preprint arXiv:2006.16356*, 2020.
- [6] D. Owerko, F. Gama, and A. Ribeiro, “Optimal power flow using graph neural networks,” Oct. 2019, [Online] Available: <https://arxiv.org/abs/1910.09658v1>.
- [7] A. Zamzam and K. Baker, “Learning optimal solutions for extremely fast AC optimal power flow,” September 2019, [Online] Available at: <https://arxiv.org/abs/2861719>.
- [8] F. Hasan, A. Kargarian, and J. Mohammadi, “Hybrid learning aided inactive constraints filtering algorithm to enhance AC OPF solution time,” Jun. 2020, [Online] Available at: <https://arxiv.org/abs/2006.01336>.
- [9] K. Baker and A. Bernstein, “Joint chance constraints in AC optimal power flow: Improving bounds through learning,” *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6376 – 6385, Mar. 2019.
- [10] K. Baker, “Learning warm-start points for AC optimal power flow,” in *IEEE Machine Learning for Signal Proc. Conf. (MLSP)*, Oct. 2019.
- [11] F. Diehl, “Warm-starting AC optimal power flow with graph neural networks,” in *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*, Dec. 2019.
- [12] F. Hasan, A. Kargarian, and A. Mohammadi, “A survey on applications of machine learning for optimal power flow,” in *2020 IEEE Texas Power and Energy Conference (TPEC)*, 2020.
- [13] K. Li and J. Malik, “Learning to optimize,” *arXiv preprint arXiv:1606.01885*, Jun. 2016.
- [14] Y. P. Zhang, W. Huang, Z. Q. Liu, J. Y. Yang, X. L. Cai, and J. H. Zhang, “Research on the relationship of the singular point for load flow jacobian matrix and the critical point of voltage collapse,” in *IEEE Power Engineering Society General Meeting (PESGM)*, Jun. 2005.
- [15] K. Baker, D. Zhu, G. Hug, and X. Li, “Jacobian singularities in optimal power flow problems caused by intertemporal constraints,” in *2013 North American Power Symposium (NAPS)*, 2013.
- [16] T. Masters, *Practical neural network recipes in C++*. Burlington, MA: Morgan Kaufmann, 1993.
- [17] J. E. Steck, “Convergence of recurrent networks as contraction mappings,” in *Intl. Joint Conference on Neural Networks (IJCNN)*, Jun. 1992.
- [18] S. Babaeinejad-sarookolae *et al.*, “The power grid library for benchmarking AC optimal power flow algorithms,” Aug. 2019, [Online] Available: <https://arxiv.org/abs/1908.02788>.
- [19] S. Fliscounakis, P. Panciatici, F. Capitanescu, and L. Wehenkel, “Contingency ranking with respect to overloads in very large power systems taking into account uncertainty, preventive, and corrective actions,” *IEEE Trans. on Power Systems*, vol. 28, no. 4, pp. 4909–4917, 2013.
- [20] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas, “MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education,” *IEEE Trans. on Power Systems*, vol. 26, no. 1, pp. 12–19, Feb 2011.
- [21] C. Bowen, T. Buennemeyer, and R. Thomas, “Next generation SCADA security: Best practices and client puzzles,” in *IEEE Workshop on Information Assurance and Security*, Jul. 2005.