Modeling and Predicting Steam Power Plant Condenser Vacuum based on Small-sized Operation Data

Aqli Mursadin
Mechanical Engineering Program, Lambung Mangkurat University, Indonesia
a.mursadin@ulm.ac.id (corresponding author)

Andi Yuwenda Iriyanto
UPK Asam-asam, PT PLN Indonesia Power, Indonesia
andiyuwenda@gmail.com

Received: 21 April 2024 | Revised: 15 May 2024 | Accepted: 25 May 2024

 DOI: https://doi.org/10.48084/etasr.7574

ABSTRACT

The condenser vacuum is an important variable in steam power plants. Monitoring and controlling this variable requires predicting its behavior. This paper develops further Autoregressive-Generalized Autoregressive Conditional Heteroscedasticity (AR-GARCH) models for this purpose, using lagged values of predictors. The predictors include the inlet temperature of the condenser cooling water and the active power of the generator. Models can be adequately trained with small-sized data, making them suitable for use in thermal plants, which are often regularly maintained with operating conditions being reset, rendering past data obsolete. Training and testing were carried out using operation data from an actual steam power plant generating unit during a period in which it faced the prospect of an emergency turbine shutdown. When the models pass all the required statistical tests, they tend to outperform other techniques, including autoregressive neural networks and support vector regression, in terms of prediction. This study also discusses an implementation scenario. The choice of training sizes and model variants can be flexible, enhancing the models' practicality for real operational situations. This study also provides additional directions for further research.

Keywords-AR-GARCH; condenser vacuum; energy infrastructure asset; off-design condition; operation and maintenance; small-sized data; statistical learning; steam power plant; turbine trip

I. INTRODUCTION

The condenser vacuum is an important variable in steam power plants both for monitoring and control [1]. While frequency and voltage are usually the main control points in power generation [2-4], it is well understood that any decrease in condenser vacuum would increase the plant heat rate and, consequently, the operation cost [5]. Based on machine learning techniques, such as regression and neural networks, various prediction methods have been proposed to monitor and control the condenser vacuum [6-8]. Among those, the Long Short-Term Memory (LSTM) network is a promising approach. However, this type of network requires a large amount of training data to work effectively. In [7], a dataset consisting of 177,058 observations on condenser vacuum was used to train an LSTM network to predict future values. Such a large size of data would require constant measurements taken every minute for around four months. On the contrary, many steam plants in practice monitor hundreds of variables, with measurements of each variable taken regularly on an hourly basis. For example, in a typical generating unit of the Asam-asam Steam Power Plant in South Kalimantan, Indonesia, daily operation data only contain 11 observations on each variable, that is, eight two-hour and two three-hour measurement intervals every day. An even smaller size of daily operation data is found in the generating units of the Pulang Pisau Steam Power Plant in Central Kalimantan, Indonesia, that is, 9 daily observations on every variable.

Recently, a statistical approach was proposed to model condenser vacuum time series data [9]. Using two external regressors (predictors), i.e., the inlet temperature of the condenser cooling water and the active power of the generator, it only required a minimum of around 100 observations per variable involved to fit condenser vacuum models effectively. This technique was based on the Autoregressive-Generalized Autoregressive Conditional Heteroscedasticity (AR-GARCH) framework [10-12], a modeling technique originally developed to model volatile time-series data from financial assets. Compared to other statistical approaches that have also been suggested for the same purpose [13-15], this approach can deal effectively with the obvious fact that observations in operation data are autocorrelated over time and that the volatility in the data is often non-homogeneous. In other words, these observations are both autocorrelated and heteroscedastic.

However, to predict the condenser vacuum values, the AR-GARCH model still needs to obtain the predicted values of the predictors. There is a possibility that the lagged values of the predictors can be used for this purpose. Lagged values are those observed several days before prediction, which may also be part of the training dataset. This study aimed to further develop the AR-GARCH modeling approach to predict the condenser vacuum values using lagged predictor values. This is subject to the model being trained using a dataset that can have a minimum of about 100 observations for each variable involved. The novelty of this work lies in the following:

- Use AR-GARCH models to predict system variables in thermal plants based on small-sized operation data
- Use a prediction technique that works for engineering systems that are often regularly maintained with operating conditions reset rendering past data, however big, obsolete, with recently acquired data becoming the actual baseline for future measurements
- Develop a potential ability to monitor and predict the volatile behavior of important variables in a thermal plant system during a critical period before a possible emergency shutdown.

II. METHOD

A. Datasets

The operation data for the condenser vacuum, the inlet temperature of the condenser cooling water, and the active power of the generator were derived from [9]. The dataset came from a generating unit (Unit 2) of the Asam-asam steam power plant, consisting of 13 days of observations, i.e., 143 observations (11 observations per day) per variable, accompanied by reference data from the initial commissioning of the unit. During this period, the unit operated at around 86% of its rated capacity of 65 MW. Table I provides three possible segmentations of the dataset, each with its own training size for model fitting and testing or prediction size (i.e., the number of values to be predicted). Further transformation into 12 equally-spaced observations a day is required for the data.

TABLE I. DATASET SEGMENTATIONS*

Segmentation	Training Size		Testing Size	
	(observations)	(days)	(observations)	(days)
1	99	9	44	4
2	110	10	33	3
3	121	11	22	2

* based on the Asam-asam steam power plant's log sheets at Unit 2

These observations represent an off-design condition that the unit underwent, recording a critical period that led to a potential emergency shutdown of the unit. After those 13 days, no complete observations were available for 8 days, indicating a possible turbine trip. When observations reappeared, the unit operated for 11 more days only at about 75.9% of its rated capacity. In such a situation, the ability to monitor and predict what would happen to the system for a short period ahead could make a big difference in terms of performance or even safety for this energy infrastructure asset.

B. Models

The AR(1)-eGARCH(1, 1) model has been proposed for the generating unit [9]. Predictors are included in the AR part only. The AR(1) part, which is the mean model, can be simplified from its general form [16] as follows:

$$(1-\phi_1 L)(y_t-\mu_t)=\varepsilon_t \tag{1}$$

$$\mu_t=\mu+\textstyle\sum_{i=1}^{k}\delta_i x_{i,t}, \tag{2}$$

where y_t is the t^{th} observation on condenser vacuum, $x_{i,t}$ is the t^{th} observation on the i^{th} predictor (i = 1, …, k), ε_t is the t^{th} observation error, $Ly_t = y_{t-1}$, and μ and δ_i are the regression coefficients to be estimated along with φ_1. At any given t, μ_t is the conditional mean. The eGARCH(1, 1), which is the variance model, can also be simplified from its general form [17] as follows:

$$log\,\sigma_t^2=\omega+\alpha_1 z_{t-1}+\gamma_1(|z_{t-1}|-E|z_{t-1}|)+\beta_1\sigma_{t-1}^2 \tag{3}$$

where ω, α_1, β_1, and γ_1 are also parameters of the model, and

$$z_t=(y_t-\mu_t)/\sigma_t \tag{4}$$

with $E|z_t|$ being the expected value of z_t. Among the choices of conditional distributions of z_t are the normal and the Student's t distributions, as well as their skew variants [18]. At any given t, σ_t^2 is the conditional variance that represents the volatility in the data.

All the above parameters are estimated simultaneously using the maximum likelihood method [12]. To ensure its validity, a fitted model is then tested for serial correlation (the Ljung-Box test) [19], autoregressive conditional heteroscedasticity (the ARCH-LM test) [10], leverage effects (the sign bias test) [20], distributional fitness (the Pearson goodness-of-fit test) [21], and parameter stability (the Nyblom test) [22]. Overfitting is taken care of by model selection using the Akaike Information Criterion (AIC) [23]. All the fitting and testing procedures were carried out using the R statistical programming language [24]. This study used the R package 'rugarch' [25] to provide the necessary functions for this purpose.

C. Prediction and Performance

The inlet temperature of the condenser cooling water and the active power of the generator are used as predictors. The inlet temperature of the cooling water is a critical variable in a condenser, as it must be sufficiently low to allow effective heat transfer between the steam and the cooling water. Low effectiveness means that not enough heat is removed from the steam to ensure a sufficiently high vacuum inside the condenser's shell. The active power of the generator is basically its load. A higher load requires more steam to be supplied to the turbine and subsequently cooled in the condenser. As the cooling load of the condenser increases, its ability to maintain a sufficiently high vacuum may decrease. Obviously, variations in both variables affect how the condenser vacuum behaves.

The lagged values of these predictors are taken from the corresponding training set, i.e., the most recent ones in the dataset. These are the closest observations to those predicted. In particular, the size of the lag is taken to equal the prediction size. Prediction performance is measured using the following metrics [26]: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Root Mean Squared Error (RMSE). The validity of the results is examined by comparing the performance to that of two other time series prediction techniques, namely, the autoregressive neural network (NNAR) and the Support Vector Regression (SVR), as implemented in R and used through the packages 'forecast' [27] and 'e1071' [28], respectively.

III. RESULTS AND DISCUSSIONS

A. Fitting the Model Variants

This study uses two AR(1)-eGARCH(1, 1) models. The first has the temperature of the condenser cooling water inlet as the only predictor. The second has two predictors: the inlet temperature of the condenser cooling water and the active power of the generator. Table II shows six variants of the models fitted by varying the conditional distributions across different dataset segmentations. Variants 4 and 5 fail the Nyblom test, meaning that they have some non-constant parameters. This may affect the resulting predictions by the corresponding variants. Adding generator active power as another predictor improves AIC for any segmentation. In general, increasing the training dataset also improves AIC.

TABLE II. MODEL VARIANTS

Variant	Predictor(s)	Segmentation	Distribution	AIC
1	T	1	Student's t	-7.8578
2	T	2	normal	-7.9201
3	T	3	Student's t	-7.9859
4	T, P	1	Student's t	-8.2393
5	T, P	2	normal	-8.2949
6	T, P	3	Student's t	-8.4511

T: condenser cooling water inlet temperature, P: generator active power

B. Prediction Results

Table III compares the prediction performance of AR(1)-eGARCH(1, 1) with that of NNAR and ε-SVR. The NNAR model consists of 50 2-2-1 networks, each with 9 weights, 1 lag, and 2 hidden layers, and is then optimized over AIC. The ε-SVR model is based on the radial kernel, $\gamma = 1$, and further tuned for ε and cost. The results suggest that increasing the training size tends to increase the metrics for variants 1, 2, and 3. This is an interesting phenomenon, assuming that a bigger size should normally result in lower performance metrics. However, given the small size of the data, such a rule may not always apply. Adding the active power of the generator as another predictor tends to reduce the performance of AR-eGARCH but improve NNAR and SVR. However, a better fit of a model does not necessarily translate into better prediction performance. However, adding this predictor appears to stabilize the performance across varying lags for all variants. After all, in terms of performance, the difference between variants with the same training size does not seem significant.

TABLE III. PREDICTION PERFORMANCE

Model	Variant	MAE+	MAPE	RMSE+
AR-eGARCH	1	0.0049	0.0067	0.0062
	2	0.0052	0.0070	0.0065
	3	0.0053	0.0073	0.0068
	4	0.0049	0.0067	0.0062
	5	0.0060	0.0081	0.0075
	6	0.0054	0.0074	0.0070
NNAR*	1	0.0053	0.0073	0.0066
	2	0.0052	0.0070	0.0066
	3	0.0070	0.0096	0.0088
	4	0.0048	0.0066	0.0058
	5	0.0054	0.0073	0.0067
	6	0.0067	0.0091	0.0085
SVR*	1	0.0047	0.0064	0.0061
	2	0.0054	0.0073	0.0067
	3	0.0057	0.0078	0.0073
	4	0.0051	0.0069	0.0064
	5	0.0055	0.0075	0.0068
	6	0.0055	0.0076	0.0071

*vary due to random weights and tuning; +in 0.1 MPa.

The performance of AR-eGARCH tends to be the highest among the others for variants 1, 2, and 3. For variants 4 and 5, the performance of AR-eGARCH looks average compared to that of NNAR and SVR. These variants have some non-constant parameters that might affect prediction badly. Nevertheless, it is certainly the highest for variants 3 and 6, that is, with the biggest training size. These comparison results also confirm the validity of the predictions made by all AR-eGARCH variants.

The results are also shown in Figure 1. In general, all variants visually demonstrate their ability to capture data dynamics. All AR-eGARCH variants can reliably produce predictions with patterns that mimic those of the actual data. Repeatability is also demonstrated by sequentially changing the prediction size from variants 1 to 3 and 4 to 6, in which the ability to capture the data dynamics remains consistent. On the other hand, it is possible to take the average of the training data and report it as a prediction with relatively high performance. However, the plot of this prediction would not show any resemblance to the actual data and cannot be considered reliable.

C. An Implementation Scenario

If the condenser vacuum is too low, which means that the pressure inside the condenser shell is too high, it becomes difficult for the incoming steam from the turbine to enter the condenser. Because of this, too much heat accumulates at the final stage of the turbine, and the temperature of the exhaust hood becomes too high for the system to operate safely. When this happens, the system has to either operate at a much lower capacity or be shut down.

The last four days, as observed in the data, can be considered critical for the generating unit in question. The lowest condenser vacuum in that period was -0.0718 MPaG at an early point on the first day, while the highest exhaust hood temperature was 70.3°C, as observed very late on the fourth day. Such a temperature is too high for normal operation. As a comparison, the average exhaust hood temperature during the first 9 days of the data was 65.9°C (or 66.1°C for the entire dataset), and was already very high. Although the effect of condenser vacuum on exhaust hood temperature might be lagged, the turbine should have been shut down even before entering the first day of the period. In fact, an exhaust hood temperature as high as 70.7°C was also observed on the first day of the training set. However, it might be too early to stop the turbine for that reason.

Had the system implemented an emergency shutdown procedure based on an AR(1)-eGARCH(1, 1) model for prediction as proposed, it would have discovered, for instance, a maximum condenser pressure of -0.07192538 MPaG on the second day using variant 1, or -0.07198573 MPaG at exactly the point where it was observed (i.e., -0.0718 MPaG) on the first day using variant 4. Both would be too high. As a comparison, the average condenser vacuum during the initial commissioning of the unit was -0.0909 MPaG. In this case, variant 4 gives an earlier warning than variant 1. For either variant, such a procedure would not wait a few more days for the lagged effects to develop into high exhaust hood temperature just to begin operating at a lower capacity or to initiate a turbine shutdown. All other variants, except variant 5, also give their predicted minimum condenser vacuum on their first prediction day.

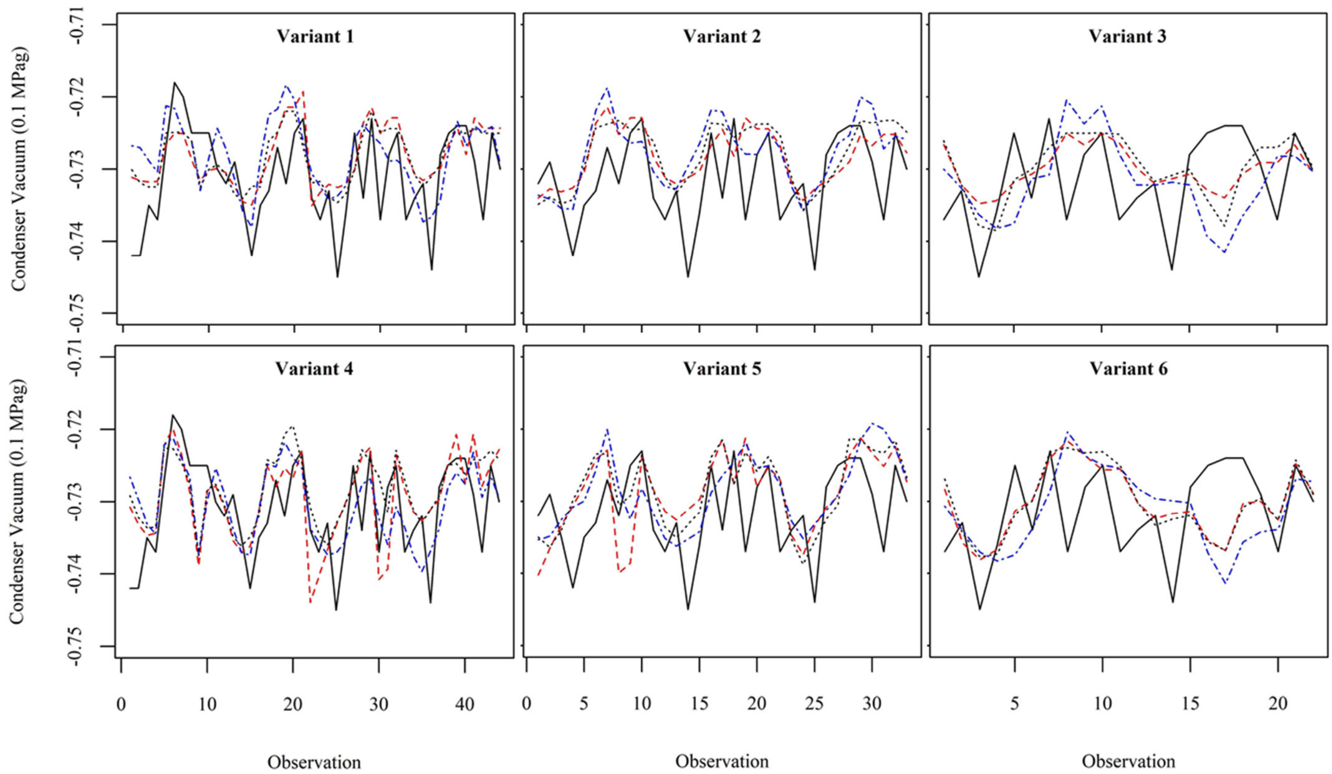

Fig. 1. Prediction plots (- actual, -- AR-eGARCH, -- NNAR, ··· SVR).

It is important to note that cool, low-pressure steam at the later stage can act as a coolant. Unfortunately, lowering steam flow cancels out such a beneficial effect, whereas the windage, which is the friction between the rotor blades and the near-stationary steam, at such low flow produces additional heat on the blades' surface. It should be noted that continuous operation at low steam flow is yet another cause of high exhaust hood temperature [5]. Therefore, an attempt to cope with the effect of low condenser vacuum on rising exhaust hood temperature by continuing to operate at lower steam flow would appear to be an attempt to remove one cause by introducing another. This can even exacerbate the problem.

Here is a real example in which an attempt to reduce steam flow to contain the above-mentioned effect of condenser vacuum improves the cause but worsens the effect. During the entire period presented by the dataset, where the average active power was roughly 87.8% of the rated capacity, the average exhaust hood temperature was 66.1°C with a maximum of 70.3°C. This was with an average main-steam flow of 98.0% of the initial commissioning value. When the observations reappeared eight days later, the average main-steam flow for the following 11 days was reduced to 89.2% of its initial commissioning value, pushing the average active power to as low as 75.9% of the rated capacity and the average exhaust hood temperature to reach 67.1°C with a maximum of 70.5°C. Although the minimum vacuum was improved from -0.0715 MPaG to -0.0723 MPaG, as already expected, it still suggests that reducing steam flow is hardly a desirable solution to the problem.

Furthermore, all variants also capture the trend in condenser vacuum in their respective prediction. For variants 1 and 4, this is shown in Figure 2. This is based on a simple moving average taken for every 11 consecutive predicted values. Along with the minimum predicted vacuum, it suggests that an emergency procedure should also be considered as soon as such a prediction becomes available.

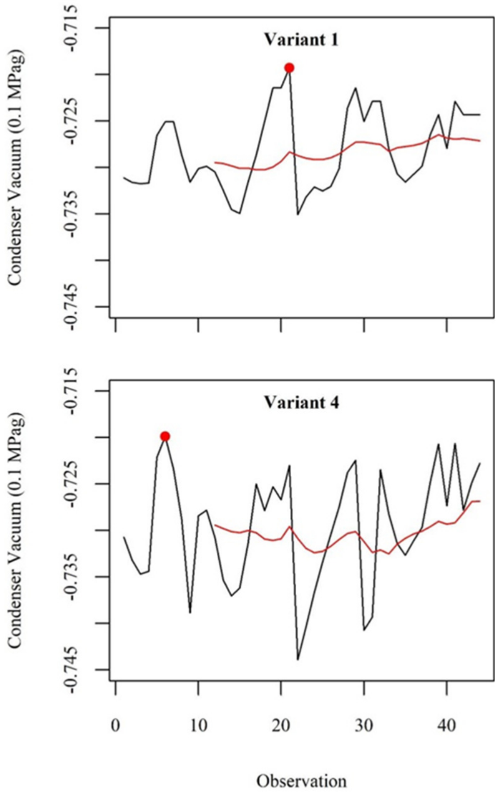

Fig. 2. AR(1)-eGARCH(1, 1) trend prediction (- predicted values, - predicted trend, ● lowest predicted value).

D. The Choice of Models and Training Sets

One main concern with model fitting in statistics is how large the standard error is for each estimate. Fortunately, the size around 100 is not too small to estimate a statistical model such as AR(1)-eGARCH(1, 1) using the maximum likelihood method. In fact, for all variants trained so far, the standard errors are still reasonably small. Another concern is the power of each statistical test carried out on the models. It is at present difficult to measure and might take a series of simulations best left for future work.

By examining the results produced by the different variants, one may notice that the difference in performance between two variants of the same training size is not so significant. Also, despite changes in performance, small changes in the training size can still maintain the models' ability to capture the correct trend in the actual values of the condenser vacuum. This offers flexibility in choosing between models with only one predictor (condenser cooling water inlet temperature) and models with two (condenser cooling water inlet temperature and generator active power). However, it is recommended that both models should be implemented simultaneously while being cautious when a model does not pass some of the tests. This also offers flexibility in choosing the training size. However, it is desirable to train a model as soon as the required size (e.g., 99 in this case) becomes available. As long as the model fits and passes all tests, it is ready to be used for prediction. Even if it does not pass the Nyblom test, for instance, it can still come up with reasonably accurate predictions (see variants 4 and 5). As more observations become available, the training size may grow.

There is no guarantee that a specified model fits a given training set or passes all required tests. However, the flexibility to expand or even shrink the training set and choose the predictors, as well as the conditional distribution or model parameters, improves the chance of finding a suitable model.

IV. CONCLUSIONS AND DIRECTIONS FOR FURTHER RESEARCH

The importance of condenser vacuum in steam power plants has previously been highlighted, and this study explored a knowledge gap in predicting it using AR-GARCH models. The novelty of this work lies in three aspects: AR-GARCH models for predicting system variables in thermal plants with small-sized operation data, a prediction technique that works for engineering systems that are often regularly maintained but increasingly under off-design conditions, and a potential capability of monitoring and predicting the volatile behavior of a thermal plant system's important variables, even during a critical period before a possible emergency shutdown. Two AR-GARCH models were developed even further for predicting condenser vacuum, which can be trained with relatively small data sizes. These models use lagged predictor values. Six variants were tested for prediction performance, showing that they perform relatively well compared to two other techniques.

An implementation scenario was discussed to an extent that covers some important engineering implications. The models can offer another way to deal with an emergency situation related to low vacuum in a steam power plant. This is based on information provided by the predicted vacuum, as well as its predicted trend. Instead of waiting for the lagged effects of low vacuum to translate into high exhaust hood temperature, emergency actions can be initiated whenever the prediction returns information on the future behavior of the condenser vacuum. The flexibility in choosing the model and the training size was discussed, showing that it can potentially improve the practicality of the technique for use in real operational situations.

One major problem with small operation data is that they seldom capture the operation's long-term seasonal pattern. While AR-GARCH models natively work with large data, the presence of such a data pattern requires a correct set of parameters to be included in their formulation before training. As for the use of lagged values, the technique relies on the closeness between the values in the time series. While closeness often leads to similarity, a rare event such as a sudden turbine trip may still occur, resulting in extremely volatile observations and rendering some lagged values unrepresentative.

Future work can focus on the use of Autoregressive Fractionally Integrated Moving Average (or ARFIMA) instead of simply AR to derive more sophisticated ARFIMA-GARCH models for the condenser vacuum. Another direction may be toward the development of a procedure for an emergency turbine trip based on the condenser vacuum predicted patterns given a small window of observations. The possible involvement of predicted patterns, as exemplified using variants 1 and 4, suggests that conformity or, rather, non-conformity can also be used in monitoring condenser vacuum in addition to exact point prediction. Specifically, this can possibly be implemented with a conformal prediction that provides prediction intervals.

REFERENCES

[1] S. Basu and A. K. Debnath, Power Plant Instrumentation and Control Handbook: A Guide to Thermal Power Plants. London, UK: Academic Press, 2014.

[2] Z. R. Labidi, H. Schulte, and A. Mami, "A Systematic Controller Design for a Photovoltaic Generator with Boost Converter Using Integral State Feedback Control," *Engineering, Technology & Applied Science Research*, vol. 9, no. 2, pp. 4030–4036, Apr. 2019, https://doi.org/10.48084/etasr.2687.

[3] D. V. Doan, K. Nguyen, and Q. V. Thai, "Load-Frequency Control of Three-Area Interconnected Power Systems with Renewable Energy Sources Using Novel PSO~PID-Like Fuzzy Logic Controllers," *Engineering, Technology & Applied Science Research*, vol. 12, no. 3, pp. 8597–8604, Jun. 2022, https://doi.org/10.48084/etasr.4924.

[4] M. Masmali, M. I. Elimy, M. Fterich, E. Touti, and G. Abbas, "Comparative Studies on Load Frequency Control with Different Governors connected to Mini Hydro Power Plant via PSCAD Software," *Engineering, Technology & Applied Science Research*, vol. 14, no. 1, pp. 12975–12983, Feb. 2024, https://doi.org/10.48084/etasr.6722.

[5] D. K. Sarkar, Ed., "Thermal Power Plant - Pre-Operational Activities," in *Thermal Power Plant*, Amsterdam, Netherlands: Elsevier, 2017.

[6] H. Kumar, Rahul, S. Verma, and S. Bera, "Analysis of Machine Learning algorithms for prediction of Condenser Vacuum in Thermal Power Plant," in *2020 International Conference on Electrical and Electronics Engineering (ICE3)*, Gorakhpur, India, Feb. 2020, pp. 778–783, https://doi.org/10.1109/ICE348803.2020.9122848.

[7] Z. Sun *et al.*, "A Novel Condenser Vacuum Degree Prediction Model Based on LSTM and MemN2N," *Journal of Physics: Conference Series*, vol. 2294, no. 1, Mar. 2022, Art. no. 012030, https://doi.org/10.1088/1742-6596/2294/1/012030.

[8] P. He *et al.*, "Condenser Vacuum Degree Prediction Model with Multi-View Information Fusion," *Journal of Physics: Conference Series*, vol. 2294, no. 1, Mar. 2022, Art. no. 012032, https://doi.org/10.1088/1742-6596/2294/1/012032.

[9] A. Mursadin and A. Y. Iriyanto, "Modeling Steam Power Plant Condenser Vacuum under Off-design Conditions: A Statistical Approach," in *2023 International Conference on Advanced Mechatronics, Intelligent Manufacture and Industrial Automation (ICAMIMIA)*, Surabaya, Indonesia, Nov. 2023, pp. 266–271, https://doi.org/10.1109/ICAMIMIA60881.2023.10427916.

[10] R. F. Engle, "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, vol. 50, no. 4, pp. 987–1007, 1982, https://doi.org/10.2307/1912773.

[11] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, Apr. 1986, https://doi.org/10.1016/0304-4076(86)90063-1.

[12] C. Francq and J. M. Zakoian, *GARCH Models: Structure, Statistical Inference and Financial Applications.* Hoboken, NJ, USA: John Wiley & Sons, 2019.

[13] L. Pan, D. Flynn, and M. Cregan, "Statistical Model for Power Plant Performance Monitoring and Analysis," in *2007 42nd International Universities Power Engineering Conference*, Brighton, UK, Sep. 2007, pp. 121–126, https://doi.org/10.1109/UPEC.2007.4468931.

[14] K. Świrski, "Power Plant Performance Monitoring Using Statistical Methodology Approach," *Journal of Power Technologies*, vol. 91, no. 2, pp. 63–76, 2011.

[15] S. Chandrasekharan, R. C. Panda, and B. N. Swaminathan, "Statistical modeling of an integrated boiler for coal fired thermal power plant," *Heliyon*, vol. 3, no. 6, Jun. 2017, Art. no. e00322, https://doi.org/10.1016/j.heliyon.2017.e00322.

[16] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control.* John Wiley & Sons, 2015.

[17] D. B. Nelson, "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, vol. 59, no. 2, pp. 347–370, 1991, https://doi.org/10.2307/2938260.

[18] C. Fernández and M. F. J. Steel, "On Bayesian Modeling of Fat Tails and Skewness," *Journal of the American Statistical Association*, vol. 93, no. 441, pp. 359–371, Mar. 1998, https://doi.org/10.1080/01621459.1998.10474117.

[19] G. M. Ljung and G. E. P. Box, "On a measure of lack of fit in time series models," *Biometrika*, vol. 65, no. 2, pp. 297–303, Aug. 1978, https://doi.org/10.1093/biomet/65.2.297.

[20] R. F. Engle and V. K. Ng, "Measuring and Testing the Impact of News on Volatility," *The Journal of Finance*, vol. 48, no. 5, pp. 1749–1778, 1993, https://doi.org/10.1111/j.1540-6261.1993.tb05127.x.

[21] H. Pham, "Basic Statistical Concepts," in *Springer Handbooks*, Springer, 2006, pp. 3–48.

[22] J. Nyblom, "Testing for the Constancy of Parameters over Time," *Journal of the American Statistical Association*, vol. 84, no. 405, pp. 223–230, Mar. 1989, https://doi.org/10.1080/01621459.1989.10478759.

[23] J. de Leeuw, "Introduction to Akaike (1973) Information Theory and an Extension of the Maximum Likelihood Principle," in *Breakthroughs in Statistics: Foundations and Basic Theory*, S. Kotz and N. L. Johnson, Eds. New York, NY: Springer, 1992, pp. 599–609.

[24] "R: The R Project for Statistical Computing." https://www.r-project.org/.

[25] A. Galanos and T. Kley, "rugarch: Univariate GARCH Models." Sep. 20, 2023, [Online]. Available: https://cran.r-project.org/web/packages/rugarch/.

[26] A. Botchkarev, "A New Typology Design of Performance Metrics to Measure Errors in Machine Learning Regression Algorithms," *Interdisciplinary Journal of Information, Knowledge, and Management*, vol. 14, pp. 045–076, Jan. 2019.

[27] R. Hyndman *et al.*, "Forecasting Functions for Time Series and Linear Models." [Online]. Available: https://pkg.robjhyndman.com/forecast/.

[28] D. Meyer *et al.*, "e1071: Misc Functions of the Department of Statistics, Probability Theory Group (Formerly: E1071), TU Wien." [Online]. Available: https://cran.r-project.org/web/packages/e1071/.