

Distribution System Voltage Prediction from Smart Inverters using Decentralized Regression

Zachary R. Atkins
Department of Mathematics
University of Kansas
Lawrence, USA
zatkins@ku.edu

Christopher J. Vogl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, USA
vogl2@llnl.gov

Achintya Madduri
Computational Engineering Division
Lawrence Livermore National Laboratory
Livermore, USA
madduri2@llnl.gov

Nan Duan
Computational Engineering Division
Lawrence Livermore National Laboratory
Livermore, USA
duan4@llnl.gov

Agnieszka K. Miedlar
Department of Mathematics
University of Kansas
Lawrence, USA
amiedlar@ku.edu

Daniel Merl
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, USA
dmerl@llnl.gov

Abstract—As photovoltaic (PV) penetration continues to rise and smart inverter functionality continues to expand, smart inverters and other distributed energy resources (DERs) will play increasingly important roles in distribution system power management and security. In this paper, it is demonstrated that a constellation of smart inverters in a simulated distribution circuit can enable precise voltage predictions using an asynchronous and decentralized prediction algorithm. Using simulated data and a constellation of 15 inverters in a ring communication topology, the COLA algorithm is shown to accomplish the learning task required for voltage magnitude prediction with far less communication overhead than fully connected P2P learning protocols. Additionally, a dynamic stopping criterion is proposed that does not require a regularizer like the original COLA stopping criterion.

Index Terms—smart grid, smart inverters, voltage prediction, distributed computing

I. INTRODUCTION

Power distribution networks are increasingly dynamic due to the growing prevalence of distributed energy resources (DERs), such as rooftop solar photovoltaic (PV) panels. The old distribution network model as a predictable set of loads is no longer accurate. In particular, it is not uncommon for distribution networks with high PV-penetration to be net producers of power during certain parts of the year. As a result, distribution networks with high PV-penetration require more complex monitoring and control objectives compared to traditional networks [1]. Meeting the demands of complex monitoring requires large-scale deployment of smart sensors and efficient processing of large volumes of data [2]. This has motivated recent research into utilizing “on-device” computation to avoid collecting data from all smart devices in one location for processing [3]. The industry standard today is to

collect sensor data from the smart devices in a central control facility, where it is analyzed for the purposes of maintaining a healthy and safe grid. Collecting data in one location, however, can incur large communication cost and can require large data storage capability [4]. Additionally, there can be cybersecurity concerns (e.g. data breaches) [5] and a host of privacy concerns ranging from intellectual property loss [6] to personal data loss [7].

Without addressing these issues, power utilities are unable to safely accept more solar generation in the grid. If the computational resources of smart grid devices, such as smart meters and smart solar inverters can be exploited, the stress on the central control facility can be alleviated by migrating some DER analyses and control decisions to either distributed energy resource management system (DERMS) controllers or the smart inverters themselves [8]. Also, in the event that a cyber attack targeting smart devices is launched [9], computational resources of smart devices can facilitate local command verification to improve the cybersecurity posture of the power grid [10]. Therefore, the development of “on-device” algorithms is essential for power utilities to harness the full potential of DER power generation.

The focus of this work is on the optimization/learning/training of a predictive model across edge devices, such as smart inverters, for the purpose of assessing the health of the grid (e.g. over-voltage issues). There is a vast array of techniques from the optimization and machine learning communities for accomplishing distributed linear regression. That said, the computing and data transfer limitations of a collection of edge devices requires a distributed algorithm that also avoids (i) a central controller/aggregator (decentralized) and (ii) global synchronization of local data (asynchronous). The *COmmunication-Efficient Decentralized Linear LeArning* (COLA) algorithm proposed in [11] is specifically designed for “column-partitioned data”, in which each edge device represents a sort of data silo from which

This work is supported by the U.S. Department of Energy’s Office of Energy Efficiency and Renewable Energy (EERE) under Solar Energy Technologies Office (SETO) Agreement Number 34229. Prepared by LLNL under Contract DE-AC52-07NA27344. LLNL-CONF-816452.

data egress is to be minimized.

This work investigates the applicability of the CoLA algorithm to linear regression training on smart inverter data. This work also suggests a stopping criterion for the CoLA algorithm based on the magnitude of the local iterate updates. Such a stopping criterion enables a broader class of global objective functions than the ones introduced in [11], e.g., those that do not require a regularizer function. The remainder of this paper is organized as follows: Section II describes the CoLA algorithm and proposes a dynamic stopping criterion; Section III formulates the inverter data regression problem; Section IV presents the numerical results of applying the CoLA algorithm to inverter voltage magnitude data; finally Section V draws conclusions.

II. DECENTRALIZED LEARNING METHOD

Consider a collection of n smart inverters collecting local voltage magnitude at regular intervals (e.g. every 15 minutes for a day), resulting in m voltage magnitude values for each inverter. Denote as $\mathbf{A} \in \mathbb{R}^{m \times n}$ the data matrix with each column representing voltage magnitude values from a single inverter. Consider a smart meter also collecting local voltage magnitude at the same times as the smart inverters. Denote as $\mathbf{b} \in \mathbb{R}^m$ the data vector containing the smart meter voltage magnitudes values. A regression model, such as linear least-squares, can be trained on $\mathbf{A}$ and $\mathbf{b}$ and used to predict future smart meter voltage magnitude values from the local smart inverter data. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be the objective/loss function to be minimized to train the regression model. With $g: \mathbb{R}^n \rightarrow \mathbb{R}$ as a convex regularization function, the training of the linear regression is modeled as an optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{Ax}) + g(\mathbf{x}). \quad (1)$$

A. The CoLA Algorithm

The CoLA algorithm [11] will be used to solve (1) in a decentralized fashion across K nodes. The algorithm requires regularizer g to be additively separable, $g(\mathbf{x}) = \sum_{i=1}^n g_i(x_i)$, and objective function f to be of the form $f(\mathbf{v})$, where $\mathbf{v} = \mathbf{Ax}$ is convex, and $(1/\tau)$ -smooth, i.e.,

$$f(\mathbf{y}) \leq f(\mathbf{z}) + \nabla f(\mathbf{z}) \cdot (\mathbf{y} - \mathbf{z}) + \frac{1}{2\tau} \|\mathbf{y} - \mathbf{z}\|_2^2 \quad \forall \mathbf{y}, \mathbf{z} \in \mathbb{R}^n.$$

The algorithm, shown as Algorithm 1, decomposes the global minimization problem (1) into a set of local subproblems, which are then solved independently of one another. The local subproblem for node k is defined as

$$\Gamma_k(\Delta \mathbf{x}_{[k]}; \mathbf{v}_k, \mathbf{x}_{[k]}) = \frac{1}{K} f(\mathbf{v}_k) + \nabla f(\mathbf{v}_k) \cdot \mathbf{A}_{[k]} \Delta \mathbf{x}_{[k]} + \frac{K}{2\tau} \left\| \mathbf{A}_{[k]} \Delta \mathbf{x}_{[k]} \right\|_2^2 + \sum_{i \in \mathcal{P}_k} g_i(x_i + \Delta x_{[k]i}), \quad (2)$$

where $\mathcal{P}_k$ is a set of indices assigned to node k that corresponds to the columns of $\mathbf{A}$ (denoted $\mathbf{A}_{[k]}$) and rows of $\mathbf{x}$ (denoted $\mathbf{x}_{[k]}$). The results from the local subproblems are mixed according to a doubly-stochastic mixing matrix $\mathcal{W} \in \mathbb{R}^{K \times K}$ in an asynchronous manner, making the CoLA

algorithm fully decentralized. For additional details on the relationship between solving the global minimization problem (1) and solving the local subproblems (2), see [11].

Inputs : Data matrix $\mathbf{A}$ partitioned column-wise into $\{\mathbf{A}_{[k]}\}_{k=1}^K$, mixing matrix $\mathcal{W}$, and iteration number $N > 0$.

Output: Solution $\mathbf{x} \in \mathbb{R}^n$ to problem (1)

```

1 Initialize:  $\mathbf{x}^{(0)} := \mathbf{0} \in \mathbb{R}^n$ ;
    $\mathbf{v}_k^{(0)} := \mathbf{0} \in \mathbb{R}^m \quad \forall k = 1 \dots K$ 
2 for  $i = 0 \dots N$  do
3   for  $k = 1 \dots K$  in parallel do
4     Update  $\mathbf{v}_k^{(i+\frac{1}{2})} = \sum_{\ell=1}^K \mathcal{W}_{k\ell} \mathbf{v}_\ell^{(i)}$ 
5     Compute  $\Delta \mathbf{x}_{[k]} \leftarrow \text{solve a local subproblem}$ 
        $\Gamma_k(\Delta \mathbf{x}_{[k]}; \mathbf{v}_k, \mathbf{x}_{[k]})$ 
6     Update  $\mathbf{x}_{[k]}^{(i+1)} := \mathbf{x}_{[k]}^{(i)} + \Delta \mathbf{x}_{[k]}$ 
7     Compute  $\Delta \mathbf{v}_k := \mathbf{A}_{[k]} \Delta \mathbf{x}_{[k]}$  as a local
       estimate update
8     Update a local estimate
        $\mathbf{v}_k^{(i+1)} = \mathbf{v}_k^{(i+\frac{1}{2})} + K \Delta \mathbf{v}_k$ 
9   end
10 end
```

Algorithm 1: CoLA: Communication-Efficient Decentralized Linear LeArning ($\gamma = 1$).

B. Stopping Criterion

As stated, Algorithm 1 will stop after N iterations regardless of whether the iterates have converged. In the case of centralized iterative algorithms, this static iteration count is usually replaced by a dynamic stopping criterion that may or may not depend on the current residual. Because the residual is a global quantity that usually requires synchronization of all nodes to be computed, a major challenge for decentralized iterative algorithms is determining a suitable stopping criterion that can be obtained in a decentralized fashion. The criterion proposed in [11] utilizes a mapping between local data and a bound on the global duality gap. The local data in that mapping includes evaluating the complex conjugate of g_i , which excludes the use of $g(\mathbf{x}) = 0$ and its corresponding complex conjugate $g^*(\mathbf{x}^*) = \sup_{\mathbf{x}^*} \mathbf{x}^* \cdot \mathbf{x}$. As such, a decentralized stopping criterion is sought that indicates progress on minimizing $f(\mathbf{Ax})$ without requiring $g^*(\mathbf{x}^*)$.

Motivated by the distributed algorithms that use the magnitude of $\Delta \mathbf{x}_{[k]}$, the relationship between $\Delta \mathbf{x}_{[k]}$ and $f(\mathbf{Ax})$ is studied, i.e., the relationship between the backward and forward error. It is found that centering and normalizing the columns of data matrix,

$$\mathbf{A}_{ij} \rightarrow \frac{\mathbf{A}_{ij} - \mu_j}{\nu_j}, \quad \text{where} \quad \begin{cases} \mu_j = \frac{1}{m} \sum_{l=1}^m \mathbf{A}_{lj} \\ \nu_j = \left(\frac{1}{m} \sum_{l=1}^m \mathbf{A}_{lj}^2 \right)^{1/2} \end{cases}, \quad (3)$$

have a significant impact on whether decreases in $\Delta \mathbf{x}_{[k]}$ correspond to decreases in $f(\mathbf{Ax})$. Because the data in the COLA algorithm is column-partitioned, (3) can be applied in a fully decentralized fashion: no communication between nodes is necessary to preprocess the data. With the proposed preprocessing, the empirical results herein indicate that the magnitude of $\Delta \mathbf{x}_{[k]}$ is a viable candidate for a stopping criterion.

III. INVERTER REGRESSION

In this section, the utility of a regression-based voltage prediction approach in simulated high-DER-penetration scenarios is demonstrated. To obtain such a network, the base model of the power distribution network shown in Figure 1 is modified to include solar panels at 80% penetration levels. An associated solar panel was added to each load on the system such that the peak power of the solar panels was equal to 80% of the peak demand of the associated load. The open source power grid simulator GridLAB-D [12] is used to obtain the corresponding voltage data. Given that only one distribution network is considered, the results herein are presented as a case study in voltage prediction using a linear regression model.




Fig. 1. Network topology of the power distribution network used for the experiments. This distribution network has a peak demand of approximately 25MW. Solar panels for a uniform 80% grid penetration are indicated as yellow parallelograms with sizes represented relative to their peak capacity. The red arrow indicates the node where the voltage is being predicted by the regression-based algorithm.

Because the node voltages in a high-PV-penetration distribution network can be heavily impacted by the power generated by the distributed solar panels, the impact of the particular training time periods on the ability of the algorithm to predict node voltages accurately is first investigated. The voltage magnitude data obtained from 5 inverters that are downstream from the node of interest are used to train a linear least-squares regression model. The regression is trained on data

obtained at 15 minute intervals for a given day of the year and then used to predict for all other days of the simulated year, where each simulated day involves real stochastic variation in solar irradiation levels, leading to variable PV generation. The process is repeated so that each calendar day is considered for training. The results in Figure 2 show that linear least-squares regression model trained on any day of the year can predict the voltage at the node of interest for any other day with a high degree of accuracy.




Fig. 2. Linear least-square regression testing statistics, where each day is individually used to train a model that is testing on the remaining 364 days (nominal voltage is 7200V).

With some confidence in the applicability of linear least-squares regression to voltage prediction, the focus moves to using the COLA algorithm to train the regression. Voltage magnitude data is now obtained from 16 solar inverters at 15 minute intervals for a given day of the year. A linear least-squares regression model is used to predict values in $\mathbf{b}$ from the data in $\mathbf{A}$:

$$f(\mathbf{Ax}) = \frac{1}{2}[(\mathbf{Ax} - \mathbf{b}) \cdot (\mathbf{Ax} - \mathbf{b})],$$

where $\mathbf{A}$ is all but one column of the dataset, and $\mathbf{b}$ is the final column. It can be shown that f is $(1/\tau)$ -smooth, with $\tau = 1$, by noting

$$\begin{aligned} \frac{1}{2}(\mathbf{y} - \mathbf{z}) \cdot (\mathbf{y} - \mathbf{z}) &= \frac{1}{2}\mathbf{y} \cdot \mathbf{y} - \frac{1}{2}\mathbf{z} \cdot \mathbf{z} - \mathbf{z} \cdot (\mathbf{y} - \mathbf{z}) \\ \Rightarrow \frac{1}{2}\|\mathbf{y}\|_2^2 &= \frac{1}{2}\|\mathbf{z}\|_2^2 + \mathbf{z} \cdot (\mathbf{y} - \mathbf{z}) + \frac{1}{2}\|\mathbf{y} - \mathbf{z}\|_2^2. \end{aligned}$$

The local subproblem (2) for the linear least-squares f can be expressed as

$$\begin{aligned} \Gamma_k(\Delta \mathbf{x}_{[k]}; \mathbf{v}_k, \mathbf{x}_{[k]}) &= \frac{1}{2} \left\| \frac{\mathbf{v}_k - \mathbf{b}}{K} + \mathbf{A}_{[k]} \Delta \mathbf{x}_{[k]} \right\|_2^2 \\ &\quad + g_{[k]}(\mathbf{x}_{[k]} + \Delta \mathbf{x}_{[k]}), \end{aligned}$$

where each $\mathbf{A}_{[k]}$ is the column of $\mathbf{A}$ containing data from inverter k .

The implementation of the COLA algorithm provided in [13] requires a non-zero regularizer g , which is a consequence of the choice of stopping criterion that uses the global duality gap in [11]. To satisfy this requirement, an elastic net regularizer is used:

$$g(\mathbf{x}) = \lambda \left[\frac{\eta}{2}(\mathbf{x} \cdot \mathbf{x}) + (1 - \eta) \sum_{j=1}^m |x_j| \right],$$

where $\eta \in (0, 1]$ and $\lambda > 0$ are adjustable parameters. As mentioned before, a ring topology network is used where

each node k exchanges $\mathbf{v}_k$ values with only two other nodes. In other words, $\mathcal{W}$ is symmetric with only three non-zero entries of $1/3$ in each row/column, with two of those entries being off-diagonal. Another approach would be to have each of the 15 (n) nodes broadcast the local data to all other nodes, so that each node has access to the full data matrix $\mathbf{A}$. The latter approach involves the communication of 210 ($n(n-1)$) column vectors of data. With the ring topology, each CoLA iteration involves the communication of 30 ($2n$) column vectors. Thus, if an acceptable regression is found in less than 7 ($(n-1)/2$) iterations, the CoLA algorithm results in less communication. The specific implementation of the CoLA algorithm used here is available in [14].

IV. RESULTS

To highlight the importance of data preprocessing, the relationship between the global minimization problem (1) and solving the local subproblems (2) is first observed. Figure 3 demonstrates that the sum of the local subproblem objective functions $\Gamma_k(\Delta\mathbf{x}_{[k]}; \mathbf{v}_k, \mathbf{x}_{[k]})$ does indeed bound the value of the global objective function $\mathcal{O}_A(\mathbf{x}) = f(\mathbf{A}\mathbf{x}) + g(\mathbf{x})$. The




Fig. 3. Local subproblems and global objective function values for the inverter voltage dataset (left: original data, right: preprocessed data).

decrease in the sum of $\Gamma_k(\Delta\mathbf{x}_{[k]}; \mathbf{v}_k, \mathbf{x}_{[k]})$ values in Figure 3 is mirrored by the decrease in the values of $\mathcal{O}_A$, regardless of whether the inverter data is preprocessed or not. That said, the preprocessing of the inverters voltage dataset appears to be necessary to avoid the stagnation of $\mathcal{O}_A$ values.

The stagnation in $\mathcal{O}_A$ values, when no preprocessing is performed, is reflected in the resulting regression quality obtained using the original data versus preprocessed data. Figure 4 shows voltage magnitude profiles from both regressions after 500 iterations, along with the corresponding train and test data. The stagnation of $\mathcal{O}_A$ values in the case of the original data is reflected in a voltage magnitude profile that overshoots the peak voltage just before the 5th hour (5 am) and undershoots the voltage dip around the 16th hour (4 pm). On the other hand, the voltage magnitude profile that corresponds to the regression using preprocessed data very accurately captures both the corresponding train and test data.

To observe the effect of data preprocessing on the proposed dynamic stopping criterion, the values $f(\mathbf{A}\mathbf{x})$ and $|\Delta\mathbf{x}_{[k]}|$ are recorded for the first 500 iterations and presented in Figure 5. With the original dataset, $f(\mathbf{A}\mathbf{x})$ exhibits a very slow linear decay: $f(\mathbf{A}\mathbf{x}) \sim \varepsilon n$, with $\varepsilon \ll 1$. Preprocessing the data




Fig. 4. Regression comparison after 500 iterations of the CoLA algorithm applied to the elastic net problem with the inverter voltage data.




Fig. 5. Local updates (dashed) vs loss function (solid) for the inverter voltage dataset (top: original data, bottom: preprocessed data)

results in at least a power-law decay in $f(\mathbf{A}\mathbf{x})$. Regarding using the values $|\Delta\mathbf{x}_{[k]}|$ as a stopping criterion, preprocessing the data appears to be necessary to avoid stagnation in the decrease of values $|\Delta\mathbf{x}_{[k]}|$, as was already seen for the decrease of values $\mathcal{O}_A$. At this point, it is worth noting that the data matrix of the inverters voltage dataset was found to be of low numerical rank: a Lasso regression leads to more zero regression coefficients than non-zero coefficients. Therefore, the stagnation behavior can be explained by the CoLA algorithm oscillating between two or more non-unique solutions of the least-square problem. While the non-uniqueness might be removed or improved by further refinement of the elastic net regularization parameters λ and η , the use of preprocessing avoids the need of parameter tuning and allows for the choice of $g = 0$.

Whereas the regression in Figure 4 is trained on a random sample of the data, a more practical setting might be to train

the regression on a contiguous section of the data. Consider, for example, a regression trained on the first 5 hours of data that is then used to predict voltages for the latter 19 hours. Furthermore, an overvoltage is created by substantially increasing the PV in the 80% distribution network. Two linear least-squares models are trained and tested as to whether the admittedly exaggerated overvoltage scenario is detected. The first model, referred to as the *decentralized model*, is trained using a single communication round of the CoLA algorithm with a Lasso approach ($\eta = 0$). The second model, referred to as the *collocated model*, is generated using a centralized Lasso regression from the SciKit-Learn Python library [15]. Figure 6 shows that CoLA outperforms the collocated regression at accurately predicting both when the overvoltage scenario begins and ends.




Fig. 6. Comparison of Lasso regression with CoLA (decentralized) prediction and sklearn (collocated) prediction (orange dots: training data, blue dots: test data).

V. CONCLUSION

For a particular distribution network augmented with 80% PV penetration, a linear least-square regression is able to accurately predict 364 days of inverter voltage magnitude data after only 1 day of training. The CoLA algorithm is shown capable of producing such a linear least-square regression in a decentralized and asynchronous fashion. Furthermore, a suitable regression for an overvoltage scenario is obtained by CoLA with substantial reduction in communication cost compared to aggregating all the data in one location. Finally, it was found that augmenting the CoLA algorithm with a data preprocessing step is crucial to obtain sustained decay in both the objective function and the local update magnitudes. An investigation into the extent that the results of this work apply to other distribution networks and a theoretical analysis how the preprocessing step affects the CoLA algorithm are subjects of future work.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal

liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

REFERENCES

- [1] W. Wang and F. de León, "Quantitative evaluation of DER smart inverters for the mitigation of FIDVR in distribution systems," *IEEE Transactions on Power Delivery*, vol. 35, no. 1, pp. 420–429, February 2020.
- [2] M. Sanduleac, G. Lipari, A. Monti, A. Voulkidis, G. Zanetto, A. Corsi, L. Toma, G. Fiorentino, and D. Federenciu, "Next generation real-time smart meters for ICT based assessment of grid data inconsistencies," *Energies*, vol. 10, no. 7, pp. 857–873, June 2017.
- [3] N. Duan, J. Cadena, P. Sotorrio, and J. Joo, "Collaborative inference of missing smart electric meter data for a building," in *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, Pittsburgh, PA, October 2019.
- [4] W. Meng, X. Wang, and S. Liu, "Distributed load sharing of an inverter-based microgrid with reduced communication," *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1354–1364, March 2018.
- [5] O. T. Soyoye and K. C. Stefferud, "Cybersecurity risk assessment for California's smart inverter functions," in *2019 IEEE CyberPELS (CyberPELS)*, Knoxville, TN, April 2019.
- [6] K. Sakurama and M. Miura, "Communication-based decentralized demand response for smart microgrids," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 6, pp. 5192–5202, June 2017.
- [7] M. A. Mustafa, N. Zhang, G. Kalogridis, and Z. Fan, "DEP2SA: A decentralized efficient privacy-preserving and selective aggregation scheme in advanced metering infrastructure," *IEEE Access*, vol. 3, pp. 2828–2846, December 2015.
- [8] S. Saxena, N. A. El-Taweel, H. E. Farag, and L. S. Hilaire, "Design and field implementation of a multi-agent system for voltage regulation using smart inverters and data distribution service (DDS)," in *2018 IEEE Electrical Power and Energy Conference (EPEC)*, Toronto, ON, October 2018.
- [9] N. Duan, N. Yee, B. Salazar, J. Joo, E. M. Stewart, and E. Cortez, "Cybersecurity analysis of distribution grid operation with distributed energy resources via co-simulation," in *IEEE Power & Energy Society General Meeting*, Virtual Conference, August 2020.
- [10] I. Zografopoulos and C. Konstantinou, "DERauth: a battery-based authentication scheme for distributed energy resources," in *2020 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, IEEE, August 2020, pp. 560–567.
- [11] L. He, A. Bian, and M. Jaggi, "COLA: Decentralized linear learning," *arXiv e-prints*, p. arXiv:1808.04883, August 2018.
- [12] D. P. Chassin, K. Schneider, and C. Gerkensmeyer, "GridLAB-D: An open-source power systems modeling and simulation environment," in *2008 IEEE/PES Transmission and Distribution Conference and Exposition*, Chicago, IL, April 2008.
- [13] L. He, A. Bian, and M. Jaggi, "GitHub repository," 2018. [Online]. Available: <https://github.com/epfml/cola>
- [14] Z. Atkins, "GitHub repository," August 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.3986158>
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.