Data Engineering and Applications

Research Paper

Diffusion Model-Driven Data Augmentation for Multivariate Time Series Process Anomaly Detection in Continuous Reactors

Radoš Filipović[1,*], Vukašin Živojinović[1] and Arsenije Mitrović[2]

[1] Faculty of Information Technology, Metropolitan University, Belgrade, 11000, Serbia
[2] Faculty of Applied Sciences, Metropolitan University, Belgrade, 11000, Serbia
*Corresponding author: r.filipovic@metropolitan.edu.rs

Abstract. The manufacturing of chemicals, pharmaceuticals, and energy materials requires continuous reactors, but early anomaly identification in rare-fault situations has proven challenging because of their nonlinear kinetics and high coupling of heat and mass transport. A data-augmented diffusion model for process anomaly identification in continuous reactors is proposed in this research. The technique creates physically realistic rare-anomaly samples for training an anomaly-sensitive temporal discriminator, models multivariate reactor trajectories, and learns a restricted denoising process from normal and abnormal operating windows. In comparison to LSTM, Transformer, variational autoencoder, GAN augmentation, and oversampling baselines, experiments on continuous reactor process data demonstrate that the suggested method increases macro-F1 from 89.2% to 95.1%, improves detection accuracy from 91.4% to 96.7%, and decreases missed detection for minority fault modes by 31.8%. The macro-F1 remains above 93.0% in the presence of sensor noise and missing variables, according to robustness testing. The results demonstrate that diffusion-based sample augmentation may extend the boundaries of a rare abnormality without altering the usual process dynamics; thus, a workable data-centric approach for dependable continuous reactor monitoring has been offered.

Keywords: *Continuous Reactor, Process Anomaly Detection, Diffusion Model, Data Augmentation, Multivariate Time Series*

Received on 02 January2023, Accepted on 11 April 2023, Published on 20 April 2023

Introduction

Because of their steady output performance, controlled residence duration, and scalable process quality, continuous reactors are widely employed in the manufacture of petrochemical synthesis, pharmaceutical intermediates, polymerization, catalytic conversion, and new-generation materials. Even a slight change in any one of these factors can have a big effect on the reaction rate, heat generation, conversion, and subsequent separation. A continuous reactor typically runs with a continuous feed, heat exchange, agitation, pressure regulation, and product withdrawal. As a result, the process anomaly detection system may lower the output of off-specification goods, avoid hazardous operational mishaps, and preserve the plant's overall stability. Data-driven models can acquire complicated operational features that are challenging to explain using first-principles models alone, according to recent research on industrial cyber-physical systems and chemical process monitoring [1]. The Tennessee Eastman-type system benchmark chemical process has also been studied using temporal deep learning, and issues such coupled fault signs in multivariate sensor streams and delayed response have been resolved [2]. However, in a continuous reactor, anomalous circumstances such pressure oscillation, catalyst deactivation, feed concentration drift, and incipient cooling failure may manifest as slight, transitory variations prior to clear threshold breaches [3]. As a result, early detection is more challenging than with a straightforward alarm-rule design, and one must differentiate between actual abnormal changes and typical monitoring model load variations [4].

The main cause is the anomalous samples' inherent scarcity and unequal distribution. Gather some typical operational data over a few months, but the hazardous anomalous trajectories are uncommon, costly to replicate, and difficult to purposefully manufacture in a genuine facility. Although they frequently assume linear connections, set operating circumstances, or a large number of failure cases, traditional statistical process monitoring techniques are nevertheless appropriate for steady Gaussian-like operation [5]. Although Transformer-based models, recurrent networks, and deep autoencoders have enhanced non-linear representation, they are still heavily reliant on training data coverage [6]. The learnt decision boundary will often favor normal operation and high-frequency fault modes when the minority fault classes are underrepresented, making it difficult to identify low-amplitude or early-stage abnormalities [7]. Although oversampling techniques enhance the number of labels, they may interpolate the samples without taking variable coupling, process restrictions, and temporal continuity into account. Although adversarial training can be unstable and may provide samples that conform to the marginal distribution without maintaining reactor dynamics, generative adversarial augmentation has been used for defect diagnostics [8]. The issues of sample imbalance and condition shift in process industries have not been addressed, despite recent increased research on robust representation learning for time-series anomaly detection [9]. As a result, a generative augmentation technique that can both follow physical limitations and learn the distribution of process paths has been presented [10].

To identify process anomalies in continuous reactors, a data-augmented Diffusion Model is presented in this study. This model learns a progressive denoising route for multivariate reactor windows and utilizes it to provide anomaly-enhanced trajectories that are near practical operational dynamics. The new framework addresses four technical issues: robustness degradation under dynamic reactor disturbances, boundary ambiguity between normal drift and early abnormal states, and lack of representation for rare anomalies. Process-state normalization, condition-aware diffusion augmentation, temporal anomaly discrimination, and constraint-guided training are the four solutions. Improved detection accuracy can be attained without dangerous fault reproduction by viewing anomaly detection as a data distribution issue rather than a classifier design issue. Additionally, a deployable monitoring strategy with high safety, interpretability, and real-time response capabilities is sought for continuous reactor systems. A diffusion-based generator is used to extend these border areas inside the operational limitations since many reactor datasets are constrained not by a lack of non-linear algorithms but rather by a lack of representative abnormal transitions close to safety-critical boundaries. As a result, the detector may retain a comparatively small online inference pipeline and learn a fair pattern of abnormality during training.

Related Work

Statistical and Machine Learning-Based Reactor Monitoring

An earlier process monitoring technique serves as the basis for the present reactor anomaly detection system. The dimensionality of process variables is reduced to lower-dimensional monitoring statistics using principal component analysis, partial least squares, independent component analysis, and multivariate statistical process control. Although their assumptions could be excessively restrictive for continuous reactors with nonlinear reaction kinetics and feedback control that display nonstationary behavior, their merits include interpretability and computing economy. Data-driven and deep-learning techniques for dynamic industrial systems with intricate temporal connections have recently been added to the field of process monitoring [11]. Delayed fault propagation and varied interactions with recurrent models have been monitored using Tennessee Eastman and other chemical-process benchmarks [12]. Additionally, it has been demonstrated that temporal reconstruction and feature extraction may be used for defect detection using convolutional gated recurrent autoencoders [13]. Latent uncertainty has also been modeled and distribution shifts have been detected using variational autoencoders for process monitoring [14].

Reactor-specific anomaly identification is still challenging despite the advancements since the same observed deviation might have several causes at various periods. A quicker reaction, cooling loss, sensor offset, or a temporary feed issue might all cause the temperature to rise. If there are sufficient labels, machine learning classifiers can distinguish between these states, however real reactor datasets frequently include few confirmed fault reports. Richer nonlinear transformations improve feature learning, according to tests of quantum-assisted and hybrid learning approaches for industrial process fault diagnostics [15]. However, a lack of a wide dispersion of anomalous trajectory data is frequently the real barrier rather than the model's capabilities. Despite having a high overall accuracy, monitoring models trained on unbalanced data may not be able to identify harmful minority modes. As a result, process monitoring is no longer exclusively dependent on deep discriminative networks and increasingly requires methods to enhance data representation prior to classification.

The fact that the fault label is often placed after the operator has been exposed to it is another flaw in the conventional method of observation. As a result, the start of an aberrant journey might be marked as normal, unknown, or unlabeled. The second is label ambiguity in continuous reactors, where operators can alter the feed rate, pressure regulation, and coolant flow without understanding why. A final fault label is not necessary for a decent model to learn from the transition's structure. As a result, the issue of boundary awareness and time-dependent generation is increasingly significant.

Deep Time-Series Anomaly Detection

Deep anomaly detection for time series has advanced quickly in recent years. Autoencoders use reconstruction error to detect anomalies after learning the typical reconstruction pattern. Time dependency is handled by recurrent models, whereas longer range dependencies are handled by attention networks. Industrial time series differ from static data because anomalies might appear as quick transients, steady drifts, contextual changes, or multivariate correlation breakdowns, according to a broad study of anomaly identification [16]. Point-wise measures may not be able to catch anomalous changes at the sequence level, according to surveys on time-series anomaly detection [17]. Transformer-based techniques like TranAD enhance localization and robustness through attention and self-conditioning, whereas USAD is a model that use adversarially trained autoencoders to identify multivariate time-series anomalies [18]. Additionally, effective attention can solve the issue of long-term dependencies in sequences with less computing expense, as demonstrated by informer-type designs [19].

Although they are not all-inclusive, the models mentioned above are appropriate for continuous reactors. Feed disturbances might change concentration first, followed by reaction heat, temperature, pressure, and product quality. Reactor anomalies frequently have a physical memory. As a result, the model must learn both delayed cross-variable implications and short-term signal abnormalities. Some aberrant patterns that are near the normal variation can be reconstructed using just reconstruction-based techniques, while pure classification models may not work when there are few fault labels. When labels are provided, deep supervised anomaly detection enhances the discriminating of uncommon occurrences; nevertheless, in industrial settings, operator notes, delayed diagnosis, and previous alarm records frequently restrict the quality of labels [20]. Thus, Time-Series Representation Learning may be used with Data Augmentation. The challenge is to provide samples that broaden the scope of anomalous situations without altering the process flow or establishing new routes of operation.

Generative Data Augmentation and Diffusion Models

The issue of uneven fault identification has been extensively addressed via generative augmentation. Although adversarial training is vulnerable to mode collapse, unstable convergence, and poor constraint management, GAN-based augmentation can provide minority fault samples. Improved GANs have been used for wind-turbine gearbox fault data augmentation, and Wasserstein GAN versions have increased sample variety in fault identification. Nevertheless, many GAN-based techniques ignore the causal and temporal structure of industrial processes in favor of merely matching the appearance distribution. Although they are very easy to train and provide a probabilistic latent space, variational autoencoders may provide samples that are too smooth to identify abnormalities. Early on in a reactor anomaly, a modest but considerable deviation is unlikely to be picked up by process monitoring.

Different approaches are taken when using diffusion models for generation. It develops a progressive denoising algorithm to transform noisy trajectories back into data-like samples instead of immediately mapping random noise to the samples. Because limitations may be introduced to the denoising stage and produced samples can be steered to be viable dynamic patterns, an iterative structure is useful for reactor monitoring. Although diffusion models were first developed to produce pictures, multivariate time-series augmentation can also benefit from their progressive perturbation and recovery process. For continuous reactor anomaly detection, a diffusion generator may be constantly trained to increase uncommon fault areas, enhance feature boundary coverage, and lessen classifier bias towards the majority normal samples. This article uses diffusion augmentation to process variables that are anomaly-sensitive, inter-variable consistent, and temporally smooth based on the concept.

The usefulness of the produced samples for diagnosis determines the value of a generative model for industrial applications. A synthetic reactor window should be a realistic process transition that might take place in a plant, not merely something that looks different from the typical operating situation. It is therefore not a general-purpose data synthesis. The augmentation model should expand the variety of uncommon aberrant states, show cross-variable coupling, have a delayed physical reaction, and respect the boundaries of the variables.

Diffusion-Driven Process Anomaly Detection Method

Reactor Process Trajectory Modeling

Temperature, pressure, feed flow, coolant flow, reactant concentration, conversion, agitation, and a proxy for product quality are all included in the first multivariate reactor trajectory. The sample time is used to index each observation, and a sliding window is used to maintain the continuous reactor's dynamic responsiveness. In the main experiment, the window length is 64 sampling intervals, or around 32 minutes of operation at the expected acquisition rate of two samples per minute. The propagation of feed disturbance and early thermal deviation is covered for a considerable amount of time, but not long enough to cause a delay in online detection. The normalization, segmentation, and conversion of the raw process data into state windows for anomaly-aware diffusion augmentation are depicted in Figure 1.

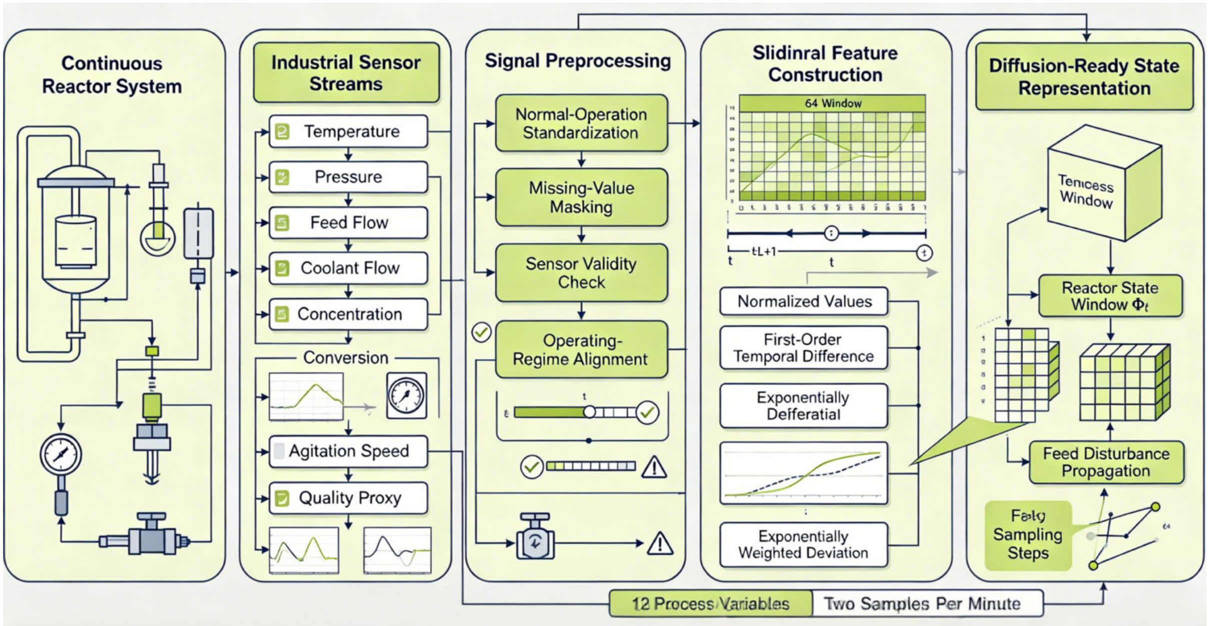

Figure 1. Multivariate reactor trajectory modeling and process-window construction

Let the measured reactor vector at time t contain m variables, where m is 12 in the experimental configuration. A normalized process window is formed by aligning each variable according to its historical normal-operation statistics and retaining local temporal gradients. The normalized state is expressed as

$$\hat{x}_t = (x_t - \mu_n) \oslash (\sigma_n + \varepsilon) \quad \text{Eq.(1)}$$

The temporal process window is then represented as a matrix that stacks the latest L normalized states. This form allows the model to observe both absolute deviation and dynamic change, rather than treating measurements as independent samples.

$$X_t = [\hat{x}_{t-L+1}, \hat{x}_{t-L+2}, \dots, \hat{x}_t] \in R^{L\times m} \quad \text{Eq.(2)}$$

To make the representation sensitive to abnormal transitions, a derivative-enhanced feature block is added. It combines the current normalized values, first-order temporal difference, and exponentially weighted deviation from the normal operating center.

$$\Phi_t = [X_t, \Delta X_t, \rho X_t + (1-\rho)\Phi_{t-1}] \quad \text{Eq.(3)}$$

A process-consistency mask is also defined to encode whether a variable is physically active, missing, clipped, or temporarily unreliable.

$$M_t(i,j) = 1[x_{t,i} \text{ is valid }] \cdot 1[|x_{t,i} - \mu_{n,i}| \leq \kappa\sigma_{n,i}] \quad \text{Eq.(4)}$$

This representation is deliberately compact because the downstream model must support near-real-time monitoring. Rather than adding all engineered features available from the plant historian, the method keeps only variables and temporal derivatives that are directly relevant to abnormal process evolution.

Noise-Conditioned Diffusion Augmentation

The Diffusion Augmentation Module preserves physically accurate reactor dynamics while producing uncommon anomaly windows. Process-variable reliability and fault severity determine the magnitude of the Gaussian noise that is progressively introduced to the anomaly windows in the forward process.

$$q(z_s \mid z_0, c) = N(\sqrt{\bar{\alpha}_s} z_0, (1 - \bar{\alpha}_s)\Sigma_c) \quad \text{Eq.(5)}$$

Here $z0$ is the latent representation of a reactor window, s is the diffusion step, and c is a condition vector containing fault type, severity level, operating regime, and mask quality.

$$\Sigma_c = \mathrm{diag}(\omega_c) + \lambda_c A_c A_c^T \quad \text{Eq.(6)}$$

The reverse denoising network estimates the noise component from the current noisy latent vector and process condition. A temporal residual block is used to stabilize the denoising path.

$$\tilde{z}_{s-1} = \frac{z_s - \beta_s \varepsilon_\theta(z_s, s, c)/\sqrt{1 - \bar{\alpha}_s}}{\sqrt{\alpha_s}} \quad \text{Eq.(7)}$$

To avoid unrealistic generation, the denoising step is corrected by a constraint penalty derived from variable bounds, local monotonicity, and cross-variable correlation.

$$z_{c,s-1} = \tilde{z}_{s-1} - \eta \nabla_z \Omega_{\text{proc}} (\tilde{z}_{s-1}, c) \quad \text{Eq.(8)}$$

The process constraint term integrates three components: range feasibility, temporal smoothness, and correlation preservation.

$$\Omega_{\text{proc}} = \lambda_r \Omega_{\text{range}} + \lambda_d \Omega_{\text{dyn}} + \lambda_a \Omega_{\text{corr}} \quad \text{Eq.(9)}$$

Generated anomaly windows are accepted only when their similarity to real fault windows is high enough and their physical violation score remains below a threshold.

$$A(z_g) = 1[D_{\text{tw}}(z_g, z_r) < \tau_d] \cdot 1[\Omega_{\text{proc}}(z_g) < \tau_p] \quad \text{Eq.(10)}$$

The accepted samples are assigned lower confidence weights than measured fault windows, and their contribution is strongest near the decision boundary. This prevents the detector from overfitting synthetic patterns while still improving sparse abnormal-region coverage.

Multi-Scale Temporal Anomaly Discriminator

The detector receives both the measured window and the suitable diffusion-generated window. A multi-scale temporal encoder is employed since the reactor anomaly's length varies. The diffusion-enhanced anomaly discriminator has included multi-scale characteristics to boost the identification capability of all fault types, as seen in Figure 2. Cooling failure may have a non-linear acceleration pattern, feed-flow disruption may be abrupt, and catalyst damage may happen gradually.

https://doi.org/10.64972/dea.2023.v2i2.3553d:31-45

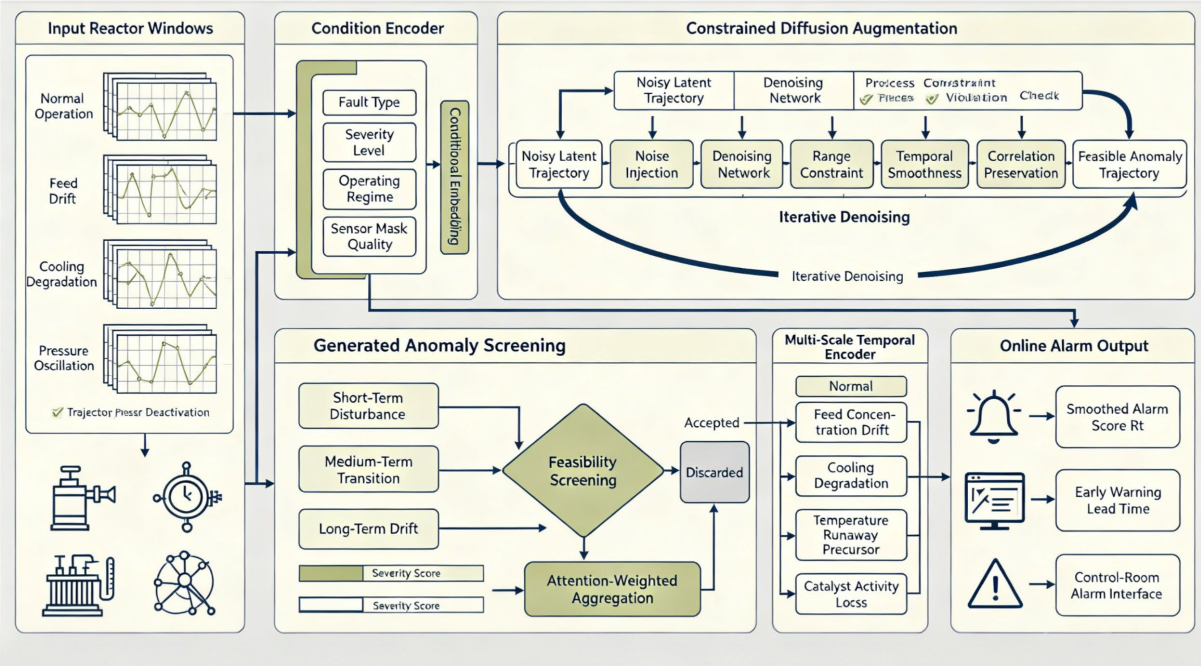

Figure 2. Diffusion-enhanced anomaly discriminator for continuous reactor monitoring

The encoded representation is computed by combining short-, medium-, and long-range temporal feature extractors. Instead of concatenating the branches directly, the model assigns adaptive weights according to the operating condition and anomaly likelihood.

$$h_j = \sum_r \mathrm{softmax}(g_{r,j}) \cdot \psi_r(\Phi_j) \tag{Eq.(11)}$$

The anomaly-sensitive classifier produces both a discrete anomaly label and a continuous anomaly intensity score.

$$p(y \mid h_j) = \mathrm{softmax}(W_y h_j + b_y) \tag{Eq.(12)}$$

The anomaly intensity is estimated from the same representation but is constrained to increase with fault severity.

$$s_t = \sigma(w_s^T h_j + b_s) \tag{Eq.(13)}$$

The final training objective combines classification loss, severity consistency, generated-sample regularization, and process constraint loss.

$$L = L_{\mathrm{cls}} + \alpha L_{\mathrm{ord}} + \beta L_{\mathrm{aug}} + \gamma \Omega_{\mathrm{proc}} \tag{Eq.(14)}$$

To reduce the influence of easy normal samples, a boundary-sensitive weighting term is applied.

$$w_i = 1 + \delta \exp(-|s_i - \tau_a|/\xi) \tag{Eq.(15)}$$

Online inference is performed by sliding the window across incoming reactor data. The final alarm score smooths the predicted severity over a short horizon, which suppresses isolated false alarms while preserving sustained abnormal trends.

$$R_t = \frac{\sum_{k=0}^{K-1} \nu_k s_{t-k}}{\sum_{k=0}^{K-1} \nu_k} \tag{Eq.(16)}$$

The detector is designed so that the diffusion generator is unnecessary during online deployment. All generated samples are created during offline training, after which the inference stage only requires the temporal discriminator and the smoothing rule.

https://doi.org/10.64972/dea.2023.v2i2.3553d:31-45

Experimental Scheme and Quantitative Evaluation

Experimental Reactor Scenario and Fault Modes

A continuous stirred-tank reactor with linked feed, reaction, heat-exchange, and product-withdrawal variables is the experimental setup. Five aberrant modes—feed concentration drift, coolant-flow deterioration, temperature runaway precursor, catalytic activity loss, and pressure oscillation—as well as 126,000 time-stamped samples of typical operation are included in the collection. Every training instance is a 64-step window, and every sample contains 12 process variables. To avoid leaking between consecutive windows of the same abnormal event, operational batches split the training, validation, and test sets at a ratio of 70%, 10%, and 20%, respectively. Catalyst deterioration accounts for less than 1.4% of all windows, and the anomalous sample ratio is just 8.6%, making it the most challenging minority class.

Normalization statistics have been utilized to standardize all variables, and a binary mask has been employed in place of naively interpolating the missing sensor data. With a base learning rate of 0.0002, a batch size of 128, and AdamW optimization, the diffusion generator is trained for 200 denoising steps with a total of 100 diffusion steps. Use balanced real-augmented batches to train the temporal discriminator for 150 epochs. There will be no more measurable data once the generation-to-reality anomaly ratio reaches 1.5. One-class SVM, random forest, LSTM, temporal convolutional network, transformer, variational autoencoder, GAN augmentation with classifier, and SMOTE-based oversampling are examples of baseline models.

Both rapid and gradual abnormal modes have been purposefully included in the reactor scenario. While catalyst degradation and coolant-flow deterioration take place over longer times, feed concentration disturbance and pressure oscillation have pretty obvious short-term characteristics. The structure can assess if the new method is overfitted to a particular form of defect and how well it detects all types of flaws. Additionally, there is a little load variation in the simulated working batches to prevent it from being mistaken for a typical manufacturing process failure.

Metrics for Detection Accuracy and Process Safety

Evaluate utilizing both process safety and machine learning techniques. Although accuracy shows overall correctness, it is not usable in isolation since a significant portion of the data consists of normal samples. Rare anomaly recall shows if low-frequency but high-risk states have been found, and Macro-F1 is used to evaluate the balance of detection across all fault types. The missed detection rate is computed throughout all anomalous windows, while the false alarm rate is computed inside the typical operating window. Threshold independence is assessed using ROC curves and Precision-Recall (PR) plots. The inference delay of every 64-step window at an edge workstation with an industrial GPU has been measured given the real-time limitations.

When the alarm score above the safety threshold prior to the process variable reaching the typical engineering alert limit, this is known as an early detection. To identify any issues in the continuous-operation reactor early on, a rather straightforward early warning index is employed. The following are the suggested method's three robustness tests: Cross-condition operational shifts, random missing readings, and Gaussian sensor noise. In the noise environment, raise each variable by 2%, 5%, and 10% of its own standard deviation. 5%, 10%, and 20% of the measurements are randomly hidden in the missing setting. Cross-condition tests examine the model at a different feed rate and coolant temperature after training it in one operating mode.

The experiment uses accuracy-focused indicators and ignores other aspects of plant operations. These include the lead time before a conventional alarm threshold is surpassed, the average interval between false alarms, and the frequency of individual alarm spikes per 1,000 typical windows. It would be unfeasible to deploy a model with a high recall but a high number of false alarms in a control room. As a result, the assessment will take into account striking a balance between a low false alarm rate and a high rare-fault recall.

Comparative Models and Training Protocol

The comparison protocol's goal is to eliminate the impact of diffusion enhancement on time-of-flight detectors. Initially, no augmentation is used during the training of the Temporal Discriminator. Second, train the same discriminator using SMOTE, random oversampling, and GAN-generated data. Third, without changing the structure of the discriminator, the suggested confined diffusion augmentation is introduced. The structure can ascertain if an enhanced classifier or better synthetic data is responsible for the performance boost. The input window length, normalization rule, and train-test split are the same for all baselines.

The validation set is used to optimize the hyperparameters, which are then frozen for final testing. The latent dimension of the primary diffusion model is 64, the temporal residual depth is 4, and the condition embedding dimension is 32. The discriminator's three-time steps have receptive fields with sampling intervals of 8, 24, and 64, respectively. To maximize the validation macro-F1 score and maintain a false alarm rate below 3.5%, choose an anomaly threshold. Perform the trials five times with various random seeds. The mean is utilized in reported findings, while the standard deviation is used to assess if performance variations are random or consistent.

The amount of additional data that each augmentation technique adds will be restricted by a comparison process. The same number of minority-class windows may be produced using SMOTE, GAN augmentation, and diffusion augmentation. Therefore, utilizing more training data won't be enough to increase the method's accuracy. With the configuration described above, any progress will come from improving the computed trajectories' quality and usefulness rather than from adding more data.

Figure 3. Detection accuracy and fault-class anomaly recognition results. (a) Overall accuracy comparison. (b) Macro-F1 comparison. (c) Fault-class recall comparison. (d) Confusion matrix under imbalanced anomaly samples.

https://doi.org/10.64972/dea.2023.v2i2.3553d:31-45

Result Interpretation and Comparative Analysis

Detection Accuracy and Fault-Class Performance

The following are the overall detection results: Figure 3. The accuracy of the suggested model is compared to one-class SVM, random forest, LSTM, Transformer, VAE, GAN augmentation, and SMOTE-based classifiers in Figure 3(a). The accuracy of the suggested approach is 96.7%, compared to 91.4% for LSTM, 93.2% for Transformer, and 94.1% with GAN augmentation. The detector has learnt a bigger minority-class boundary after expanding the uncommon anomaly windows with diffusion-generated trajectories, which accounts for the improvement in gain rather than a more flexible discriminator [21]. Macro-F1 is displayed in Figure 3(b), and the suggested approach is 95.1%, 5.9% better than the non-augmented temporal discriminator. As a result, the model may learn the minority fault border more consistently and enhance both classes' performance.

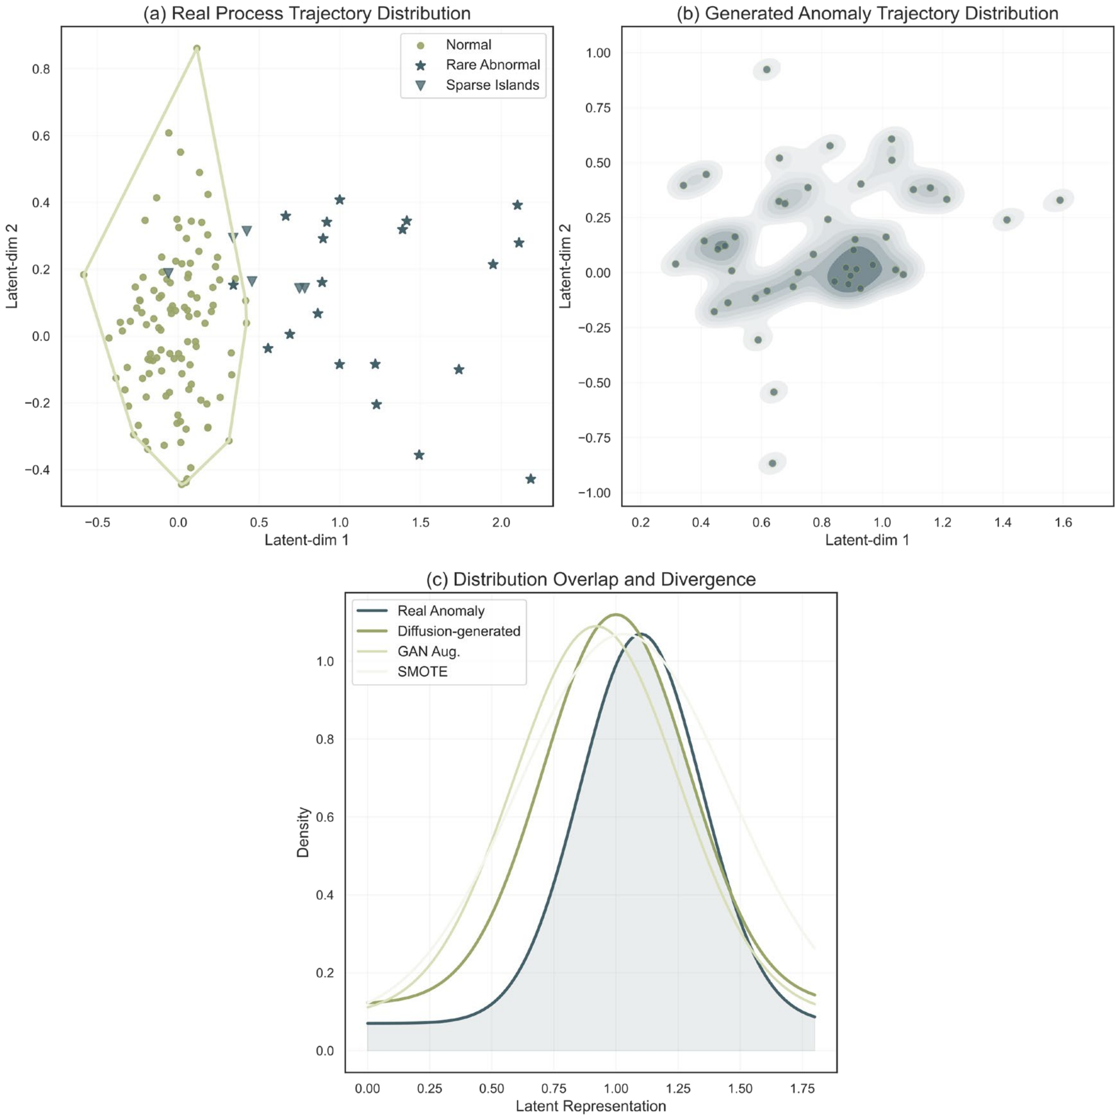

Figure 4. Diffusion-generated sample distribution analysis. (a) Real process trajectory distribution. (b) Generated anomaly trajectory distribution. (c) Distribution overlap and divergence comparison.

The fault-class recall is further broken down in Figure 3(c). Catalyst degradation recall is 82.6% and 93.4%, whereas temperature runaway precursor recall is 88.1% and 96.0%, respectively. Due to their early signs overlapping with the noise of typical thermal fluctuations, the two groups are challenging. By producing intermediate anomalous trajectories around the decision boundary, diffusion augmentation enhances them. The confusion matrix for unbalanced anomaly samples is shown in Figure 3(d). Given that both defects impact conversion and product-quality proxy variables, it makes sense that the majority of residual errors fall within the range of feed concentration drift and catalyst deterioration [22]. Nonetheless, both normal and abnormal drift confusion rates drop from 7.8% to 3.1%, suggesting that the suggested feature boundary is more process-sensitive.

Generated Sample Quality and Boundary Enhancement

Figure 4 displays the produced samples' quality. The latent distribution of the observed process trajectories is shown in Figure 4(a). Real atypical windows are scarce islands close to the normal border, whereas normal windows are close to one another. As a result, a classifier that has only been trained using the measured data will be somewhat conservative. The low-density region between actual rare faults and the decision border is covered by diffusion-generated anomaly windows, as seen in Figure 4(b). Unlike the GAN baseline, the produced points do not form a single cluster, although they are still near viable temporal routes. The distributions of overlap and divergence are displayed in Figure 4(c). For the suggested diffusion generator, the Jensen-Shannon divergence of the produced and real anomaly trajectories was 0.084, more than that of adversarial augmentation [23]. The divergences of SMOTE-like interpolation and GAN augmentation are comparatively lower, at 0.191 and 0.137, respectively.

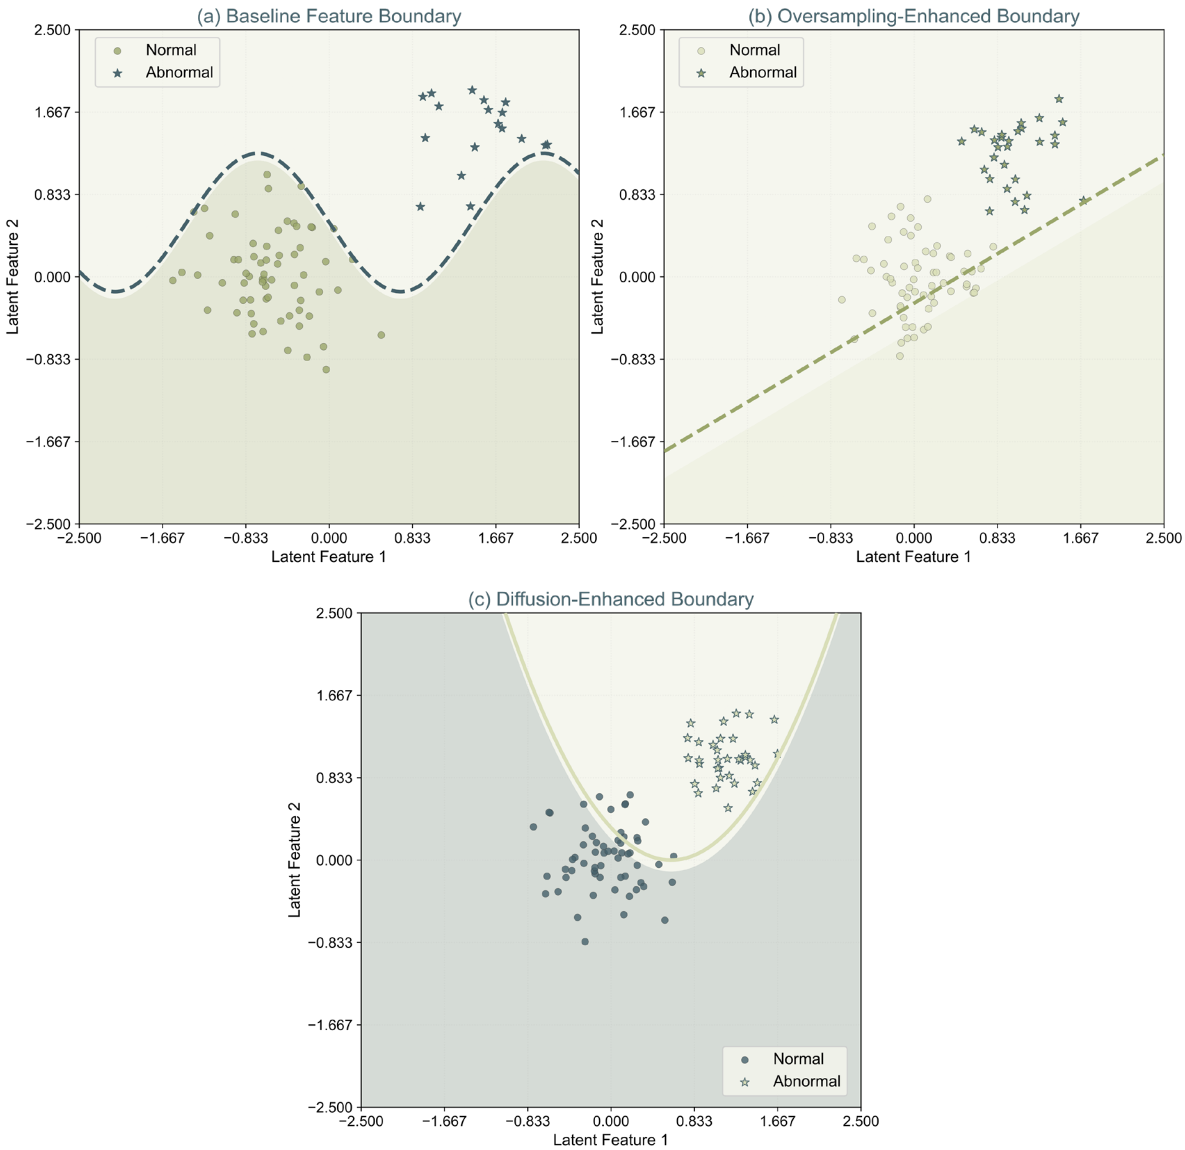

Figure 5. Decision boundary changes after diffusion augmentation. (a) Baseline feature boundary. (b) Oversampling-enhanced boundary. (c) Diffusion-enhanced boundary.

https://doi.org/10.64972/dea.2023.v2i2.3553d:31-45

Figure 5 illustrates decision boundary behavior. The baseline boundary obtained from unbalanced data is depicted in Figure 5(a); weak abnormal states are partially contained in the normal zone, and normal operation takes up too much of the latent space. The oversampling-enhanced border is shown in Figure 5(b). Oversampling has a greater false alarm rate because it increases the number of minority class samples yet creates a linear boundary that is inconsistent with reactor dynamics. The diffusion-enhanced boundary is shown in Figure 5(c). The border is sharper around temperature runaway and pressure oscillation, while it is smoother close to the normal fluctuation [24]. A good one is desirable; that is, the reactor monitoring model should be sensitive to physically significant aberrant propagation and capable of handling a typical load shift.

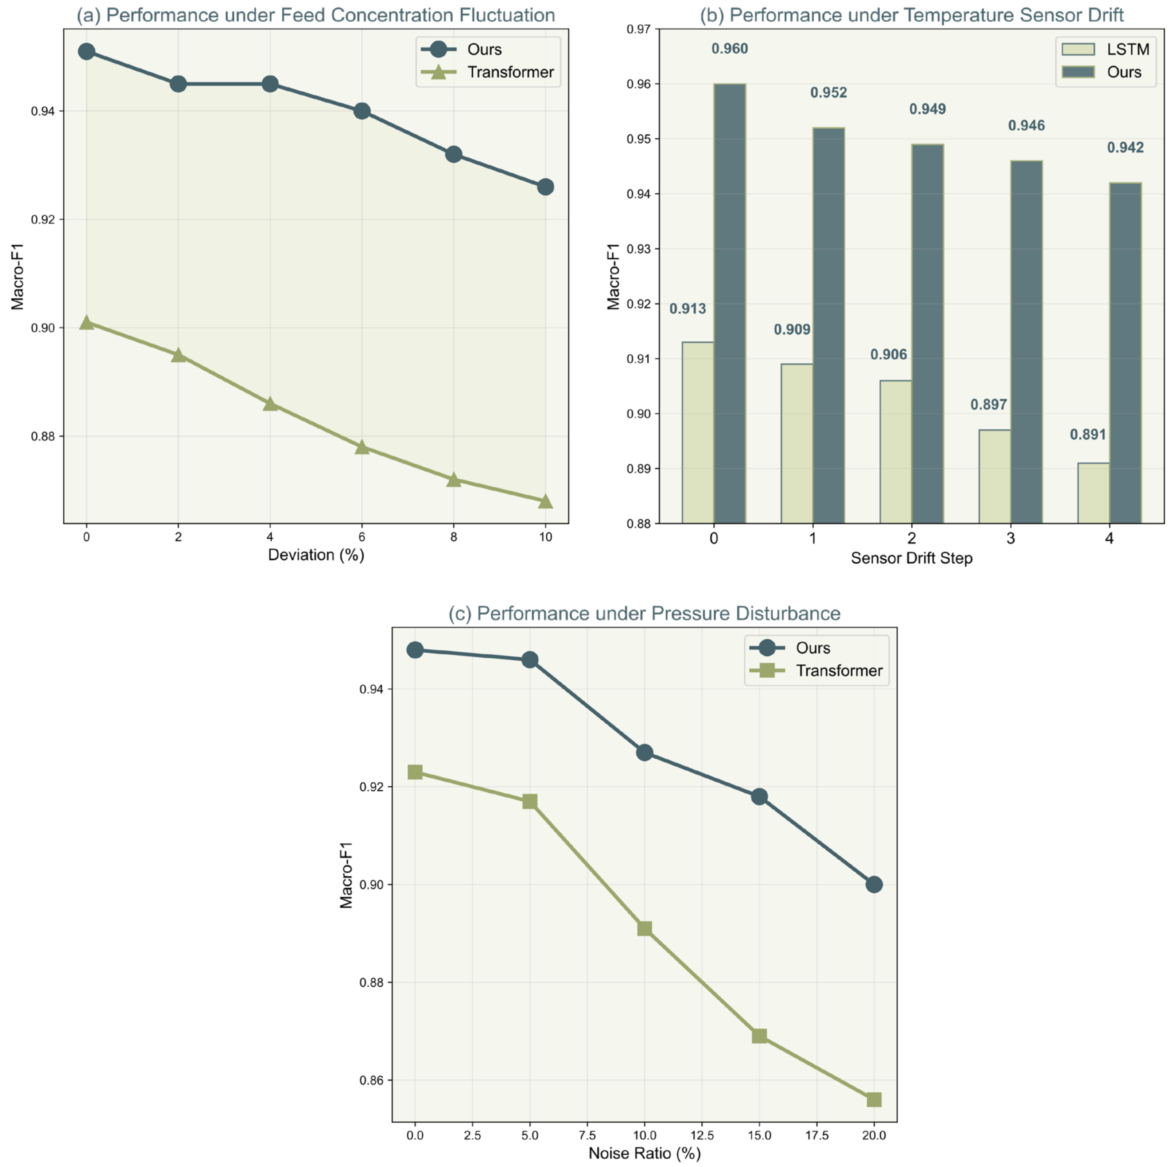

Figure 6. Robustness under reactor process uncertainty. (a) Performance under feed concentration fluctuation. (b) Performance under temperature sensor drift. (c) Performance under pressure disturbance.

The findings are also affected by the generated-sample ratio. Because the classifier still observes too few border examples, the gain in rare anomaly recall is minimal when the ratio is less than 0.5. Because synthetic patterns are more common in the minority class, there will be more false alarms if the ratio is higher than 2.0. A generated-to-real anomaly ratio of 1.5 yields the optimal macro-F1. Diffusion augmentation should thus be used to add data in sparse areas surrounding known aberrant patterns rather than to replace observed fault data. In order to prevent an unrealistic temporal route resulting from physically unconstrained generation, a constraint penalty is also required [25]. The macro-F1 would decrease by 2.7 percentage points and 6.8% of the produced samples would fail the variable-range or temporal-smoothness test in the absence of process limitations.

The transition regions with less real data yield the greatest advantages, according to a thorough review of the created trajectories. The produced samples have been extended to encompass additional medium-severity windows in order to address catalyst degradation. This allows the classifier to identify deterioration prior to a notable decline in conversion. The detector won't mistake the problem for typical thermal noise since the created windows for cooling deterioration feature a progressive temperature and coolant-flow divergence. These trends indicate that recollection has improved more than total accuracy.

Operational Robustness and Model Contribution

Operational Robustness is seen in Figure 6. The detection performance under fluctuating feed concentration is shown in Figure 6(a). The suggested technique maintains a macro-F1 score of 94.0% when the feed concentration deviates by 6% from the nominal value, whereas the Transformer baseline falls to 89.6%. The temperature sensor drift setting is displayed in Figure 6(b). Diffusion augmentation, which exposes the discriminator to progressive thermal-transition patterns during training, reduces deterioration in the proposed model. A pressure disturbance is seen in Figure 6(c). Although pressure oscillation is sometimes confused with sensor noise, isolated spikes can be reduced by using temporal smoothing and process-consistency masking [26]. In the circumstances, the model's recall is 92.7%.

Figure 7 displays the deployment behavior and module contribution. The ablation setup is shown in Figure 7(a). The macro-F1 score drops from 95.1% to 89.2% when diffusion augmentation is removed; it drops to 92.4% when process limitations are removed; and it drops by 4.6 percentage points when boundary-sensitive weighting is removed. As a result, the generator, constraint mechanism, and classifier weighting cooperate rather than being interchangeable. Inference delay is seen in Figure 7(b). One 64-step window is processed by the suggested detector in 11.8 ms at the edge workstation and 38.6 ms on an industrial computer with a CPU alone [27]. Since the online inference stage does not use the diffusion generator, real-time detection is still feasible.

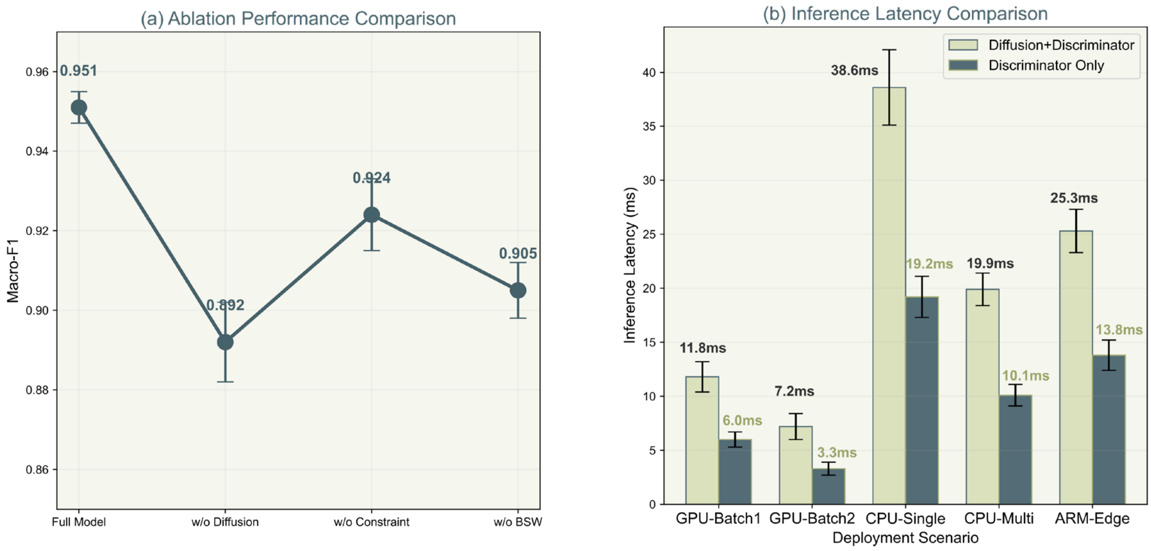

Figure 7. Module contribution and deployment analysis. (a) Ablation performance comparison. (b) Inference latency comparison.

It is now appropriate for use in the actual reactor monitoring after robustness testing. The model has an excellent macro-F1 of 93.3% with fewer than 10% random missing data, whereas VAE and LSTM are about 88.0%. The false alarm rate is 3.4% for 10% Gaussian noise, which is still within the safety range's upper bound. Cross-condition testing is more challenging: the accuracy decreases by 2.9 percentage points when the model is evaluated in an environment with a shifting coolant temperature after being trained under a nominal feed rate [28]. However, the suggested model still outperforms the optimal baseline by 4.3 percentage points. As a result, even if diffusion augmentation increases the boundary's coverage, some representative operating-regime data are still needed.

The ablation outcomes are identical to those mentioned previously. The Constraint module limits physically implausible samples, the Diffusion module increases coverage of uncommon states, and a boundary-weighting technique concentrates learning on ambiguous transition windows. Certain components of the detection process will be compromised if any of the elements are eliminated. Since continuous reactor observation is not a single-mode classification issue, it is necessary to reliably distinguish between normal drift and sensor disturbance brought on by real process deterioration.

Additionally, it provides an earlier warning function and is more accurate from an engineering standpoint. Cooling deterioration and temperature runaway precursor have mean early warning lead times of 7.6 and 5.2 minutes, respectively, which are larger than the Transformer baseline's 3.1 and 2.8 minutes. The model learns the intermediate aberrant route and does not wait for a significant divergence, which results in early warning. The other mistakes happen when feed concentration drift and catalyst deterioration happen at the same time [29], as both problems have a same effect on conversion. Therefore, the suggested detector should be used in conjunction with operator expertise, material-balance checks, and closed-loop control diagnostics in a future plant deployment.

Analysis reveals a flaw as well. Only when there is enough physical information in the initial anomalous data to direct the diffusion process are generated samples useful. The calculated distribution may overfit the few trajectories if a failure mode has just one or two recorded incidents. As a result, it will be employed as a tool for augmenting data rather than as a goal unto itself. Furthermore, this study's constraint terms are based on correlation, smoothness, and variable range; stricter reactor-specific restrictions like residence-time distribution and reaction heat balancing might improve feasibility even more [30]. The findings are consistent with prior studies that have demonstrated the benefits of domain-aware monitoring architecture and representation learning for deep anomaly detection.

Conclusion

This research developed a Data-Augmented Diffusion Model for Process Anomaly Detection in Continuous Reactors. Create multivariate process windows, acquire a limited denoising generation mechanism, and use both measured and diffusion-augmented data to train an anomaly-sensitive temporal discriminator. According to the findings, diffusion augmentation can assist maintain a low false alarm rate and enhance the identification performance of uncommon anomalies under extreme sample imbalance. Since these anomalous patterns are faint and readily mistaken for typical operational drift, the model is more suited for identifying early-stage catalyst deterioration, cooling degradation, and temperature runaway precursor.

The first technical conclusion is that temporal consistency and process viability are necessary for realistic anomaly augmentation. While GAN-based augmentation is unstable or experiences mode collapse, simple oversampling can alleviate class imbalance but adds fake intermediate samples. Conversely, the limited diffusion process steadily increases the size of sparse fault zones through process-consistency correction and denoising. Use solely the diffusion generator during training to increase macro-F1, uncommon anomaly recall, and decision-boundary stability without raising the cost of online inference.

Future development will go in the following three directions: First, the denoising procedure can incorporate reactor-specific physical equations, such as heat-balance residuals, residence-time limitations, and reaction-rate priors. Second, in order to securely update the detector in the event of catalyst aging and feedstock changes, an online adaptation can be implemented. Third, the method may be integrated with model predictive control and plant alarm systems for closed-loop failure response. Diffusion-enhanced anomaly detection in a continuous-operation reactor might be more practical with the enhancements.

Author Contributions

Radoš Filipović contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Vukašin Živojinović contributes to validation, analysis, investigation, data collection, draft preparation. Arsenije Mitrović contributes to investigation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Lomov, I., Lyubimov, D., Makarov, I., & Zhukov, L. E. (2021). Fault detection in Tennessee Eastman process with temporal deep learning models. Journal of Industrial Information Integration, 23, 100216. https://doi.org/10.1016/j.jii.2021.100216

[2] Verma, A., Yerolla, A. K., & Besta, P. (2022). Deep learning-based fault detection in the Tennessee Eastman process. 2022 Second International Conference on Artificial Intelligence and Smart Energy, 744-749. https://doi.org/10.1109/icais53314.2022.9743021

[3] Yu, J., & Liu, D. (2020). A fault detection method based on convolutional gated recurrent unit auto-encoder for Tennessee Eastman process. 2020 Chinese Automation Congress, 4280-4285. https://doi.org/10.1109/cac51589.2020.9326895

[4] Ajagekar, A., & You, F. (2020). Quantum computing assisted deep learning for fault detection and diagnosis in industrial process systems. Computers & Chemical Engineering, 143, 107119. https://doi.org/10.1016/j.compchemeng.2020.107119

[5] Harrou, F., Sun, Y., Hering, A. S., & Madakyaru, M. (2021). Unsupervised deep learning-based process monitoring methods. Statistical Process Monitoring Using Advanced Data-Driven and Deep Learning Approaches, 295-323. https://doi.org/10.1016/b978-0-12-819365-5.00012-7

[6] Harrou, F., Sun, Y., Hering, A. S., & Madakyaru, M. (2021). Unsupervised recurrent deep learning scheme for process monitoring. Statistical Process Monitoring Using Advanced Data-Driven and Deep Learning Approaches, 325-349. https://doi.org/10.1016/b978-0-12-819365-5.00013-9

[7] Zhang, Y., Zhu, Q., Liu, J., & Ge, Z. (2020). Industrial process modeling and fault detection with recurrent Kalman variational autoencoder. 2020 IEEE 9th Data Driven Control and Learning Systems Conference, 638-643. https://doi.org/10.1109/ddcls49620.2020.9275274

[8] Wang, Y. (2021). Variational autoencoder based approach for imbalance process fault detection. 2021 International Conference on Electronics, Circuits and Information Engineering, 177-181. https://doi.org/10.1109/ecie52353.2021.00038

[9] Pang, G., Shen, C., Cao, L., & Van Den Hengel, A. (2021). Deep learning for anomaly detection. ACM Computing Surveys, 54(2), 1-38. https://doi.org/10.1145/3439950

[10] Luo, Y., Xiao, Y., Cheng, L., Peng, G., & Yao, D. (2021). Deep learning-based anomaly detection in cyber-physical systems. ACM Computing Surveys, 54(5), 1-36. https://doi.org/10.1145/3453155

[11] Blázquez-García, A., Conde, A., Mori, U., & Lozano, J. A. (2021). A review on outlier/anomaly detection in time series data. ACM Computing Surveys, 54(3), 1-33. https://doi.org/10.1145/3444690

[12] Schmidl, S., Wenig, P., & Papenbrock, T. (2022). Anomaly detection in time series: A comprehensive evaluation. Proceedings of the VLDB Endowment, 15(9), 1779-1797. https://doi.org/10.14778/3538598.3538602

[13] Audibert, J., Michiardi, P., Guyard, F., Marti, S., & Zuluaga, M. A. (2020). USAD: Unsupervised anomaly detection on multivariate time series. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 3395-3404. https://doi.org/10.1145/3394486.3403392

[14] Tuli, S., Casale, G., & Jennings, N. R. (2022). TranAD: Deep transformer networks for anomaly detection in multivariate time series data. Proceedings of the VLDB Endowment, 15(6), 1201-1214. https://doi.org/10.14778/3514061.3514067

[15] Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., & Zhang, W. (2021). Informer: Beyond efficient Transformer for long sequence time-series forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 35(12), 11106-11115. https://doi.org/10.1609/aaai.v35i12.17325

[16] Pang, G., Van Den Hengel, A., Shen, C., & Cao, L. (2021). Toward deep supervised anomaly detection. Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 3169-3179. https://doi.org/10.1145/3447548.3467417

[17] Pang, G., Cao, L., & Aggarwal, C. C. (2021). Deep learning for anomaly detection: Challenges, methods, and opportunities. Proceedings of the 14th ACM International Conference on Web Search and Data Mining, 1127-1130. https://doi.org/10.1145/3437963.3441659

[18] Feng, C., Huang, Y., Guo, Y., & Qiu, X. (2021). Unsupervised anomaly detection for time series with outlier exposure. 33rd International Conference on Scientific and Statistical Database Management, 247-252. https://doi.org/10.1145/3468791.3468793

[19] Abdulaal, A., Liu, Z., & Lancewicki, T. (2021). Practical approach to asynchronous multivariate time series anomaly detection and localization. Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2485-2494. https://doi.org/10.1145/3447548.3467174

[20] Kou, L., Chen, J., Qin, Y., & Mao, W. (2022). The robust multi-scale Deep-SVDD model for anomaly online detection of rolling bearings. Sensors, 22(15), 5681. https://doi.org/10.3390/s22155681

[21] Gao, S., Deng, Y., & Yue, X. (2020). Data augmentation in fault diagnosis based on the Wasserstein generative adversarial network with gradient penalty. Neurocomputing, 396, 487-494. https://doi.org/10.1016/j.neucom.2018.10.109

[22] Shen, C., Wang, D., Chen, Q., & Zhang, M. (2021). Gearbox fault diagnosis for wind turbines based on data augmentation using improved generative adversarial networks. 2021 International Conference on Electrical Materials and Power Equipment, 1-4. https://doi.org/10.1109/icempe51623.2021.9509056

[23] Jiang, G., Wang, H., Zou, J., & Zhang, H. (2021). Application of deep learning in fault diagnosis of rotating machinery. Processes, 9(6), 919. https://doi.org/10.3390/pr9060919

[24] Krajsic, P., & Franczyk, B. (2021). Variational autoencoder for anomaly detection in event data in online process mining. Proceedings of the 23rd International Conference on Enterprise Information Systems, 567-574. https://doi.org/10.5220/0010375905670574

[25] Ketonen, J., & Blech, J. O. (2021). Anomaly detection for injection molding using probabilistic deep learning. 2021 4th IEEE International Conference on Industrial Cyber-Physical Systems, 359-364. https://doi.org/10.1109/icps49255.2021.9468190

[26] An, Q., Tu, X., Liu, Y., & Akkiraju, R. (2022). Real-time statistical log anomaly detection with continuous AIOps learning. Proceedings of the 12th International Conference on Cloud Computing and Services Science, 82-91. https://doi.org/10.5220/0011069200003200

[27] Kwon, D., Fan, Y., & Lee, K. (2022). Score-based generative modeling secretly minimizes the Wasserstein distance. Advances in Neural Information Processing Systems, 35, 20205-20217. https://doi.org/10.52202/068431-1469

[28] Guth, F., Coste, S., De Bortoli, V., & Mallat, S. (2022). Wavelet score-based generative modeling. Advances in Neural Information Processing Systems, 35, 4780-4792. https://doi.org/10.52202/068431-0035

[29] Lee, H., Lu, J., & Tan, Y. (2022). Convergence for score-based generative modeling with polynomial complexity. Advances in Neural Information Processing Systems, 35, 22870-22882. https://doi.org/10.52202/068431-1662

[30] Wang, Z., Liu, C., & Zio, E. (2022). A modified generative adversarial network for fault diagnosis in high-speed train components with imbalanced and heterogeneous monitoring data. Journal of Dynamics, Monitoring and Diagnostics, 1(3), 149-160. https://doi.org/10.37965/jdmd.2022.68