Advance Publication by J-STAGE

Mechanical Engineering Journal

DOI:10.1299/mej.20-00534

Received date : 15 December, 2020

Accepted date : 24 June, 2021

J-STAGE Advance Publication date : 3 July, 2021

Anomaly sign detection by monitoring thousands of process values using a two-stage autoencoder

Susumu NAITO*, Yasunori TAGUCHI*, Yuichi KATO*, Kouta NAKATA*, Ryota MIYAKE**,
Isaku NAGURA**, Shinya TOMINAGA** and Toshio AOKI**

* Corporate Research & Development Center, Toshiba Corporation

1 Komukai-Toshiba-cho, Saiwai-ku, Kawasaki 212-8582, Japan

E-mail: susumu.naito@toshiba.co.jp

** Isogo Nuclear Engineering Center, Toshiba Energy Systems & Solutions Corporation

8 Shinsugita-cho, Isogo-ku, Yokohama 235-8523, Japan

Abstract

In a large-scale plant such as a nuclear power plant, thousands of process values are measured for the purpose of monitoring the plant performance and the health of various systems. It is difficult for plant operators to constantly monitor all of the process values. We present a data-driven method to comprehensively monitor a large number of process values and detect early signs of anomalies, including unknown events, with few false positives. In order to learn the complex changing internal state of a nuclear power plant and accurately predict the normal process values, we created a two-stage autoencoder composed of a time window autoencoder and a deviation autoencoder, which is a deep learning network structure corresponding to the characteristics of the process values. We assessed performances of the two-stage autoencoder with simulated process values of a nuclear power plant, a 1,100 MW boiling water reactor having 3,100 analog process values. In situations where it is difficult to predict the normal state (rapid operation mode change, transient state, and small fluctuations in the process values), the two-stage autoencoder properly predicted the normal process values and showed excellent performances with early detection and zero false positives, except for one case. The two-stage autoencoder would be an effective solution for comprehensive plant monitoring and early detection of anomaly signs.

Keywords : Plant monitoring, Big data, Anomaly sign detection, Machine learning, Deep learning, Autoencoder

1. Introduction

In a large-scale plant such as a nuclear power plant, thousands of process values are measured to monitor the plant performance and the health of various systems. It is difficult for plant operators to constantly monitor all of the process values. We aim to comprehensively monitor a large number of process values and detect early signs of anomalies, including unknown events, with few false positives, without limiting the monitoring targets to specific components and operating conditions (Fig. 1). With the decrease in the number of skilled operators and the shortage of manpower, there is a growing need for such technology to reduce the workload on the site and to operate the plant safely and efficiently.

Numerous studies have been conducted on plant monitoring with data-driven methods (Ayo-Imoru and Cilliers, 2018), such as principal component analysis (PCA) for sensor fault detection and support vector regression (SVR) for component degradation monitoring. Data-driven models predict the future state of the plant depending on past patterns of the system. They try to define relationships between variables by means of data fitting. Many of the previous studies, including the examples above, limited the target to specific anomalies, components, or small systems that can be handled with a small number of variables. The objective of our study is to detect early signs of anomalies from the dynamic correlation of a large number of process values without limiting the target, using machine learning suitable for processing large amounts of data.

Our study corresponds to the first process in the series processes of anomaly detection, diagnosis, and countermeasures. Our method is used to detect early signs of anomalies. Expert knowledge and highly reliable physical model-based analysis then determine if the anomalies are true. This procedure aims to detect anomalies at the early stage or before the spread of damage.

Anomalies are detected from differences between the process values and predicted values. The predicted values should agree with the normal process values and be different from the anomalous process values. However, a plant system is a complex system composed of various systems and components. Even under normal conditions, the internal state of a plant system complexly changes. The features of the plant process data are described below.

- It is a multivariate problem with a large number of variables.
- There are rapid changes of the process values that are difficult to predict from their immediate behaviors, such as a change of an operation mode, pump switching, valve operation, etc. These cause false positives.
- Multiple process values with different response times change in tandem with each other. This is seen in the transient states.
- Many process values have steady and minute fluctuations. In the initial stage of an anomaly, the anomaly is buried in the minute fluctuations, which makes its early detection difficult. From the analysis of actual data, we deduce that the minute fluctuations are due to external factors (exogenous noise and vibration, etc.).

These cause learning difficulties to arise in machine learning. In order to accurately predict the normal process values, it is necessary to be able to learn (1) complex correlations of a large number of process values, (2) rapid changes of the process values such as a change of an operation mode, (3) changes of the process values over time such as the transient states, and (4) minute fluctuations of the process values due to external factors, i.e., external disturbance.

In order to address issues (1) – (4), we created a two-stage autoencoder composed of a time window autoencoder learning the major fluctuations of the time series process values and a deviation autoencoder learning the minute fluctuations (Naito, et al., 2020). The details are described in the next section. This is a deep learning network structure corresponding to the characteristics of the process values that accurately predicts the normal process values. With simulation data of a nuclear power plant (1,100 MW boiling water reactor, 3,100 analog process values), we assessed performances of the two-stage autoencoder: a predicting ability for rapid changes of the process values for issue (2), detection performances in the transient state for issue (3), and those under realistic noisy conditions for issue (4).

(a) Monitoring of systems.

(b) Monitoring of each process value.

Fig. 1 Example of screens of anomaly sign detection system. (a) Top screen and (b) deviation ranking in one system and trend data of process values and the corresponding predicted values.

2. Two-stage autoencoder

Figure 2 shows the configuration of the two-stage autoencoder. Figures 3 (a) to (c) are explanatory drawings of autoencoders. As described in Sec. 1, in order to accurately predict the normal process values, it is necessary to be able to learn (1) complex correlations of a large number of process values, (2) rapid changes of the process values such as a change of an operation mode, (3) changes of the process values over time such as the transient states, and (4) minute fluctuations of the process values due to external factors. To address issues (1) and (2), we adopted an autoencoder

(Hinton and Salakhutdinov, 2006), one of the deep learning algorithms (Chalapathy and Chawla, 2019). Figure 3 (a) shows a simple autoencoder. An autoencoder is an unsupervised algorithm that learns neural networks so as to set the inputs equal to the outputs. The autoencoder predicts the normal process values by learning the past normal process values. Since the autoencoder can learn the high-dimensional and non-linear correlations, it can learn complex correlations across the multiple process values. Moreover, a rapid change of the process values at time t , which is difficult to predict from its past time series data (time $t-1$, $t-2$, ...), can be predicted from correlations of multiple process values at time t . This has an important bearing on reducing false positives.

Next, we explain how to deal with issue (3). In anomaly detection of multivariate time series data in the field of machine learning, which is becoming increasingly active, a time window for cutting out the data (Fig. 3 (b)) is generally introduced to learn the temporal correlation with an autoencoder (Audibert, et al., 2020). This is because it is difficult to learn a large amount of the time series data for a long period of time as it is. For the same reason, we also adopted the time window. In a Time window AutoEncoder (TAE, Fig. 3 (b)), the data cut out in a time window are used as the input. With this method, the waveforms (the temporal fluctuations) of the process values in the time window can be learned, and by setting the time window width appropriately, the characteristic temporal fluctuations included in the data can be learned. The time window is used by sliding it, and in plant monitoring, the output values at the right end of the time window (the latest predicted values) are used as the predicted values.

However, the minute fluctuations of the process values (issue (4)) cannot be learned well with a single TAE. A TAE with one intermediate layer (shallow TAE) cannot learn the minute fluctuations due to lack of expressiveness. When the TAE is made into a multi-layered deep network (deep TAE) to learn the minute fluctuations, the predicted values become in good agreement with the normal process values including the minute fluctuations, but they also become in agreement with the minute anomalies in the process values. The anomalies cannot be detected. This is because the minute fluctuations have a very small contribution to the value of the loss function, so they are ignored as negligible smallness in the shallow TAE, and they tend to be learned as identity functions in the deep TAE. In the design of the loss function, it is possible to increase the contribution of the minute fluctuations, but there is a trade-off relationship between increasing the contribution of the minute fluctuations and decreasing the learning accuracy of non-small fluctuations. Therefore, based on the analysis of the actual data in several plants, we focused on the fact that the minute fluctuations can be extracted as the differences (deviations) between the process values and the predicted values from the shallow TAE, and that there is almost no correlation between the predicted values and the deviations (the minute fluctuations). The predicted values from the shallow TAE reproduce well the behavior of the normal process values except for the minute fluctuations. The minute fluctuations are deduced to be mainly due to external factors (exogenous noise and vibration, etc.) that have little relation to the dynamic characteristics of the plant. Because of this uncorrelativeness, as a model for learning the minute fluctuations well, we devised a configuration in which only the deviations are learned by another independent autoencoder, the outputs are added to the predicted values from the shallow TAE, and the added values are used as the final predicted values. We named this configuration two-stage autoencoder.

In the study of anomaly detection of multivariate time series data in the machine learning field, the most recent studies are UnSupervised Anomaly Detection (USAD) (Audibert, et al., 2020) as an autoencoder, Omni Anomaly (OA) (Su, et al., 2019) as a Recurrent Neural Network (RNN), and LSTM-VAE (Park, et al., 2018) as a hybrid of autoencoder and RNN. In the USAD, an autoencoder architecture for adversely training in an unsupervised way was proposed. In the OA, a stochastic recurrent neural network with a stochastic variable connection and a planar normalizing flow was proposed. The LSTM-VAE combines the LSTM with a variational autoencoder (VAE) by replacing the feed-forward network in a VAE with a LSTM. These are studies that enhance the expressiveness of a single model for learning robust multivariate time series' representations from the perspective of versatility in a wide range of fields. On the other hand, our study is specialized in power plant process values. Using two autoencoders with different learning targets, based on the features found from the analysis of the actual plant data, is an approach not found in other studies.

We explain a two-stage autoencoder (Fig. 2). By making the first-stage TAE a shallow network with one intermediate layer, the output becomes a curve that captures the major temporal fluctuation of the normal process value. Minute fluctuations in the process values are extracted as the differences (deviations) between the process values and the outputs. By learning only the deviations in the second-stage autoencoder (Deviation AutoEncoder (DAE) in Fig. 3 (c)), the minute fluctuations of the process values can be learned well. In this way, the two-stage autoencoder learns signals with different properties to predict normal process values with high accuracy.

In Fig 3 (c), the DAE has no time window structure. It copes with the external disturbance with no temporal correlation, such as exogenous noise. When the external disturbance has temporal information, the DAE with the time window is used.

Fig. 2 Configuration of two-stage autoencoder.

(a) Simple autoencoder

(b) Time window autoencoder

(c) Deviation autoencoder

Fig.3 Explanatory drawings of (a) simple autoencoder, (b) time window autoencoder, and (c) deviation autoencoder.

3. Evaluation conditions

3.1 Methods

We assessed the predicting ability for rapid changes (issue (2)) and the anomaly detection performances in the transient state (issue (3)) and those under the external disturbance (issue (4)) with simulated process values of a nuclear power plant. For comparison, the simple autoencoder in Fig. 3 (a) was also assessed.

3.2 Datasets

The target plant was a 1,100 MW boiling water reactor with 3,100 analog process values. Table 1 shows seven groups of the process values classified by system. 2,720 process values were selected, excluding the electrical and the auxiliary systems, because the electrical system has no signs of anomalies (sudden occurrences) and the auxiliary system is less important. The analog process values include temperature, pressure, fluid flow, vibration, radiation, valve opening, etc.

Table 2 shows target cases. Cases 1 – 3 were in the transient state (issue (3)). Cases 4 – 6 were in the steady state under the external disturbance (issue (4)). Case 7 was a normal operation mode change in the steady state (issue (2)). The process value of interest is the process value that changes because of the anomaly. In the transient state, we selected anomalies that can actually occur in the transient state. In the steady state, the leakage and rupture events were selected. Their early detection is advantageous in that the plant can be shut down before the damage expands, a grace period can be obtained before shutdown, and the cause investigation can be conducted in the operating state. In the operation mode change, RSW pump switching was selected, which is often done.

Table 3 shows the training and evaluation data. The simulation data were created at 1-minute intervals using a full-scope nuclear power plant simulator (Yonezawa, et. al., 2017). As the external disturbance, artificial noise was added to the simulation data. The anomalies were generated with the malfunction occurrence function.

Next, we explain the external disturbances. The simulation data have no minute fluctuations. Therefore, the simulation data are the learning targets for the TAE. As mentioned in Sec. 2, the minute fluctuations are deduced to be mainly due to external factors that have little relation to the dynamic characteristics of the plant. Therefore, as learning targets for the DAE that simulate the actual data, the external disturbances (artificial noise) that have independent correlation with the simulation data were added to the simulation data. The artificial noise was generated with an amplitude according to a normal distribution. The standard deviation σ of the normal distribution was set to be a median value (0.01) of the (σ / d) distribution of the actual process values, where d is the difference between the maximum and the minimum values of the normal process value in the training data. We also assessed under the highly noisy condition of the maximum value (0.1) of the (σ / d) distribution in order to make it easier to understand the behavior of the two-stage autoencoder. The artificial noise was assumed to be synchronized between 10 process values (one group of 10 process values). Within each group, the same normalized noise amplitude (σ / d) was used. As an example, Figure 4 shows heat maps of correlation coefficients of the process values for the F system in Table 1. In Fig. 4 (a), it can be confirmed that the artificial noise synchronized between the 10 process values. In the actual data, the groups of the process values are not diagonally aligned, and groups with different numbers of the process values are in different locations on the heat map. In the core idea of the two-stage autoencoder, what is important is not the location, but the uncorrelativeness between the minute fluctuations and other fluctuations, as found in the actual data. The superimposed data in Fig. 4 (b) simulated this uncorrelativeness, and are suitable for the evaluation of the two-stage autoencoder.

3.3 Model parameters

Table 4 shows the model parameters, which were all the same regardless of the target cases. The intermediate layer of the simple autoencoder was set to one layer. Since the physical relationship between the systems is small, in order to enhance the learning accuracy, we created one model for one system. There were seven models of the two-stage autoencoder and the simple autoencoder in total. One model handled hundreds of process values.

3.4 Detection thresholds, metrics, and criteria

Anomaly is detected when the threshold is continuously exceeded three times for the time-series deviation value (= the actual value – the predicted value). This is in accordance with the fact that in many cases of actual anomaly detection, an anomaly is judged when the threshold is exceeded several times in succession in order to reduce false positives due to noise and make a careful judgment. For each process value, the threshold value was set to be a multiple (n) of standard deviation (σ) of the differences between the training data and the corresponding predicted values of the simple autoencoder or the two-stage autoencoder. That is, the threshold is $\pm n\sigma$. The σ corresponds to the model accuracy. The multiple (n) was common to all process values.

Metrics were as follows. “Detection timing”: time from the anomaly occurrence to first detection. “Total false positive”: total number of false positives in the target system. The number of false positives was counted as one for the period from when the deviation exceeds the threshold to when it falls below the threshold regardless of the number of times exceeding the threshold value in a row.

We explain the evaluation criteria, except for issue (2) regarding the superiority of autoencoder itself. If the detection timing of the two-stage autoencoder is faster than that of the simple autoencoder in each case corresponding to issues (3) and (4), or if the total false positive of the two-stage autoencoder is smaller, the two-stage autoencoder is judged to have good detection performance. Since the detection timing and the total false positive depend on the threshold, as a threshold-independent evaluation method, we drew Receiver Operating Characteristic (ROC)-like “detection timing vs. total false positive” curves, and judged the performance according to the positions of the curves.

Table 1 Groups of process values.

System	Name	Number of process values
A	Neutron monitoring system	196
B	Emergency core cooling systems	267
C	Reactor systems	884
D	Containment and radiation monitoring systems	155
F	Condensate and feedwater systems	669
G	Generator systems	106
T	Turbine systems	443
Total		2720

Table 2 Target cases.

Case	State	Event	Process value of interest	System
1	Transient	Recirculation flow control hunting	Core flow rate [t/h]	C
2	Transient	TD-RFP control value failure	Reactor feedwater flow rate [t/h]	F
3	Transient	Pressure control system malfunction	Reactor pressure [MPa]	C
4	Steady, noisy condition	Feedwater line rupture	Low-pressure condensate pump (A) outlet pressure [MPa]	F
5	Steady, noisy condition	Circulating water line leak	Condenser (A) hotwell outlet conductivity [mS/m]	F
6	Steady, noisy condition	Main condenser vacuum reduction	Dehumidifier-cooler outlet flow rate [m ³ /h]	F
7	Steady, noisy condition	RSW pump switching	RSW pump (A) outlet pressure [MPa]	C

Table 3 Training and evaluation data

State	Data	Period	Sampling interval	Number of samplings	Normalized artificial noise amplitude [σ/d]
Transient	Training data *1	68.4 hours	1 minute	4106	0
	Evaluation data (normal period)	1.5 hours	1 minute	90	0
	Evaluation data (anomalous period)	1.5 hours	1 minute	90	0
Steady, noisy condition	Training data	18.5 days	1 minute	26175	0.01 *3 0.1 *4
	Evaluation data (normal period)	18.5 days	1 minute	26174	0.01 *3 0.1 *4
	Evaluation data *2 (anomalous period)	-	1 minute	-	0.01 *3 0.1 *4
Steady, operation mode change	Training data	16.3 days	1 minute	23470	0
	Evaluation data (normal period)	1.8 hours	1 minute	110	0

*1) 12 patterns of transient states with different slopes. *2) Each event has a different duration.

*3) Noisy condition. *4) Highly noisy condition.

Table 4 Model parameters.

Model		Number of nodes in intermediate layer	Time window width
Simple autoencoder	-	1/2 of number of nodes in input layer	No time window
Two-stage autoencoder	Time window autoencoder	1/2 of number of nodes in input layer	30 minutes in steady state. 10 minutes in transient states.
	Deviation autoencoder	1/10 of number of nodes in input layer	No time window

Fig. 4 Heat maps of correlation coefficients of process values in the F system. (a) Artificial noise (σ/d : 0.01), and (b) superimposed data.

4. Results and discussion

This section shows our evaluation results: the predicting ability for the operation mode changes for issue (2), detection performances in the transient state for issue (3), and those under realistic noise amplitude conditions for issue (4).

4.1 Prediction of rapid changes of the process values (Issue (2))

Figure 5 show the process values and the corresponding predicted values of the two-stage autoencoder when switching RSW pumps (Case 7). The two-stage autoencoder could predict the rectangular waves of the outlet pressures, which change rapidly and are difficult to predict, without time delay. This shows that the autoencoder has a good prediction ability for rapid changes.

Fig. 5 Time-series data of process values in the change of the operation mode and the corresponding predicted values of the two-stage autoencoder. Red and blue lines show process values and the corresponding predicted values, respectively. Dashed lines show the maximum and the minimum values of training data.

4.2 Anomaly detection in the transient state (Issue 3)

First, we explain the results qualitatively using trend graphs. Figures 6 (a) to (d) show the process values of interest in the transient state (Cases 1 - 3) and the corresponding predicted values and deviations between the process values and the predicted values. The anomaly can be detected with few false positives if there is a clear difference between the deviations in the normal period and those in the anomalous period. In Fig. 6 (a), the predicted value of the simple autoencoder almost agreed with the process value not only in the normal period but also in the anomalous period. As shown in Table 5 below, it is difficult to detect the anomaly and the anomaly cannot be detected early. Since the simple autoencoder did not learn the temporal correlation and the instantaneous values similar to the anomalous values existed in the training data, the simple autoencoder recognized the anomalous values as normal. On the other hand, in Fig. 6 (b), since the TAE in the two-stage autoencoder learned the temporal correlation, the predicted value became a value clearly different from the anomalous value at the same time as the anomaly occurs. A clear difference was apparent in the deviations between the normal period and the anomalous period. The anomaly can be detected early. Also in Figs. 6 (c) and (d), the clear differences were apparent.

Next, for quantitative explanation, Figs. 7 (a) to (c) show detection timing vs. total false positive in the transient state (Case 1 - 3). These figures were created by continuously changing n (n of $\pm n\sigma$) of all threshold values with the same n . In each figure, the higher the position of the curve is at the upper left, the better the detection performance is. It indicates that an anomaly can be detected with earlier detection timing or with a smaller total number of false positives.

In all figures, the two-stage autoencoder was at the upper left of the simple autoencoder. This indicates that the two-stage autoencoder is superior to the simple autoencoder. Table 5 shows detection timing and total false positive at the minimum false positives, which correspond to the green circles in Figs. 7 (a) to (c). In Cases 1 and 3, for the two-stage autoencoder, the anomalies were detected early within minutes of their occurrence without false positives. In Case 2, the total false positive was not zero. It is difficult to detect the anomaly with the threshold. In order to detect the anomaly, it is necessary to monitor the trend of the process value. The predicted value of the two-stage autoencoder becomes a good guideline. In conclusion, the two-stage autoencoder has good detection ability in the transient state.

Fig. 6 Time-series data of the process values of interest in transient state (Cases 1 – 3) and the corresponding predicted values and deviations of the simple autoencoder and the two-stage autoencoder. Red, blue, and green lines show process value, the corresponding predicted value, and deviation between the process value and the predicted value, respectively. Dashed lines show the maximum and the minimum values of training data.

Fig. 7 Detection timing vs. total false positive in transient state (Cases 1 - 3). Red and blue lines show the results of the two-stage autoencoder and the simple autoencoder, respectively. Green circles show detection timings at the minimum false positives.

Table 5 Detection timings and total false positive at the minimum false positives (Cases 1 - 3).

Case	Model	Threshold [σ]	Detection timing [min]	Total false positive [n/day]
1	Two-stage autoencoder	20.2	8	0.0
	Simple autoencoder	13.2	61	0.0
2	Two-stage autoencoder	9.4	40	72.9
	Simple autoencoder	24.0	3	5170.9
3	Two-stage autoencoder	51.9	3	0.0
	Simple autoencoder	192.0	7	0.0

4.3 Anomaly detection under the external disturbance (Issue (4))

First, we explain the results qualitatively using trend graphs. Figures 8 (a) to (f) show the process values of interest in the steady state under the external disturbance (Cases 4 - 6) and the corresponding predicted values and deviations between the process values and the predicted values. As the external disturbance, the artificial noise was added to the process values, as shown in Table 3. In order to make the difference in behavior between the simple autoencoder and the two-stage autoencoder easy to grasp, the results under the highly noisy condition (Figs. 8 (a) and (b)) are also shown in Case 4. Case 4 (Figs. 8 (a) to (d)) is a case where the process value slightly decreased in the anomalous period. In Figs. 8 (a) and (c), the simple autoencoder did not predict the fine vibration of the artificial noise in the normal period. The predicted values tended to approach the process values in the anomalous period. On the other hand, in Figs. 8 (b) and (d), the two-stage autoencoder predicted the fine vibration in the normal period. This indicates the effect that the DAE learned the fine vibration. The predicted values (Figs. 8 (b) and (d)) were more clearly separated from the process values in the anomalous period. The anomalous small decrease in the process value during the anomalous period is a change in the simulated data that is uncorrelated with the artificial noise. Therefore, this separation indicates the effect that the TAE learned the major fluctuations. The predicted values indicate that the DAE and the TAE worked respectively. The reason why the difference between the process value and the predicted value in the anomalous period was larger in the two-stage autoencoder than in the simple autoencoder is presumably because the change of the normal process value over time was learned in the TAE. Also in Figs. 8 (e) and (f), the clear differences were obtained.

Next, for quantitative explanation, Figs. 9 (a) to (c) show detection timing vs. total false positive under the external disturbance (Case 4 - 6). In all figures, the two-stage autoencoder was at the upper left of the simple autoencoder. This indicates that the two-stage autoencoder is superior to the simple autoencoder. Table 6 shows detection timing and total false positive at the minimum false positives, which correspond to the green circles in Figs. 9 (a) to (f). The noisy and the highly noisy conditions in Table 6 have different artificial noise amplitudes as shown in Table 3. In Table 6, in all cases, the threshold of the two-stage autoencoder is smaller than that of the simple autoencoder. Because the minute fluctuations (the artificial noise) can be predicted by the DAE, the two-stage autoencoder is less likely to cause false positives due to the minute fluctuations. This makes it possible to set a smaller threshold, resulting in faster detection timing and fewer false positives. In the noisy condition of Case 4 (Fig. 8(c)), the oscillation of the deviation immediately after the start of the anomalous period is larger than the artificial noise amplitude in the normal period, so it seems possible to set the threshold to detect the anomaly without false positives due to the minute fluctuations. However, under the detection condition of "exceeding three times in succession" described in Sec. 3.4, the oscillation of the deviation was not detected until the threshold was reduced until false positives due to the minute fluctuations occurred. In all cases, there were no false positives in the two-stage autoencoder. The anomalies were detected early within minutes of their occurrence without false positives. The two-stage autoencoder showed a good detection performance.

(a) Simple autoencoder, Case 4: Feedwater line rupture, highly noisy condition.

(b) Two-stage autoencoder, Case 4: Feedwater line rupture, highly noisy condition.

(c) Simple autoencoder, Case 4: Feedwater line rupture, noisy condition.

(d) Two-stage autoencoder, Case 4: Feedwater line rupture, noisy condition.

(e) Two-stage autoencoder, Case 5: Circulating water line leak, noisy condition

(f) Two-stage autoencoder, Case 6: Main condenser vacuum reduction, noisy condition.

Fig. 8 Time-series data of the process values of interest in external disturbance (Cases 4 – 6) and the corresponding predicted values and deviations of the simple autoencoder and the two-stage autoencoder. Red, blue, and green lines show process value, the corresponding predicted value, and deviation between the process value and the predicted value, respectively. Dashed lines show the maximum and the minimum values of training data.

(a) Case 4: Feedwater line rupture (highly noisy condition)

(b) Case 4: Feedwater line rupture

(c) Case 5: Circulating water line leak

(d) Case 6: Main condenser vacuum reduction

Fig. 9 Detection timing vs. total false positive (Cases 4 - 6). Red and blue lines show the results of the two-stage autoencoder and the simple autoencoder, respectively. Green circles show detection timings at the minimum false positives.

Table 6 Detection timings and total false positive at the minimum false positives (Cases 4 - 6).

Case	Model	Threshold [σ]	Detection timing [min]	Total false positive [n/day]
4 *1	Two-stage autoencoder	48	3	0.0
	Simple autoencoder	94	7	6.0
4 *2	Two-stage autoencoder	41	3	0.0
	Simple autoencoder	51	3	16.0
5 *2	Two-stage autoencoder	40	3	0.0
	Simple autoencoder	164	5	0.0
6 *2	Two-stage autoencoder	36	7	0.0
	Simple autoencoder	52	183	6.0

*1) highly noisy condition. *2) noisy condition.

5. Conclusion

We aim to comprehensively monitor a large number of process values and detect early signs of anomalies, including unknown events, with few false positives, without limiting the monitoring targets to specific components and operating conditions. A plant system is a complex system composed of various systems and components. In order to learn the complex changing internal state of a plant and accurately predict the process values in the normal state, there were four learning issues: (1) complex correlations of a large number of process values, (2) rapid changes of the process values, (3) changes of the process values over time, and (4) minute fluctuations of the process values. In order to address issues (1) – (4), we created a two-stage autoencoder, which is a deep learning network structure corresponding to the characteristics of the process values. With simulation data of a nuclear power plant, we assessed performances of the two-stage autoencoder: a predicting ability for rapid changes of the process values for issue (2), detection performances in the transient state for issue (3), and those under realistic noisy conditions for issue (4).

For the rapid change (issue (2)), the two-stage autoencoder predicted the rectangular waves of the outlet pressures when switching RSW pumps, which are difficult to predict, without time delay. For the transient state (issue (3)), the two-stage autoencoder predicted normal process values due to the effect of a time window autoencoder, recognized the anomalous value as not normal, and detected the anomalies early, with no false positives, except for one case. For external disturbance (issue (4)), the two-stage autoencoder predicted normal fine vibration under realistic noisy conditions due to the effect of a deviation autoencoder, recognized the anomalous small reduction in the process value as not normal due to the effect of the time window autoencoder, and detected the anomalies early, with no false positives. In conclusion, the two-stage autoencoder would be an effective solution for comprehensive plant monitoring and early detection of anomaly signs. This technology supports daily plant monitoring and health check activities by informing the notices obtained by comprehensively monitoring the relationships of thousands of process values. We hope that this technology will reduce the burden on users and enhance operation availability, maintainability, and safety.

References

- Audibert, J., Michiardi, P., Guyard, F., Marti, S. and Zuluaga, M. A., USAD : UnSupervised Anomaly Detection on Multivariate Time Series, Proceedings of 26th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2020)(2020), pp.3395–3404.
- Ayo-Imoru, R.M. and Cilliers, A.C., A survey of the state of condition-based maintenance (CBM) in the nuclear power industry, Annals of Nuclear Energy, Vol. 112 (2018), pp.177–188.
- Chalapathy, R. and Chawla, S., Deep Learning for Anomaly Detection: a Survey, arXiv:1901.03407v2 (2019).
- Hinton, G. E. and Salakhutdinov, R. R., Reducing the Dimensionality of Data with Neural Networks, Science, Vol. 313 (2006), Paper No. 5786.
- Naito, S., Taguchi, Y., Kato, Y., Nakata, K., Miyake, R., Nagura, I., Tominaga, S. and Aoki, T., A New Data Driven

Method for Monitoring a Large Number of Process Values and Detecting Anomaly Signs with a Two-stage Model Composed of a Time Window Autoencoder and a Deviation Autoencoder, Proceedings of the 2020 International Conference on Nuclear Engineering (ICONE-2020) (2020), Paper No. ICONE20-16150.

Park, D., Hoshi, Y. and Kemp, C. C., A Multimodal Anomaly Detector for Robot-Assisted Feeding Using an LSTM-Based Variational Autoencoder, IEEE Robotics and Automation Letters, Vol. 3, No.3 (2018), pp.1544–1551.

Su, J., Zhao, Y., Niu, C., Liu, R., Sun, W. and Pei, D., Robust Anomaly Detection for Multivariate Time Series through Stochastic Recurrent Neural Network, Proceedings of 25th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2019)(2019), pp.2828–2837.

Yonezawa, H., Ueda, H. and Kato, T., A Full Scope Nuclear Power Plant Simulator for Multiple Reactor Types with Virtual Control Panels, Proceedings of The 2017 International Congress on Advances in Nuclear Power Plants (ICAPP 2017) (2017), Paper No. 17142.