Explainable Temporal Attribution for Time-Series-Based Quality Prediction in Injection Molding

Yulin Wang[1], Elke Deckers[2], Konstantinos Gryllias[3]

[1,2,3] *Mecha(tro)nic System Dynamics (LMSD), Department of Mechanical Engineering, KU Leuven, 3001 Leuven, Belgium*

[1,2,3] *Flanders Make@KU Leuven, Belgium*
yulin.wang@kuleuven.be
elke.deckers@kuleuven.be
konstantinos.gryllias@kuleuven.be

Abstract

Time-series sensor signals collected during injection molding contain rich information about process dynamics and final part quality. However, quality prediction models based on high-dimensional temporal signals are often difficult to interpret, and direct attribution in the original signal domain may produce unstable explanations. This paper proposes an explainable temporal attribution framework for time-series-based quality prediction in injection molding. The framework first learns compact latent representations of multivariate process signals using a one-dimensional convolutional autoencoder, and then trains a latent-space prediction model for defect classification or dimensional quality regression. To recover process-level interpretability, SHapley Additive exPlanations (SHAP) are computed in the latent space and propagated back to the original temporal signal domain through the decoder Jacobian. The proposed framework is evaluated on two injection molding datasets. The results show that the learned latent representations preserve dominant temporal process characteristics and provide useful information for downstream quality prediction. The temporal attribution results reveal that pressure-related signals and the filling-to-packing transition are especially important for quality prediction, while phase-aware analysis indicates that packing and cooling stages play dominant roles in sink mark prediction. Robustness analysis further demonstrates that the decoder-based temporal attribution framework provides more stable and consistent explanations than direct SHAP analysis in the original signal domain. The proposed method offers an interpretable way to connect latent predictive features with physically meaningful molding process dynamics, supporting process understanding and quality-related decision making.

1. Introduction

Injection molding is one of the most widely used manufacturing processes for the production of polymer components because of its high productivity, low unit cost, and capability for mass production. However, the quality of molded parts is strongly affected by the transient process condition, including melt flow, cavity pressure evolution, and thermal dynamics during filling, packing, and cooling stages. Variations in these process dynamics may lead to defects and mechanical performance degradation, even when the machine setting parameters remain unchanged. In recent years, the increasing availability of in-mold sensors and machine monitoring systems has enabled the collection of high-frequency process data, which provides new opportunities for data-driven quality prediction.

To make use of these signals, many studies have applied machine learning and deep learning methods to predict part quality directly from process parameters. In particular, time-series-based models have shown strong capability in capturing nonlinear process behavior and extracting hidden patterns from cavity pressure and temperature signals. For example, Gordon et al. developed prediction models for part thickness, dimensions, weight, and tensile strength based on real-time data obtained from multiple pressure and temperature sensors (Gordon, Kazmer, Tang, Fan, & Gao, 2015). Finkeldey et al. combined real and simulated time-series sensor data and employed an extreme gradient boosting (XGBoost) model to predict part thickness and weight using only a limited amount of experimental data (Finkeldey, Volke, Zarges, Heim, & Wiederkehr, 2020). Deng et al. developed a multi-feature fusion prediction model, which indicates that the introduction of time-series data significantly improved the model performance (Deng, Xiang, Lin, Zheng, & Yang, 2025). These methods can often achieve satisfactory prediction accuracy. However, most existing studies focus mainly on improving predictive performance, while limited attention has been paid to understanding which temporal process behaviors contribute to the final prediction. For industrial applications, prediction accuracy alone is often insufficient. Process engineers are usually more concerned with identifying the origin of quality variation, determining which processing stage is responsible for abnormal behavior, and understanding how process conditions should be adjusted.

Explainability methods such as SHapley Additive exPlanations (SHAP), Integrated Gradients, and saliency-based approaches have been increasingly adopted to interpret machine learning models. In injection molding, explainability methods are often applied to prediction models with statistical feature inputs, such as maximum injection pressure, maximum holding pressure, cooling time, and extracted features from time-series data (Hong et al., 2025; Muaz et al., 2025; Presciuttini et al., 2025). Although these methods can estimate feature importance, their application to industrial raw time-series data remains challenging.

Injection molding transient process condition signals are typically high-dimensional, highly correlated in time, and affected by measurement noise. As a result, direct attribution on raw signals often produces unstable or fragmented explanations that are difficult to associate with physically meaningful process behavior. In many cases, the attribution results highlight isolated time points without providing a clear interpretation of the underlying process dynamics. Recent studies on time-series explainability have also pointed out that existing saliency methods may suffer from inconsistency and limited interpretability when applied to sequential industrial data (Kechris, Dan, & Atienza, 2025). TimeSHAP extends feature attribution to event-level explanations by perturbing sequential inputs (Bento, Saleiro, Cruz, Figueiredo, & Bizarro, 2021), but it still operates in the original input domain and may struggle to capture high-level latent process dynamics in complex industrial systems.

Representation learning provides a potential solution to this problem. Recent studies have increasingly explored representation learning methods for injection molding process monitoring and quality prediction, particularly through autoencoder-based architectures and latent feature extraction techniques. Existing works mainly aim to encode high-dimensional cavity pressure or in-mold sensor signals into compact latent representations that can improve prediction robustness and reduce the dependency on manually designed features. Ke et al. employed an autoencoder to encode cavity pressure curves into latent features for multi-quality prediction and reported that locally encoded pressure representations achieved higher prediction stability than conventional global pressure features (Ke, Wang, & Nian, 2024). More recently, Yan et al. proposed a variational autoencoder framework for automated process monitoring in injection molding, where pressure signals were projected into a low-dimensional latent space for anomaly detection and machine setting parameter optimization (Yan et al., 2024). Wang et al. encoded multi-location in-mold pressure signals using autoencoders to capture local flow dynamics for multi-quality prediction, demonstrating that latent representations could improve robustness and stability in data-driven quality modeling (J.-C. Wang, Chang, & Ke, 2026). Ding et al. introduced a GRU-attention autoencoder framework for soft sensing in injection molding and showed that temporal latent features extracted from in-mold sensing signals significantly improved part weight prediction performance compared with conventional handcrafted features (Ding, Ali, Gao, Zhang, & Gao, 2025). Although these studies demonstrate that representation learning can effectively capture injection molding process dynamics and improve prediction accuracy, most existing approaches mainly treat latent representations as compressed numerical features for downstream prediction tasks. Limited attention has been paid to interpreting how prediction-relevant latent variables correspond to physically meaningful temporal process behavior during filling, packing, and cooling stages. Consequently, the learned latent space often remains difficult to interpret from a manufacturing perspective, which limits its usefulness for process understanding, root-cause analysis, and process optimization.

To address these limitations, this paper proposes an explainable temporal attribution framework for time-series-based quality prediction in injection molding. The proposed method first learns latent representations of process signals through an encoder-decoder structure and performs quality prediction in the latent space. Prediction-relevant latent attributions are then projected back to the temporal signal domain through the local decoder Jacobian, which enables the interpretation of influential process dynamics in the original time series. In this way, the proposed framework connects latent representation learning with process-aware temporal interpretation and provides explanations that can be directly associated with different processing stages.

The main contributions of this work are summarized as follows:

1. A latent-space explainable quality prediction framework is proposed for injection molding time-series data, integrating convolutional autoencoder-based representation learning with downstream classification and regression prediction tasks.
2. A decoder-Jacobian-based temporal attribution strategy is developed to project prediction-relevant latent attributions back to the original multivariate time-series domain, enabling the identification of influential sensor channels and critical temporal regions.
3. The proposed framework is evaluated on two injection molding datasets involving sink mark classification and dimensional quality regression, demonstrating its capability to generate process-aware explanations associated

with filling, packing, and cooling stages.

4. A robustness analysis is conducted against direct SHAP attribution in the original signal domain, showing that the decoder-based attribution mapping provides more stable and consistent temporal explanations under noise perturbations.

2. Materials and Methods

In this study, a temporal attribution framework is developed to analyze the relationship between process dynamics and product quality in injection molding. The proposed framework consists of three main components, namely latent representation learning, latent-space quality prediction, and decoder-based temporal attribution mapping. First, high-frequency sensor signals are encoded into a compact latent representation through an encoder-decoder architecture, where the latent space is expected to capture the underlying process behavior during the molding cycle. A prediction model is then trained in the latent space to estimate the target quality variable. To enable process-aware interpretation, latent attributions are subsequently projected back to the temporal signal domain through the local decoder Jacobian, thereby establishing the relationship between prediction-relevant latent variables and temporal process dynamics. The overall workflow of the proposed framework is illustrated in Figure 1. This section systematically describes the main components of the framework and introduces injection molding datasets that are used for the evaluation and validation of the proposed methodology.

2.1. Injection molding dataset

In-mold pressure and temperature signals were selected as the primary process signatures because they directly reflect the thermo-mechanical history experienced by the polymer melt inside the cavity. Cavity pressure captures the filling, packing, and material compensation behavior, and is widely used for process monitoring, defect diagnosis, and quality prediction in injection molding (Q. Wang et al., 2021). Meanwhile, in-mold or mold-surface temperature governs melt viscosity, cooling rate, solidification, shrinkage, residual stress formation, and thus dimensional and mechanical properties (Kurt, Saban Kamber, Kaynak, Atakok, & Girit, 2009). In addition, recent machine-learning studies demonstrated that pressure-derived quality indices and cavity pressure profiles can effectively predict and interpret part quality (Párizs, Török, Ageyeva, & Kovács, 2022; Gim & Rhee, 2021). Therefore, this research mainly focused on these signals, which provide a physically meaningful basis for time-series-based quality prediction and subsequent temporal attribution analysis.

2.1.1. Data description

Two datasets are employed for the classification and regression tasks, respectively. A detailed description of both datasets is provided in Table 1.

Dataset 1, used for the classification task, originated from the research project “ProBayes”, which was conducted by SKZ (German Plastics Center) and Fraunhofer IPA between 2021 and 2022. The selected subset, consisting of ABS injection-molded parts, contains a total of 240 molding cycle records, with sink mark used as the target label. Since in-mold temperature signals were unavailable in this dataset, injection pressure, injection speed, and position of the injection unit were selected as surrogate process signals. Injection pressure reflects the flow resistance and packing behavior of the polymer melt, while injection speed influences shear heating, flow-front propagation, and filling dynamics. In addition, the position of the injection unit provides information about screw movement and melt advancement during the filling and packing stages. Together, these signals capture the thermo-mechanical evolution of the molding process and its relationship to final part quality.

Dataset 2, used for the regression task, originated from the 4-th Industrial Big Data Competition organized by Foxconn. The dataset was collected from real-world industrial injection molding production systems and associated sensors. Although the original download link is no longer available, the dataset was reorganized and republished by Deng et al. (Deng et al., 2025). The dataset contains approximately 16,600 molding cycle records and includes high-frequency time-series sensor data, statistical process control data, and product quality labels represented by part dimensions. In this work, the analysis mainly focused on cavity pressure and cavity temperature signals, which directly reflect the thermo-mechanical history of the polymer melt during molding. Screw position signals were additionally included to provide complementary information regarding process dynamics and melt flow progression.

2.1.2. Data preprocessing

To ensure data quality and reduce computational cost, signals exhibiting low variance during the sampling period were removed, as such signals contained limited dynamic process information and contributed minimally to quality prediction. In addition, all time-series signals were resampled to 512 sampling points in order to unify the input dimensions and facilitate batch-based model training. This resampling procedure also preserved the overall temporal evolution and characteristic patterns of the molding process while reducing the computational burden associated with varying sequence lengths.

All input signals and prediction labels for the regression task were further normalized using z-score normalization to improve training stability and convergence performance. Since

Table 1. Description of datasets.

Dataset	Feature name	Explanation	Range	Unit
Dataset 1	DXP_MldCavPrs1Act	Mold cavity pressure	[-1,266]	bar
	DXP_Inj1PrsAct	Injection pressure	[-72,649]	bar
	DXP_Inj1VelAct	Injection speed	[-90,18]	mm/s
	DXP_Inju1PosAct	Position of injection unit	[-1,6]	mm
	LBL_SinkMarks	Manual visual defect inspection label	0 or 1	-
Dataset 2	Sensor 1	Mold cavity pressure 1	[3,74]	bar
	Sensor 2	Mold cavity pressure 2	[-10,213]	bar
	Sensor 3	Mold cavity pressure 3	[1,72]	bar
	Sensor 5	Mold temperature 1	[82,258]	°C
	Sensor 6	Mold temperature 2	[81,286]	°C
	SP	Screw position	[16,107]	mm
	Size	Dimension of a specific section of the molded product	[199.836,200.191]	mm

the collected process signals originated from different sensors with different physical units and value ranges, normalization prevented features with larger numerical scales from dominating the optimization process. Moreover, z-score normalization helped the neural network learn relative temporal variations rather than absolute magnitudes, thereby improving model robustness and generalization capability. The normalization process is defined as follows:

$$z = \frac{x - \mu}{\sigma} \tag{1}$$

where x represents the original value, μ denotes the mean value of the corresponding feature, and σ represents the standard deviation of the corresponding feature.

2.2. Temporal Explainable Quality Prediction Model

Figure 1 illustrates the overall architecture of the proposed temporal explainable quality prediction framework, which consists of two main components: (A) representation learning-based quality prediction and (B) latent attribution and temporal interpretation. During the injection molding process, high-frequency time-series signals were continuously recorded from process-related sensors. After preprocessing and resampling, the signals were provided as inputs to the encoder network, which compressed the high-dimensional temporal data into a compact latent representation space $\mathbf{z} \in \mathbb{R}^d$. The extracted latent features were subsequently fed into a prediction model to estimate the final product quality.

To improve the interpretability of the prediction model, the SHAP method was applied in the latent representation space to quantify the contribution of each latent feature to the prediction result. The obtained latent feature attributions were then propagated back to the original temporal domain through the decoder network, which was jointly trained with the encoder during the representation learning stage. By aggregating the mapped attribution values over time, temporal importance distributions were obtained, allowing the critical process stages and signal regions contributing to product quality variation to be identified. In this way, the proposed framework established a connection between latent feature representations and physically meaningful temporal process behavior, enabling interpretable quality prediction for injection molding processes. The proposed framework was implemented using PyTorch in Python.

2.2.1. Latent representation learning

A one-dimensional convolutional autoencoder was developed to learn compact latent representations from high-frequency process signals. The model consisted of an encoder for temporal feature extraction and dimensionality reduction, and a decoder for signal reconstruction. The learned latent representation was subsequently utilized for downstream quality prediction and explainability analysis.

The autoencoder took multivariate process signals as input, including in-mold sensor signals and machine-related temporal process variables. All signals were resampled to a fixed sequence length of 512 time steps prior to training.

The encoder consisted of three one-dimensional convolutional (1D-CNN) layers with progressively increasing numbers of feature channels (32, 64, and 128). The three convolutional layers used kernel sizes of 7, 5, and 5, respectively, with a stride of 2 for each layer. Each convolutional layer was followed by batch normalization and a ReLU activation function. The strided convolutions simultaneously performed feature extraction and temporal downsampling, reducing the sequence length from 512 to 64 time steps. Appropriate padding was applied to preserve the expected dimensional reduction at each convolutional stage. The resulting feature maps were then flattened and passed through a fully connected layer to obtain a compact latent representation. The dimensionality of the latent representation was set to 32 in this study.

The decoder mirrored the encoder architecture and reconstructed the original temporal signals from the latent representation. First, the latent vector was expanded through a fully connected layer and reshaped into feature maps. Subsequently, three transposed convolutional layers were applied to progres-

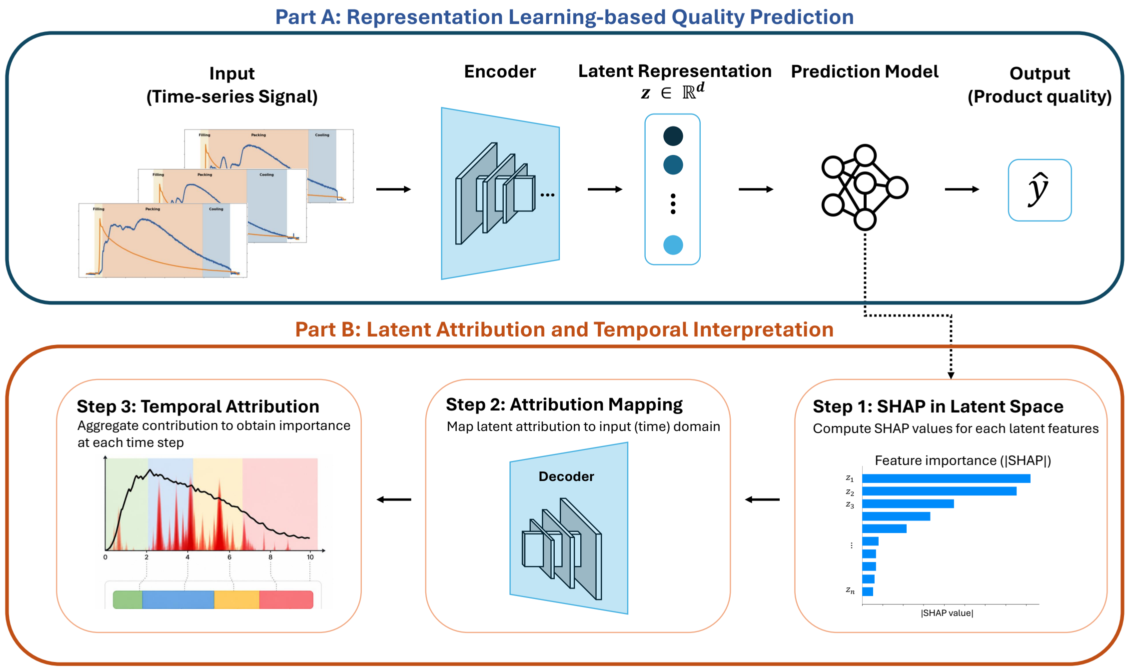

Figure 1. Framework overview.

sively recover the temporal resolution from 64 to 512 time steps. ReLU activation and batch normalization were applied between intermediate decoding layers. The final layer reconstructed the multivariate temporal signals with the original input dimensions.

The autoencoder was trained in a self-supervised manner using a reconstruction objective. The mean squared error (MSE)

$$L_{\text{rec}} = \frac{1}{N} \sum_{i=1}^{N} \|x_i - \hat{x}_i\|^2 \tag{2}$$

between the reconstructed signal and the original input signal was minimized during training.

The reconstruction objective encourages the latent representation to preserve the essential temporal dynamics and process-related information contained in the original high-dimensional signals.

The input signals were randomly split into training, validation, and test sets at a ratio of 70:15:15, with a random seed of 42 to ensure reproducibility. A batch size of 64 was used during training. To avoid data leakage, the normalization parameters were computed exclusively from the training set and subsequently applied to the validation and test sets. The model was trained using the Adam optimizer with an initial learning rate of 0.001 for a maximum of 100 epochs. Early stopping was employed based on the validation loss with a patience of 15 epochs to mitigate overfitting.

During each training iteration, the reconstruction loss was computed between the reconstructed output and the original input signals. The model parameters corresponding to the lowest validation loss were retained for subsequent analysis.

2.2.2. Latent space quality prediction

After the latent representation learning stage, the trained encoder was employed as a feature extractor to generate compact latent embeddings from the multivariate temporal process signals. These latent representations were subsequently utilized as inputs for downstream quality prediction tasks. A multi-layer perceptron (MLP) predictor was developed to establish the mapping between the latent feature space and the target quality variables.

The network consisted of two fully connected hidden layers with 64 and 32 neurons, respectively, followed by ReLU activation and dropout regularization (dropout rate = 0.2). The output layer contained a single neuron, producing either a continuous quality value for regression tasks or a logit value for binary classification tasks. The predictor was trained using the latent representations extracted by the frozen encoder, enabling the mapping between the learned latent space and the target quality variable.

The encoder parameters were frozen after the autoencoder training stage, and the extracted latent vectors were treated as fixed process-aware representations for the downstream prediction models. This two-stage learning strategy separated temporal representation learning from downstream quality prediction, enabling the latent space to preserve the temporal process characteristics learned during the reconstruction task.

Depending on the downstream objective, the latent-space prediction framework was configured either for regression or binary classification tasks. For regression tasks, the model produced continuous quality values and was trained using the mean squared error (MSE) loss:

$$L_{\text{reg}} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{3}$$

where y_i and $\hat{y}_i$ denote the ground-truth and predicted quality values, respectively, and N represents the number of samples.

For classification tasks, the model produced binary defect probabilities and was trained using the weighted binary cross-entropy loss to account for class imbalance:

$$L_{\text{cls}} = -\frac{1}{N}\sum_{i=1}^{N}\left[w_p y_i \log(\sigma(\hat{y}_i)) + (1 - y_i)\log(1 - \sigma(\hat{y}_i))\right] \tag{4}$$

where $\sigma(\cdot)$ denotes the sigmoid activation function, w_p is the positive class weight introduced to compensate for the imbalance between defective and normal samples, and $\hat{y}_i$ represents the predicted logit output.

The input data were randomly split into training, validation, and test sets at a ratio of 70:15:15, with a random seed of 42 to ensure reproducibility. A batch size of 64 was used during training. The MLP predictors for both regression and classification tasks were optimized using the Adam optimizer with an initial learning rate of 1×10^{-3}. The maximum number of training epochs was set to 100. Early stopping based on the validation loss was employed with a patience value of 15 epochs to mitigate overfitting. The model parameters corresponding to the minimum validation loss were retained for subsequent evaluation and explainability analysis.

Due to the limited amount of data available for the classification task, five-fold cross-validation was employed to provide a more robust evaluation of the predictive performance and generalization capability of the proposed model. The dataset was randomly partitioned into five mutually exclusive subsets of approximately equal size. In each fold, four subsets were used for model training, while the remaining subset was used for validation. This procedure was repeated five times so that each subset served as the validation set once. The final performance metrics were calculated by averaging the results obtained across the five folds.

2.2.3. Latent Feature Attribution using SHAP

To interpret the contribution of latent representations to the downstream quality prediction, SHAP is applied to the latent-space prediction model. After the latent representations are extracted by the trained encoder, a SHAP DeepExplainer is employed to compute the contribution of each latent feature to the prediction output. Since the downstream quality predictor is implemented as an MLP, a wrapper module is introduced to adapt the predictor for SHAP-based analysis.

Given a latent representation vector $\mathbf{z} \in \mathbb{R}^d$, the SHAP method estimates the contribution of each latent feature z_j to the model prediction $f(\mathbf{z})$. The prediction can be expressed as:

$$f(\mathbf{z}) = \phi_0 + \sum_{j=1}^{d}\phi_j \tag{5}$$

where ϕ_0 denotes the expected model output over the background samples, and ϕ_j represents the SHAP value corresponding to the j-th latent feature.

The latent-space SHAP values quantified the contribution of each latent dimension to the predicted product quality or defect probability. Positive SHAP values indicated a positive contribution to the prediction output, whereas negative SHAP values indicated an inhibitory effect. The mean absolute SHAP values were further utilized to evaluate the global importance of latent features across multiple samples.

The obtained latent feature attributions were subsequently propagated back to the temporal domain through the decoder structure to recover interpretable temporal attribution maps corresponding to the original process signals.

2.2.4. Decoder-based temporal attribution mapping

Although SHAP analysis in the latent space provides feature-level interpretability, the compact latent representation does not directly reveal the temporal regions contributing to the prediction results. To recover temporal interpretability, the latent SHAP values are mapped back to the original time domain through the decoder structure of the autoencoder.

In this work, the decoder is regarded as a differentiable mapping function from the latent representation space to the reconstructed temporal signals. Given the decoder function $D(\mathbf{z})$, the sensitivity between latent features and temporal signals is estimated using the decoder Jacobian:

$$\mathbf{J}_D = \frac{\partial D(\mathbf{z})}{\partial \mathbf{z}} \tag{6}$$

where $\mathbf{J}_D$ represents the Jacobian matrix of the decoder with respect to the latent representation.

The temporal attribution map is obtained by propagating the latent SHAP values through the decoder Jacobian:

$$\mathbf{A} = |\mathbf{J}_D|\,\phi \tag{7}$$

where ϕ denotes the latent SHAP vector and $\mathbf{A}$ represents the reconstructed temporal attribution map in the original signal domain.

The framework supports both signed and magnitude-based attribution analyses. The resulting temporal attribution maps are visualized as heatmaps and can be overlaid on the original time-series process signals, enabling the identification of the most influential temporal regions contributing to the model predictions.

To further investigate the contribution of different molding stages, phase-aware attribution analysis was performed by aggregating the temporal attributions according to the molding phases. This enabled the evaluation of the relative importance of filling, packing, and cooling stages in the downstream quality prediction process.

The results were compared with those obtained by directly computing feature attribution on high-dimensional temporal signals regarding the attribution complexity and the robustness of the interpretation process.

2.2.5. Robustness Evaluation of Attribution Maps

To assess the robustness of the generated explanations, a perturbation-based evaluation strategy was employed. Gaussian noise with varying amplitudes was added to the original input signals, and attribution maps were recomputed using both the proposed decoder-based attribution framework and a direct SHAP analysis applied in the original signal domain.

The robustness of an explanation method can be characterized by its ability to preserve attribution patterns under moderate input perturbations. Ideally, explanations should remain stable when the underlying process characteristics are unchanged and only small random disturbances are introduced.

For a given sample, let $\mathbf{A}^{(0)}$ denote the attribution map obtained from the original signal and $\mathbf{A}^{(\sigma)}$ denote the attribution map generated from the perturbed signal with Gaussian noise level σ. Two complementary similarity measures were used to quantify explanation consistency.

First, cosine similarity CosSim was employed to evaluate the overall similarity of attribution patterns:

$$\mathrm{CosSim} = \frac{\mathbf{A}^{(0)} \cdot \mathbf{A}^{(\sigma)}}{\left\|\mathbf{A}^{(0)}\right\|_2 \left\|\mathbf{A}^{(\sigma)}\right\|_2} \tag{8}$$

where $\cdot$ denotes the dot product and $\|\cdot\|_2$ represents the Euclidean norm. A cosine similarity value close to 1 indicates that the overall attribution structure remains unchanged after perturbation.

Second, Spearman rank correlation ρ_s was used to evaluate the consistency of the attribution ranking:

$$\rho_s = 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)} \tag{9}$$

where d_i is the difference between the ranks of the i-th attribution element in the original and perturbed attribution maps, and n denotes the total number of attribution elements.

Unlike cosine similarity, which measures global structural similarity, Spearman correlation evaluates whether the relative importance ordering of temporal regions is preserved. Therefore, the two metrics provide complementary perspectives on explanation robustness.

Higher values of both cosine similarity and Spearman correlation indicate greater robustness and consistency of the attribution method under input perturbations.

3. Results and Discussion

3.1. Reconstruction and latent representation analysis

To evaluate the capability of the proposed autoencoder in capturing the underlying temporal process dynamics, the reconstruction performance was assessed on both datasets. Figures 2 and 3 present representative examples of the original and reconstructed signals randomly selected from the classification task and regression task test sets respectively.

For the classification dataset, the autoencoder achieved a reconstruction MSE of 0.0075, while a lower reconstruction MSE of 0.0013 was obtained for the regression dataset. The low reconstruction errors indicate that the learned latent representations preserve the majority of information contained in the original multivariate temporal signals.

For the classification task, the reconstructed signals successfully preserve the major temporal patterns and process-state transitions observed in the original signals. For the injection pressure (S2), injection velocity (S3), and screw position (S4) signals, the reconstructed curves closely follow the original trajectories, accurately reproducing the characteristic step changes and plateau regions. The reconstruction quality is particularly high for S3, where the major process stages can be clearly identified from the reconstructed signals.

However, small high-frequency fluctuations are visible in several reconstructed signals, particularly in the plateau regions. These fluctuations are likely introduced by the decoder during the signal reconstruction process and may be attributed to reconstruction errors or limitations of the latent representation. Nevertheless, the magnitude of these fluctuations remains relatively small compared with the overall signal variation and does not affect the identification of the major process dynam-

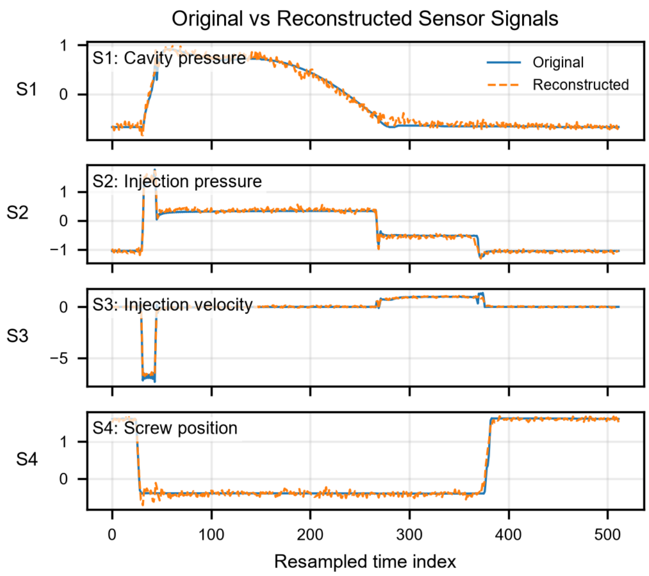

Figure 2. Comparison of original and reconstructed sensor signals (classification task).

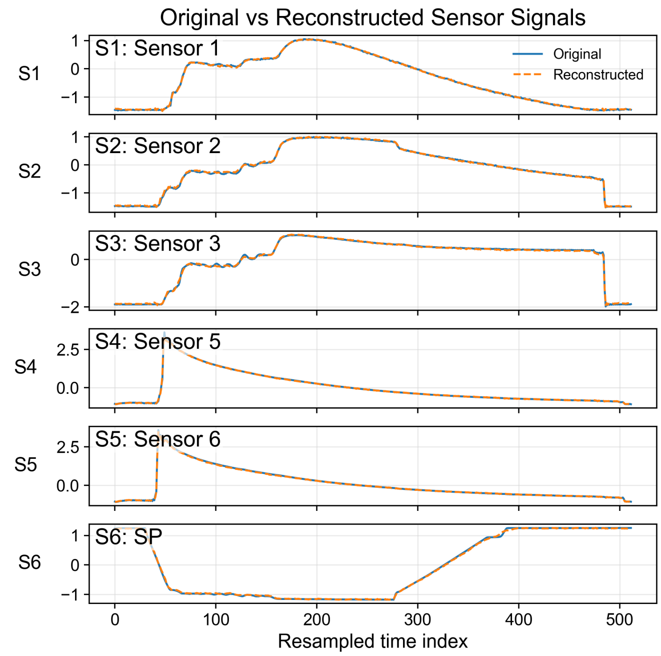

Figure 3. Comparison of original and reconstructed sensor signals (regression task).

ics.

For the regression task, the reconstructed signals exhibit an almost complete overlap with the original measurements across all sensor channels. Both the transient responses and the long-term temporal trends are accurately recovered by the decoder, resulting in a reconstruction MSE of only 0.0013.

The reconstruction results provide indirect evidence regarding the quality of the learned latent representations. Despite the substantial dimensionality reduction introduced by the encoder, the decoder is capable of recovering the essential temporal characteristics of the original signals with high fidelity. These observations suggest that the latent space contains compact yet informative representations of the molding process dynamics. Consequently, the latent embeddings provide a suitable feature space for downstream quality prediction and explainability analysis.

Furthermore, the successful reconstruction of key process transitions indicates that the latent representations retain process-relevant information associated with different molding stages, which is particularly important for the subsequent decoder-based temporal attribution mapping framework.

3.2. Quality prediction performance

For the classification task, the proposed framework achieved an accuracy of 1.0, an F1-score of 1.0, and an ROC-AUC of 1.0 on the test dataset. As shown in Figure 4(a), all good and defective samples were correctly identified.

For the five-fold cross-validation, the proposed classification model achieved consistently high predictive performance across the different folds. The mean test accuracy reached 0.990 with a standard deviation of 0.022. The mean recall and F1-score were 0.960 ± 0.089 and 0.978 ± 0.050, respectively. In addition, the mean ROC-AUC reached 0.990 ± 0.023, suggesting a strong capability to discriminate between normal and defective samples. The relatively small variation across folds indicates that the model performance was generally stable despite the limited dataset size. The autoencoder also showed consistent reconstruction performance, with an average test reconstruction MSE of 0.0122 ± 0.0013.

For the regression task, the proposed framework achieved an R^2 score of 0.749, an RMSE of 0.026, and an MAE of 0.021 on the test dataset. As illustrated in Figure 4(b), the predicted-versus-actual plot demonstrates a clear positive correlation between the predicted and measured quality values, indicating that the latent representations successfully preserve quality-relevant process information. Most samples are distributed close to the ideal prediction line, suggesting satisfactory predictive performance.

Nevertheless, a tendency toward prediction smoothing can be observed. Samples with relatively high target values are slightly underestimated, whereas samples with relatively low target values tend to be overestimated. This behavior indicates that the model effectively captures the overall trend of quality variation but exhibits limited sensitivity to extreme deviations.

Overall, the regression task demonstrates that the learned latent representations contain substantial information related to part quality, while the classification task suggests the presence of defect-related features within the latent space. Although the predictive performance is not the primary focus of this work, these results indicate that the encoder successfully captures meaningful process dynamics that can be utilized for both quality prediction and subsequent explainability analysis.

The prediction models therefore provide a suitable basis for investigating the proposed decoder-based temporal attribution framework presented in the following sections.

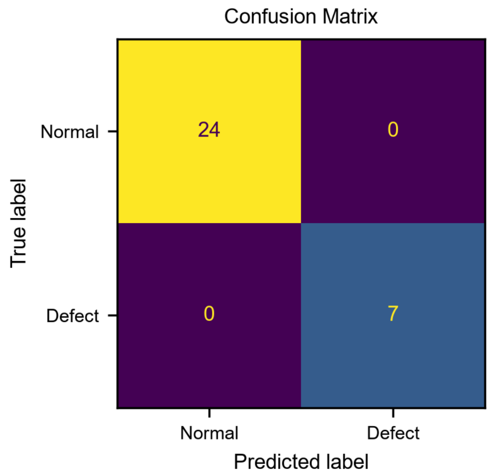

(a) Classification task.

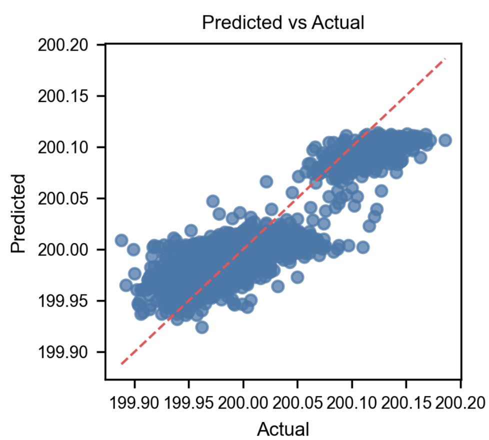

(b) Regression task.

Figure 4. Prediction performance of the proposed framework for (a) the classification task and (b) the regression task.

3.3. Temporal Attribution Analysis

3.3.1. Visualization of temporal attribution

The decoder-based temporal attribution heatmaps for the classification and regression tasks are presented in Figures 5(a) and 5(b), respectively. The attribution values represent the contribution of individual temporal regions and sensor channels to the final prediction output after propagating the latent SHAP values through the decoder Jacobian.

For the sink mark classification task, the attribution is concentrated in specific temporal regions rather than being uniformly distributed throughout the molding cycle. High-attribution regions are mainly observed in the cavity pressure signal and the injection pressure signal between approximately 50 and 300 resampled time indices. According to the corresponding process signals shown in Figure 2, this period roughly corresponds to the filling-to-packing transition and the subsequent packing phase. In addition, elevated attribution values can be observed in the early stage of the injection velocity signal, which corresponds to the velocity-controlled filling stage. A localized high-attribution region is also identified in the later stage of the screw position signal, corresponding to the end of the packing phase. These results suggest that sink mark prediction is influenced not only by pressure-related process dynamics inside the cavity but also by machine-side process variables associated with material delivery and packing control.

For the dimension prediction task, the attribution heatmap reveals a clear dominance of the cavity pressure signals. According to Table 1, Sensor1, Sensor2, and Sensor3 correspond to pressure measurements acquired at different locations inside the mold cavity. The highest attribution values are concentrated within these pressure signals, particularly during the early-to-middle portion of the molding cycle, with Sensor3 exhibiting the strongest contribution. This observation indicates that dimensional quality is primarily governed by the pressure evolution within the cavity. Differences in attribution magnitude among the three pressure sensors further suggest that sensor location may influence the relevance of the captured process information. However, because the original dataset does not provide detailed information regarding sensor placement, a more comprehensive analysis of spatial pressure propagation and sensor localization effects cannot be performed in this study.

Compared with the pressure signals, Sensor5 and Sensor6, corresponding to mold temperature measurements, exhibit weaker but still noticeable attribution values in localized temporal regions. In contrast, the screw position signal (SP) contributes relatively little to the final prediction. These results indicate that dimensional quality prediction is mainly driven by cavity pressure evolution, while mold temperature provides complementary information related to thermal history and shrinkage behavior. The influence of screw position appears to be more indirect and is likely reflected through its effect on the pressure development inside the cavity.

Overall, the obtained attribution patterns demonstrate that the proposed decoder-based attribution mapping successfully transforms latent-space explanations into physically interpretable temporal importance distributions. The resulting heatmaps not only identify the most influential sensor channels but also reveal the critical process periods that contribute to defect for-

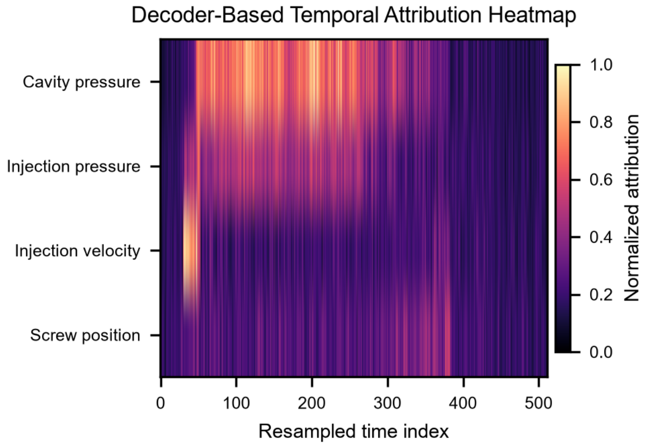

(a) Classification task

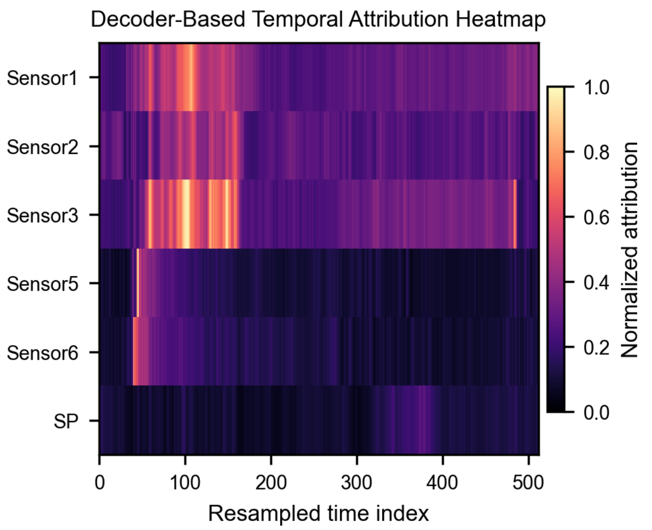

(b) Regression task

Figure 5. Decoder-based temporal attribution heatmaps for (a) the classification task and (b) the regression task.

mation and dimensional variation, thereby providing valuable insights into the underlying process–quality relationships.

To further quantify the relative importance of different process signals, the temporal attribution values were aggregated over the entire molding cycle for each sensor channel. The resulting sensor contribution ranking for the classification task is shown in Figure 6(a).

Among the four process variables, the cavity pressure signal exhibits the highest average attribution, followed by the injection pressure signal, screw position, and injection velocity. Although the cavity pressure signal shows the strongest contribution, the differences between the four signals remain relatively small.

This observation suggests that the sink mark prediction model does not rely on a single dominant process variable. Instead, the prediction is driven by the combined influence of multiple process signals that jointly characterize the filling, pressure transmission, and material compensation behaviors during the molding cycle.

The highest contribution associated with cavity pressure is physically meaningful, since cavity pressure directly reflects the pressure evolution inside the mold cavity and is closely related to volumetric shrinkage, which is one of the primary mechanisms responsible for sink mark formation. The substantial contribution of injection pressure further indicates that the pressure transfer process between the machine and the cavity plays an important role in determining defect occurrence.

Meanwhile, the non-negligible contributions of screw position and injection velocity imply that machine-side process dynamics also provide useful information for defect prediction. These signals indirectly characterize the filling behavior and material delivery process, which influence the subsequent pressure development inside the cavity.

For the regression task, the sensor contribution ranking in Figure 6(b) shows that the three mold cavity pressure signals dominate the dimensional quality prediction. According to Table 1, Sensor 1, Sensor 2, and Sensor 3 correspond to mold cavity pressure measurements at different sensing locations. Among them, Sensor 3 exhibits the highest contribution, followed by Sensor 1 and Sensor 2. This result indicates that the regression model relies primarily on in-mold pressure evolution to estimate the final part dimension, which is consistent with the physical role of cavity pressure in material compression, pressure transmission, and shrinkage compensation.

The mold temperature signals, represented by Sensor 5 and Sensor 6 in Table 1, show lower but still observable contributions. This suggests that thermal conditions also influence dimensional quality by affecting melt viscosity, solidification behavior, and cooling-induced shrinkage. In contrast, the screw position signal has one of the lowest contribution values. This is reasonable because screw position mainly reflects machine-side material delivery, whereas the cavity pressure sensors directly capture the in-mold state experienced by the polymer melt. Overall, the regression sensor ranking supports the heatmap-based observation that dimensional quality prediction is mainly driven by pressure-related process dynamics, with temperature providing complementary information and screw position contributing indirectly.

3.3.2. Process-phase-aware interpretation

To quantitatively evaluate the contributions of different molding stages, the temporal attribution values were aggregated over the filling, packing, and cooling phases. Since stage annotations were only available for the classification dataset, the phase-level analysis was conducted exclusively for this dataset. The resulting phase-level contribution shares for the sink mark

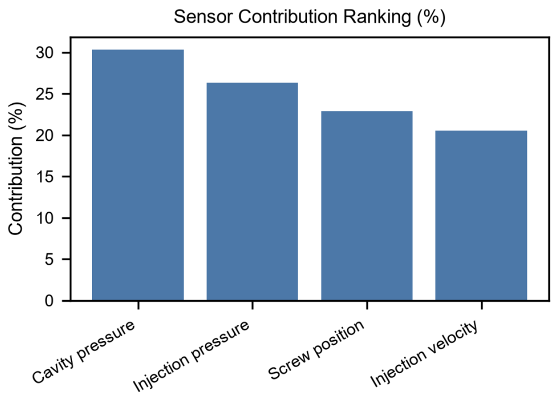

(a) Classification task.

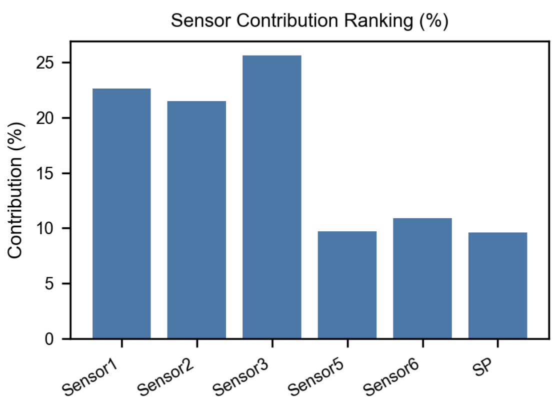

(b) Regression task.

Figure 6. Sensor contribution ranking obtained from the decoder-based attribution framework for (a) the classification task and (b) the regression task.

classification task are presented in Figure 7.

The packing phase exhibits the highest contribution, accounting for approximately 43.5% of the total attribution, followed by the cooling phase (37.6%) and the filling phase (18.9%). This result indicates that the prediction of sink mark defects is primarily determined by process dynamics occurring after cavity filling.

The relatively low contribution of the filling phase suggests that the occurrence of sink marks is not predominantly governed by the initial flow behavior. Instead, the subsequent material compensation and thermal shrinkage processes play a more important role. The dominant contribution of the packing phase is consistent with the physical mechanism of sink mark formation, as insufficient packing pressure or inadequate pressure transmission may lead to volumetric shrinkage within the molded part.

Furthermore, the substantial contribution observed during the cooling phase indicates that thermal history remains relevant for defect prediction. This observation suggests that the evolution of residual shrinkage after packing termination also

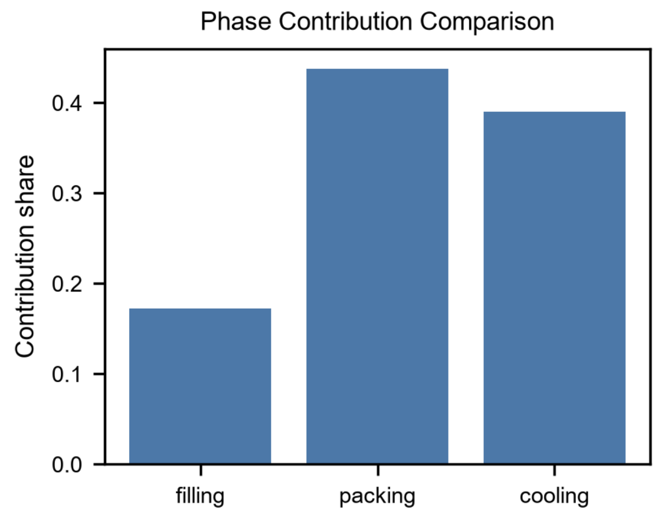

Figure 7. Phase contribution comparison (classification task).

influences the final occurrence of sink marks.

3.3.3. Interpretation of Process-related Temporal Dynamics

Beyond identifying important process phases, the proposed framework enables the interpretation of specific process variables contributing to quality prediction. The physical plausibility of the obtained attribution results was examined by relating the identified important temporal regions and sensor channels to established injection molding mechanisms.

For the sink mark classification task, high attribution values are observed in the cavity pressure and injection pressure signals during the filling-to-packing transition and throughout the packing phase. This finding is consistent with the well-established mechanism of sink mark formation. During the packing stage, additional molten material is supplied into the cavity to compensate for volumetric shrinkage caused by cooling and solidification. Insufficient packing pressure or ineffective pressure transmission can result in local material deficiency and subsequent surface depressions. Previous studies have shown that packing pressure is one of the most influential process variables affecting shrinkage, sink marks and residual stress in injection-molded parts (Chen & Gao, 2003; Kariminejad, Tormey, Huq, Morrison, & McAfee, 2021).

The phase-level attribution analysis further indicates that the packing phase contributes the largest share of the overall explanation. This result agrees well with injection molding theory, where the quality of the final part is largely determined after cavity filling, during the material compensation process before gate freeze-off (Kariminejad et al., 2021). The attribution framework therefore successfully identifies the process stage that is physically responsible for sink mark formation rather than merely detecting statistical correlations.

For the dimension prediction task, the attribution analysis indicates that cavity pressure signals contribute more strongly than temperature and screw position signals. This observation is physically meaningful because cavity pressure directly reflects the actual thermo-mechanical state of the polymer melt inside the cavity and integrates the combined effects of multiple processing variables. Speranza et al. demonstrated that cavity pressure measurements can be used to estimate solidification history and shrinkage behavior (Speranza, Vietri, & Pantani, 2013), while recent reviews have highlighted cavity pressure as one of the most informative indicators for dimensional stability and product quality monitoring (Liew, Peng, Huang, & Su, 2022; Zhao, Liu, Su, & Xu, 2024).

Particularly, the high attribution values observed around the filling-to-packing transition and the early packing stage correspond to the period when cavity pressure rapidly builds up and material compensation occurs. According to recent measurement and process-monitoring studies, pressure and temperature conditions after the velocity-to-pressure switchover significantly influence polymer shrinkage and dimensional variation (Zhao et al., 2024; Gaspar-Cunha, Melo, Marques, & Pontes, 2025). Therefore, the temporal regions identified by the proposed framework coincide with the physically critical stages governing final part dimensions.

Temperature signals also exhibit localized importance immediately after cavity filling. This behavior suggests that thermal conditions affect melt viscosity, solidification behavior, and cooling-induced shrinkage. During the cooling stage, the gradual decrease of cavity pressure and the development of thermal gradients lead to differential shrinkage within the molded part. Such effects have been widely recognized as major contributors to dimensional variation and defect formation (Speranza et al., 2013; Zhao et al., 2024).

Compared with cavity pressure and temperature signals, screw position exhibits a relatively lower attribution. This is reasonable from a process-physics perspective because screw position primarily describes machine-side material delivery, whereas cavity pressure directly reflects the resulting in-mold process state. Once cavity pressure measurements are available, part of the information carried by screw position is implicitly represented through the pressure response of the melt. Similar observations have been reported in process-monitoring studies, where in-mold measurements generally provide stronger correlations with final product quality than machine-side process variables (Liew et al., 2022; Zhao et al., 2024).

Overall, the attribution results indicate that the prediction models primarily rely on process information associated with pressure build-up, material compensation, solidification, and cooling-induced shrinkage. The agreement between the identified important regions and established injection molding knowledge supports the physical plausibility of the proposed decoder-based attribution framework.

Nevertheless, several important physical factors are not explicitly represented in the current explanation results. Sink marks and dimensional deviations are also influenced by geometric characteristics such as wall thickness distribution, rib geometry, gate location, and local cooling efficiency. For example, sink marks frequently occur in thick sections where the cooling rate is lower and volumetric shrinkage is more pronounced (Gruber, Berger, Pacher, & Friesenbichler, 2011). Similarly, pressure propagation along the flow path and spatial temperature gradients may affect local quality variations. Since the present framework only utilizes temporal sensor signals, these geometric and spatial effects can only be captured indirectly through their influence on the measured pressure and temperature responses. Future work could integrate mold geometry information, simulation-derived process variables, or distributed sensor measurements to further improve the physical interpretability of the generated explanations (Párizs, Török, Ageyeva, & Kovács, 2023; Araújo, Pereira, Dias, Marques, & Cruz, 2023).

The consistency between the attribution results and established injection molding mechanisms suggests that the proposed explainability framework not only identifies statistically important features but also captures physically meaningful process dynamics associated with defect formation and dimensional variation.

3.4. Robustness and Consistency Analysis

To evaluate the robustness of the proposed decoder-based attribution framework, Gaussian noise with different amplitudes was added to the input signals, and the resulting attribution maps were compared with the original explanations. The robustness was quantified using cosine similarity and Spearman correlation between the perturbed and reference attribution maps.

The proposed decoder-based attribution mapping was compared against a direct SHAP-based explanation approach applied in the original signal domain. A robust explanation method is expected to maintain a consistent attribution pattern under moderate input perturbations, resulting in high similarity and correlation values.

For the sink mark classification task, the proposed decoder-based attribution framework exhibits substantially higher robustness against input perturbations than the direct SHAP approach. As shown in Figure 8, the cosine similarity of the decoder-based explanations remains close to unity across all noise levels, decreasing only from 1.00 to approximately 0.99 when the noise level increases from 0% to 50%. In contrast, the cosine similarity of the direct SHAP explanations drops significantly to approximately 0.67 under the same perturbation.

A similar trend can be observed for the Spearman correlation. While the direct SHAP explanations experience a substantial deterioration from 1.0 to 0.48, the decoder-based explanations maintain a correlation above 0.92 even at the highest noise level. These results indicate that the proposed framework preserves both the overall attribution direction and the relative ranking of important temporal regions under severe signal perturbations.

Similar observations can be made for the dimension prediction task. Although both methods achieve high cosine similarity values, the decoder-based framework consistently outperforms direct SHAP under increasing noise levels.

More importantly, the difference becomes pronounced when considering the Spearman correlation. The proposed method maintains correlation values close to 1.0 across all perturbation levels, whereas the direct SHAP explanations exhibit a continuous degradation from 1.0 to 0.41.

This result suggests that direct SHAP explanations remain globally similar under perturbation but suffer from substantial changes in the ranking of important temporal regions. In contrast, the decoder-based framework preserves both the global attribution structure and the local importance ordering, indicating a higher degree of explanation consistency.

The superior robustness of the proposed framework can be attributed to the latent representation learning stage. Instead of directly computing explanations in the high-dimensional signal space, the proposed approach first extracts compact latent features that capture the dominant process dynamics while suppressing irrelevant fluctuations and noise.

Consequently, the attribution analysis is performed in a lower-dimensional and more structured representation space, reducing the sensitivity of the explanations to perturbations in the raw signals. The decoder subsequently propagates the latent attributions back to the temporal domain, enabling interpretable visualization while preserving the robustness inherited from the latent representation.

Overall, the results demonstrate that the proposed decoder-based temporal attribution framework produces explanations that are considerably more stable and consistent than direct SHAP analysis, thereby providing a more reliable basis for process interpretation and decision support in injection molding applications.

4. Conclusion

This paper presents an explainable temporal attribution framework for time-series-based quality prediction in injection molding. The proposed method combines latent representation learning, latent-space quality prediction, and decoder-based temporal attribution mapping. By first encoding multivariate process signals into a compact latent space and then propagating latent SHAP attributions back to the original temporal domain through the decoder Jacobian, the framework links prediction-relevant latent variables with physically meaningful process dynamics.

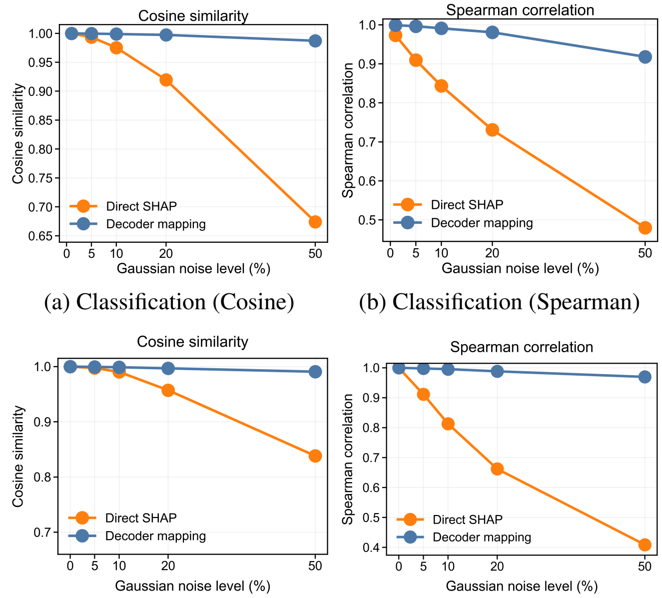

(a) Classification (Cosine) (b) Classification (Spearman)

(c) Regression (Cosine) (d) Regression (Spearman)

Figure 8. Comparison of attribution robustness under Gaussian noise perturbations. Cosine similarity and Spearman rank correlation are reported for the direct SHAP and decoder-based attribution mapping methods in both classification and regression tasks.

The experimental results on both classification and regression tasks show that the convolutional autoencoder preserves the dominant temporal characteristics of the original process signals with low reconstruction errors. The learned latent representations provide useful information for downstream quality prediction, especially for dimensional regression, where the model achieves satisfactory predictive performance. For sink mark classification, the model successfully identifies defective samples but produces a relatively high number of false positives, indicating that further improvement is still needed for defect discrimination under limited and imbalanced data conditions.

The attribution analysis demonstrates that the proposed framework can transform latent-space explanations into interpretable temporal importance distributions. For sink mark prediction, high attribution values are mainly observed around the filling-to-packing transition and the packing stage, while phase-aware analysis indicates that packing and cooling contribute more strongly than filling. For dimensional prediction, cavity pressure signals dominate the attribution results, which is consistent with the physical relationship between pressure evolution, shrinkage behavior, and final part dimensions. These findings suggest that the proposed explanations are not only model-based but also aligned with injection molding process knowledge.

Compared with direct SHAP analysis in the original high-dimensional signal domain, the decoder-based attribution framework exhibits higher robustness and consistency under noise perturbations. This indicates that performing attribution in a structured latent space and mapping the results back to the temporal domain can reduce sensitivity to local signal fluctuations while retaining process-level interpretability. Overall, the proposed framework provides a promising approach for interpretable quality prediction, root-cause analysis, and process understanding in injection molding.

Future work will focus on improving classification performance under imbalanced data conditions, validating the framework on larger industrial datasets, and extending the phase-aware attribution analysis toward closed-loop process optimization. In addition, further studies will investigate the effectiveness of the proposed framework from multiple perspectives, including cross-dataset generalization, robustness under varying process conditions, consistency across different explainability methods, and agreement with established process physics. The incorporation of spatial sensor information and physics-informed constraints will also be explored to establish a more comprehensive spatio-temporal explainability framework for injection molding processes.

REFERENCES

Araújo, C., Pereira, D., Dias, D., Marques, R., & Cruz, S. (2023). In-cavity pressure measurements for failure diagnosis in the injection moulding process and correlation with numerical simulation. *The International Journal of Advanced Manufacturing Technology*, *126*(1), 291–300. doi: 10.1007/s00170-023-11100-1

Bento, J., Saleiro, P., Cruz, A. F., Figueiredo, M. A. T., & Bizarro, P. (2021). TimeSHAP: Explaining Recurrent Models through Sequence Perturbations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (pp. 2565–2573). doi: 10.1145/3447548.3467166

Chen, X., & Gao, F. (2003). A study of packing profile on injection molded part quality. *Materials Science and Engineering: A*, *358*(1), 205–213. doi: 10.1016/S0921-5093(03)00290-9

Deng, X., Xiang, W., Lin, W., Zheng, Z., & Yang, Y. (2025). Few-shot injection molding quality prediction method integrating deep transfer and multi-task learning. *The International Journal of Advanced Manufacturing Technology*, *141*(1), 609–632. doi: 10.1007/s00170-025-16651-z

Ding, W., Ali, H., Gao, K., Zhang, Z., & Gao, F. (2025). Novel deep learning based soft sensor feature extraction for part weight prediction in injection molding processes. *Journal of Manufacturing Systems*, *78*, 58–68. doi: 10.1016/j.jmsy.2024.11.011

Finkeldey, F., Volke, J., Zarges, J.-C., Heim, H.-P., & Wiederkehr, P. (2020). Learning quality characteristics for plastic injection molding processes using a combination of simulated and measured data. *Journal of Manufacturing Processes*, *60*, 134–143. doi: 10.1016/j.jmapro.2020.10.028

Gaspar-Cunha, A., Melo, J., Marques, T., & Pontes, A. (2025). A Review on Injection Molding: Conformal Cooling Channels, Modelling, Surrogate Models and Multi-Objective Optimization. *Polymers*, *17*(7), 919. doi: 10.3390/polym17070919

Gim, J., & Rhee, B. (2021). Novel Analysis Methodology of Cavity Pressure Profiles in Injection-Molding Processes Using Interpretation of Machine Learning Model. *Polymers*, *13*(19), 3297. doi: 10.3390/polym13193297

Gordon, G., Kazmer, D. O., Tang, X., Fan, Z., & Gao, R. X. (2015). Quality control using a multivariate injection molding sensor. *The International Journal of Advanced Manufacturing Technology*, *78*(9), 1381–1391. doi: 10.1007/s00170-014-6706-6

Gruber, D. P., Berger, G., Pacher, G., & Friesenbichler, W. (2011). Novel approach to the measurement of the visual perceptibility of sink marks on injection molding parts. *Polymer Testing*, *30*(6), 651–656. doi: 10.1016/j.polymertesting.2011.04.013

Hong, J., Hong, Y., Baek, J.-W., Kang, S.-W., Hong, J., Hong, Y., … Kang, S.-W. (2025). Enhancing the Product Quality of the Injection Process Using eXplainable Artificial Intelligence. *Processes*, *13*(3). doi: 10.3390/pr13030912

Kariminejad, M., Tormey, D., Huq, S., Morrison, J., & McAfee, M. (2021). Ultrasound Sensors for Process Monitoring in Injection Moulding. *Sensors*, *21*(15), 5193. doi: 10.3390/s21155193

Ke, K.-C., Wang, J.-C., & Nian, S.-C. (2024). Data-driven quality prediction in injection molding: An autoencoder and machine learning approach. *Polymer Engineering & Science*, *64*(9), 4520–4538. doi: 10.1002/pen.26866

Kechris, C., Dan, J., & Atienza, D. (2025). *Time series saliency maps: Explaining models across multiple domains.* https://arxiv.org/abs/2505.13100v3.

Kurt, M., Saban Kamber, O., Kaynak, Y., Atakok, G., & Girit, O. (2009). Experimental investigation of plastic injection molding: Assessment of the effects of cavity pressure and mold temperature on the quality of the final products. *Materials & Design*, *30*(8), 3217–3224. doi: 10.1016/j.matdes.2009.01.004

Liew, K.-F., Peng, H.-S., Huang, P.-W., & Su, W.-J. (2022). Injection Barrel/Nozzle/Mold-Cavity Scientific Real-Time Sensing and Molding Quality Monitoring for Different Polymer-Material Processes. *Sensors*, *22*(13), 4792. doi: 10.3390/s22134792

Muaz, M., Sajid, S., Schulze, T., Liu, C., Klasen, N., & Drescher, B. (2025). Explainable AI for correct root cause analysis of product quality in injection moulding. *Journal of Manufacturing Processes*, *145*, 371–380. doi: 10.1016/j.jmapro.2025.03.114

Párizs, R. D., Török, D., Ageyeva, T., & Kovács, J. G. (2022). Machine Learning in Injection Molding: An Industry 4.0 Method of Quality Prediction. *Sensors*, *22*(7), 2704. doi: 10.3390/s22072704

Párizs, R. D., Török, D., Ageyeva, T., & Kovács, J. G. (2023). Multiple In-Mold Sensors for Quality and Process Control in Injection Molding. *Sensors*, *23*(3), 1735. doi: 10.3390/s23031735

Presciuttini, A., Cantini, A., Cramer, S., Huber, M., Wolfschläger, D., Schmitt, R. H., & Portioli-Staudacher, A. (2025). An Integrated Framework for Predictive Quality in Injection Molding: Combining Explainable AI and Time Series Analysis in a German Industry Case Study. *IFAC-PapersOnLine*, *59*(10), 1677–1682. doi: 10.1016/j.ifacol.2025.09.282

Speranza, V., Vietri, U., & Pantani, R. (2013). Monitoring of Injection Moulding of Thermoplastics: Adopting Pressure Transducers to Estimate the Solidification History and the Shrinkage of Moulded Parts. *Strojniški vestnik – Journal of Mechanical Engineering*, *59*(11), 677–682. doi: 10.5545/sv-jme.2013.1000

Wang, J.-C., Chang, C.-T., & Ke, K.-C. (2026). Data-driven representation learning and stability validation for multi-quality prediction in smart injection molding. *The International Journal of Advanced Manufacturing Technology*. doi: 10.1007/s00170-026-17655-z

Wang, Q., Zhao, X., Zhang, J., Zhang, P., Wang, X., Yang, C., … Wu, Z. (2021). Research on Quality Characterization Method of Micro-Injection Products Based on Cavity Pressure. *Polymers*, *13*(16), 2755. doi: 10.3390/polym13162755

Yan, P., Abdulkadir, A., Aguzzi, G., Schatte, G. A., Grewe, B. F., & Stadelmann, T. (2024). Automated process monitoring in injection molding via representation learning and setpoint regression. In *2024 11th IEEE Swiss Conference on Data Science (SDS)* (pp. 138–145). doi: 10.1109/SDS60720.2024.00027

Zhao, N.-Y., Liu, J.-F., Su, M.-Y., & Xu, Z.-B. (2024). Measurement techniques in injection molding: A comprehensive review of machine status detection, molten resin flow state characterization, and component quality adjustment. *Measurement*, *226*, 114163. doi: 10.1016/j.measurement.2024.114163

BIOGRAPHIES

Yulin Wang is a Ph.D. researcher in the Department of Mechanical Engineering at KU Leuven, Belgium. He is a researcher in the MPRO Core Lab (Motion Products) within Flanders Make. He received the M.Sc. degree in Mechanical Engineering from RWTH Aachen University, Germany, in 2024. His research interests include data-driven modeling, process monitoring, quality prediction and explainable artificial intelligence for advanced manufacturing systems.

Elke Deckers is Associate Professor in the Faculty of Engineering Technology at KU Leuven, Belgium, and leads the DMMS-D Core Lab (Dynamics of Mechanical and Mechatronic Systems) within Flanders Make. She received the M.Sc. and Ph.D. degrees in Mechanical Engineering from KU Leuven, Belgium, in 2008 and 2012, respectively. Her research interests include vibro-acoustic engineering, structural dynamics, linking polymer product design to product performance while including manufacturing aspects, with applications in vibration control, fault diagnosis and lightweight material design.

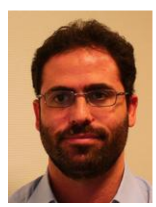

Konstantinos Gryllias received his Diploma and Ph.D. degrees in Mechanical Engineering from National Technical University of Athens, Athens, Greece, in 2004 and 2010, respectively. He is currently Professor of Vibro-acoustics of machines and transportation systems at the Department of Mechanical Engineering, KU Leuven, Leuven, Belgium. He also serves as the Manager of the University Core Lab Flanders Make@KU Leuven Motion Products, Belgium and Associate Editor of Mechanical Systems and Signal Processing. His research interests include condition monitoring, signal processing, prognostics, health management of mechanical and mechatronic systems, 3d printing and injection molding process monitoring.