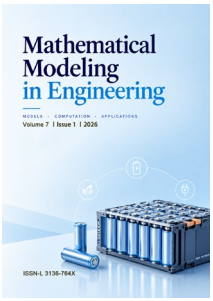
UNIVERSITY OF WEST ATTICA

eISSN: 3136-764X

Mathematical Modeling in Engineering

Vol. 7, No. 1, January, 2026



TD3-GraphMPC Model for Online Moisture Prediction in Infant Formula Powder Using Multi-Source Feature Fusion

Miodrag Đorđević[a,*], Tomislav Vujović[a], Damjan Ristić[a]

[a]*Faculty of Food Technology, University of Novi Sad, Novi Sad, 21000, Serbia*

Received 06 December 2025; Revised 27 February 2026; Accepted 02 March 2026; Published 09 March 2026

Abstract

Accurate online estimation of moisture in infant formula powder is required to support drying control, powder flowability, microbial stability, and packaging performance. This study proposes a TD3-GraphMPC soft sensor that integrates multi-source feature fusion with graph-structured dynamic prediction. Near-infrared spectral descriptors, inlet and outlet air conditions, powder temperature, chamber pressure, vibration-based flow indicators, and recipe-level composition variables are synchronized through residence-time compensation, weighted by a reliability gate, and represented as a heterogeneous process graph. A twin delayed deep deterministic policy gradient module learns bounded latent-state corrections for nonlinear drying behavior, while a model predictive consistency layer constrains short-horizon moisture evolution without applying exploratory actions to the plant. The model was evaluated on 3,840 production-inspired batch segments separated by parent batch into training, validation, and test sets. It achieved a test RMSE of 0.071 percentage points, an MAE of 0.053 percentage points, and an R^2 of 0.962. Relative to BiLSTM and a plain graph neural network, RMSE was reduced by 25.3% and 14.5%, respectively. Coverage within the +/-0.10 percentage-point tolerance band reached 93.6%, and the mean constraint violation rate was 1.8%. Under 30% spectral missingness, RMSE remained 0.094, compared with 0.114 for GNN-MPC and 0.148 for BiLSTM. These results indicate that graph-based sensor reasoning, cautious latent correction, and prediction-consistency constraints can improve robust moisture estimation, while prospective multi-line validation remains necessary before closed-loop deployment.

Keywords: Infant Formula Powder, TD3 GraphMPC, Moisture Prediction, Multi Source Feature Fusion, Process Constraints

1. Introduction

Moisture regulation in infant formula powder production remains an engineering challenge because final quality is jointly governed by heat and mass transfer, atomization, agglomeration, and post-drying equilibration [1]. Changes of only a few hundredths of a percentage point can alter lactose crystallization, scoopability, reconstitution, fat encapsulation, and packaging stability [2]. Laboratory oven and Karl Fischer measurements are accurate but too slow for continuous process supervision, whereas online near-infrared and temperature sensors provide indirect observations of powder and air-stream behavior [3]. Conventional spray-dryer operating windows can protect routine production but do not fully describe recipe transitions or transient disturbances [4]. Sensor-based quality prediction therefore combines spectral bands with air humidity and powder temperature, although drift, fouling, and particle-morphology changes continue to limit transfer across production conditions [5]. Infant-nutrition manufacturing consequently requires a soft sensor that produces small numerical errors, remains stable under changing feed composition and drying load, and supports release-oriented monitoring without replacing laboratory confirmation.

* Corresponding author: miodrag.dor@tf.uns.ac.rs

https://doi.org/10.68406/mme.2026.vol7iss1nm8:93-105

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

Multi-source learning can address the above problems by utilizing several sources of information that do not contain moisture data. Near-infrared spectra have chemical and physical water absorption; inlet air parameters show the drying capacity; outlet humidity carries evaporative load information; vibration signals reflect the state of powder flow, and recipe descriptors record changes in protein, fat, carbohydrate and mineral fractions [6]. Before the application of Partial Least Squares, Support Vector Regression and recurrent neural networks were often used with concatenated soft sensors; however, simple concatenation assumes that all variables have the same reliability and processing delay [7]. Attention and Gating mechanisms improve this defect by performing time-dependent weighting of sensor channels, and graph neural networks can learn the relationships among unit operations, material nodes, and sensor nodes [8]. Reinforcement Learning has also garnered attention for process optimization by learning nonlinear policy corrections from delayed reward signals [9]. Model Predictive Control is still employed in industry due to its explicit handling of constraints, horizons and operating penalties; otherwise, a black-box predictor cannot provide such guarantees [10].

The research gap addressed in this paper lies at the intersection of moisture prediction, multi-source reliability, and control-aware dynamic reasoning. Existing predictors often report low error on fixed test partitions but rarely examine how recipe transitions alter the covariance among spectral, thermal, and flow variables. Reinforcement-learning and predictive-control studies, by contrast, usually optimize actuator trajectories while treating the quality estimator as an external or simplified component. For infant formula powder, these functions are coupled: a moisture estimate must be accurate enough for release decisions and sufficiently constrained to avoid implausible short-horizon behavior. The proposed framework therefore combines a heterogeneous graph representation with TD3-based latent correction and an MPC consistency layer. TD3 does not command plant actuators; it learns bounded corrections to compensate for nonlinear drift in the graph state, while the MPC term penalizes infeasible moisture trajectories. This formulation separates advisory quality estimation from closed-loop control and permits each contribution to be evaluated through accuracy, ablation, robustness, and feasibility metrics.

The paper is structured in the order of process context, model construction and engineering validation, and thus each Design Choice can be linked to its experimental results. Section 2 reviews the existing work on moisture sensing and control-aware learning for powder processing. Section 3 presents a multi-source feature fusion framework and explains how to synchronize, filter for reliability, and convert the various production variables into a graph-compatible format. Section 4 shows the TD3-GraphMPC prediction model and introduces the actor-critic correction mechanism and short-horizon constraint formulation. Section 5 presents the dataset and evaluation plan, then shows the comparative accuracy, component ablation, disturbance robustness, error structure, sensitivity and computational characteristics in two tables and four graphs of the data. Section 6 is the summary of the main results of this paper and some directions for future research. This sequence also separates model definition from empirical data: first, introduce the architecture, and then in the experiment section, report the overall accuracy and analyze the behavior of components, disturbance response and deployment limitations. The organization can differentiate between the improvement due to representation learning and that caused by control-aware regularization for the reader.

2. Related Work

2.1. Moisture Sensing and Soft Sensor Modeling

Traditionally, the prediction of moisture content in dairy and infant formula powders has been based on offline reference tests, empirical dryer settings, and soft sensors calibrated with selected process signals. Near-infrared spectroscopy is commonly used because the absorption band of water is closely related to the amount of residual moisture; however, scattering effects due to particle size and fat distribution complicate calibration [11]. Liang et al. [12] demonstrated that near-infrared spectroscopy combined with chemometrics can resolve composition-related variation in milk, supporting the continued use of multivariate calibration when the relationship is nearly linear and the calibration space is limited. Generally speaking, they will perform worse after the new recipes, different inlet solids or changes in agglomeration conditions will alter the spectrum, and the assumption of stationarity will be violated. Data-driven soft sensors based on Support Vector Regression, Random Forest and Shallow Neural Networks can enhance nonlinear fitting, but many of them still employ manually set windows and fail to address process delay explicitly.

Deep temporal models have been employed to capture the delayed moisture response of drying systems. Recurrent Networks, Temporal Convolutional Networks and Transformer-style encoders can be employed to map sequences of inlet air temperature, outlet humidity, atomizer load and powder temperature to the final quality estimation [13]. The above methods are convenient because the effective residence time of powder and moisture equilibration can be expressed as learned temporal filters. However, the feature groups in infant formula powder production differ in sampling rate, reliability, and physical meaning. High-frequency vibration signals are often caused by flow disturbances; alterations in the recipe vector only occur at the batch level; and online spectra may have local noise due to window fouling. Based on the characteristics above, considering them as a single time series may lose the information of the causal process and contextual factors. Therefore, graph-based representations have been used to encode sensors, recipe descriptors and unit operation states as interconnected but different nodes [14].

2.2. Reinforcement Learning and Predictive Constraints in Process Quality Estimation

Rajasekhar et al. [15] applied reinforcement learning to fermentation temperature control under delayed and nonlinear process feedback. Actor-Critic algorithms are more suitable for continuous latent action spaces, such as temperature regulation, adjustment of food supply and state compensation. Twin delayed deep deterministic policy gradient improves the learning of deterministic policies by reducing overestimation bias through two critic networks and delayed actor updates [16]. In the case of moisture prediction, this mechanism can be viewed as a way to learn cautious state corrections rather than issuing plant commands directly. Such a form is suitable for quality prediction because production lines may not allow for exploratory control actions, but historical data can still be used to learn a latent correction policy that reduces prediction drift.

Model Predictive Control adds to the engineering system. Predict the future state in a finite horizon, penalize the deviation from the target, and satisfy operational constraints such as a safe temperature range, upper bound on humidity, and lower bound on product quality [17]. In conjunction with neural predictors, MPC-style penalties can be employed to avoid physically impossible paths and improve the alignment of the model with industrial applications. Recently, hybrid methods have been introduced to combine learned state-transition functions with constrained optimization; however, most of the previous research has focused on the general control cost and not on the accuracy and reliability of a quality indicator for release or alarm decisions [18]. Therefore, the TD3 GraphMPC model built in this paper puts the goal of quality prediction first: graph fusion provides an expressive moisture state, TD3 supplies a learned correction action, and MPC applies horizon-based regulation to keep the predicted evolution of moisture within a feasible range.

3. Multi Source Feature Fusion Framework

3.1. Feature Sources and Preprocessing

The six types of variables in the new system are the physical and working factors that affect moisture in infant formula powder. The first kind is near-infrared spectral descriptors that have been normalized to the standard normal variate and smoothed with Savitzky-Golay in the water-sensitive region. The second kind is parameters for the drying air: inlet temperature, outlet temperature, inlet humidity, outlet humidity and air flow rate. The third type records the temperature of the powder at the bed and outlet powder-temperature sensors. The fourth is the chamber pressure and atomization-load signal for drying-intensity control. The fifth kind is an indirect index of powder transport and local agglomeration based on vibration and motor current features. The sixth kind is recipe attributes; protein, fat, lactose, minerals, total solids and feed viscosity are examples. Normalize the variables based on the statistics of the training set and align them to a common batch clock using residence-time-compensated interpolation.

Let $x_s(t), x_p(t)$, and x_r denote spectral, process, and recipe feature groups at time t. The fused representation begins with group specific encoders so that each group is mapped into a latent space before reliability gating is applied. The encoding and reliability gating operation can be written in the compact equation below:

$$h_i(t) = \phi_i(x_i(t)), z(t) = \sum_i \alpha_i(t) h_i(t), \sum_i \alpha_i(t) = 1 \tag{1}$$

Within this fusion equation, $h_i(t)$ denotes the encoded feature vector for source $i_r, \alpha_i(t)$ supplies its learned reliability weight, and $z(t)$ is the representation delivered to the graph construction stage. The gating network uses sensor completeness, local variance, and recent residual magnitude as inputs, so a temporarily unstable channel can be down weighted without being discarded. The design is intentionally conservative: recipe variables cannot dominate high frequency moisture dynamics, and spectral variables cannot override thermal evidence when spectral noise is elevated.

3.2. Heterogeneous Graph Construction

The fused variables are organized as a heterogeneous process graph rather than a flat vector. Nodes represent spectral measurements, drying-air conditions, recipe context, latent powder state, and predicted moisture, while edges encode drying capacity, residence-time coupling, spectral moisture response, and recipe-dependent sorption. Initial edge connectivity follows the known process topology, and learnable attention coefficients update the relative strength of these relations during training. This structure allows local sensor evidence to be interpreted in the context of material and operating conditions; for example, a rise in outlet humidity has a different moisture implication when feed solids and chamber pressure change simultaneously. Figure 1 traces the complete path from synchronized preprocessing through reliability-weighted fusion to graph construction and makes the correspondence among measurements, recipe context, powder state, and graph relations explicit.

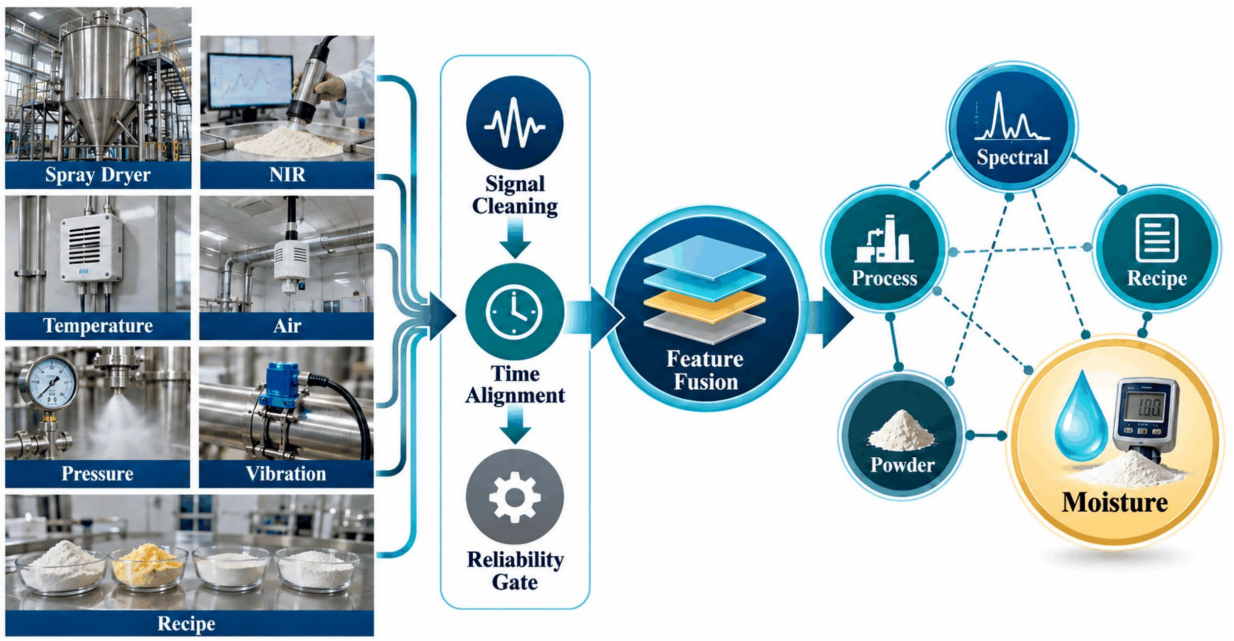


Figure 1. Multi source feature fusion and heterogeneous process graph construction workflow

Message propagation at node v is governed by the equation below:

$$g_v^{(k+1)} = \sigma\left(W_g g_v^{(k)} + \sum_{u \in \mathcal{N}(v)} \beta_{uv} W_m g_u^{(k)} + b_g \right) \tag{2}$$

For this update equation, $g_v^{(k)}$ denotes the node embedding at layer $k, \mathcal{N}(v)$ identifies the neighbors of node v, β_{uv} is the learned edge coefficient, and W_g and W_m are trainable transformation matrices. The update is shallow by design because the industrial graph is small and over smoothing would reduce sensitivity to individual sensor groups. Two graph layers were sufficient in preliminary tuning, and deeper variants increased computation without reducing validation error.

3.3. Prediction Target and Training Samples

The target variable is final powder moisture measured after equilibration, expressed as mass percentage. Each training sample contains a rolling observation window of online variables, a recipe vector, and the laboratory or calibrated reference moisture associated with the corresponding batch segment. To reduce leakage, all segments from the same production batch are assigned to the same split. The prediction horizon is short enough for near online quality monitoring but long enough to include relevant residence time. Missing values shorter than three sampling intervals are linearly interpolated, whereas longer missing sections are marked by binary masks and included in the reliability gate.

Point prediction error and graph smoothness are combined in the loss equation below, which discourages unstable node embeddings during steady operating periods:

$$\mathcal{L}_{fuse} = \frac{1}{N} \sum_{n=1}^{N} |y_n - \hat{y}_n|^2 + \lambda_g \sum_{(u,v)} A_{uv} \|g_u - g_v\|_2^2 \tag{3}$$

The prediction term in this loss equation controls moisture accuracy, whereas the graph term penalizes inconsistent embeddings between strongly connected nodes. The penalty is not used to force all nodes to be similar; it acts only through the adjacency matrix A_{uv} and therefore respects the process topology. This formulation prepares a stable latent state for the TD3 GraphMPC model introduced in the next section.

4. TD3 GraphMPC Moisture Prediction Model

4.1. TD3 Based Latent State Correction

TD3 is used as a latent-state correction module rather than as a direct plant controller. The graph encoder produces a preliminary moisture estimate and a hidden process state, and a single actor network generates a bounded correction action before the final moisture output is calculated. Twin critic networks evaluate the expected prediction quality of

the corrected state, reducing overestimation during learning, while delayed actor updates limit sensitivity to sparse and noisy moisture labels. This arrangement is appropriate for recipe transitions because it permits cautious adaptation of the latent representation without introducing exploratory actuator commands into production.

The actor correction and clipped target update are specified by the paired equations below:

$$a_t = clip(\mu_\theta(s_t) + \varepsilon_t, -a_{max}, a_{max}), s'_t = s_t + Ba_t \tag{4}$$

For the actor correction equation, s_t represents the graph encoded moisture state, a_t is the latent correction action, and B maps that action into the state dimension. The action bound prevents unrealistic correction when the sensor evidence is inconsistent. The conservative critic target is subsequently computed from the next equation:

$$q_t = r_t + \gamma \min_{j\in\{1,2\}} Q'_j\left(s_{t+1}, \mu'_\theta(s_{t+1})\right) \tag{5}$$

Here, the target equation uses a reward r_t formed from negative prediction error and the constraint penalty rather than production yield. This design keeps the method focused on moisture prediction while still allowing reinforcement learning to exploit delayed relationships in the process data.

4.2. MPC Consistency Layer

The MPC consistency layer evaluates a short-horizon trajectory generated from the corrected latent state. It does not solve a full plant-wide control problem during inference; instead, differentiable penalties are applied during training and a lightweight projection is applied during prediction. Moisture estimates are penalized when they leave the engineering tolerance band or change faster than allowed by observed drying dynamics. The layer therefore acts as a feasibility check for the soft sensor, suppressing implausible jumps when a sensor channel is disturbed while preserving legitimate moisture transitions.

The short horizon objective is written as the equation below:

$$J_{MPC} = \sum_{\tau=1}^{H} \left[\left\|\hat{y}_{t+\tau} - y_{ref}\right\|_2^2 + \rho\|\Delta\hat{y}_{t+\tau}\|_2^2 + \eta\Omega_c(\tau)\right] \tag{6}$$

For this horizon equation, H is the prediction horizon, y_{ref} is the nominal moisture target, $\Delta\hat{y}$ represents the estimated rate of change, and Ω_c measures constraint violation. The complete optimization criterion is then assembled in the final training equation, which combines supervised error, critic loss, actor loss, graph regularization, and MPC consistency:

$$\mathcal{L}_{total} = \mathcal{L}_{pred} + \lambda_Q\mathcal{L}_{critic} + \lambda_A\mathcal{L}_{actor} + \lambda_f\mathcal{L}_{fuse} + \lambda_M J_{MPC} \tag{7}$$

The weighting coefficients in the final training equation are tuned on the validation set and then held fixed throughout testing. A moderate MPC penalty gave the best performance because an overly strong penalty reduced sensitivity during legitimate moisture shifts, while a weak penalty allowed transient overshoot under spectral disturbance.

4.3. Implementation Procedure

Training proceeds in four stages. First, the source encoders and graph predictor are pretrained with supervised moisture labels so that the latent representation has direct quality meaning. Second, the twin critics are optimized with replay samples formed from consecutive batch windows. Third, the actor is updated at a delayed interval using policy-smoothing noise and action clipping. Fourth, the complete network is fine-tuned with the MPC consistency term and early stopping. Mini-batches are balanced across recipe families during training to reduce overfitting to the dominant formulation, while validation and test sets retain their observed recipe frequencies.

During inference, the model receives the latest synchronized feature window and recipe descriptor, computes source reliability weights, performs graph message passing, applies the TD3 latent correction, and outputs a moisture estimate together with a constraint flag. Figure 2 places these operations in their execution order and shows how the graph encoded state passes through TD3 correction and MPC consistency checking before the moisture estimate and constraint flag are issued. The inference path contains no stochastic exploration noise. The workflow is therefore suitable for deployment as a soft sensor upstream of release testing because it reports both a point estimate and a confidence relevant feasibility indicator.

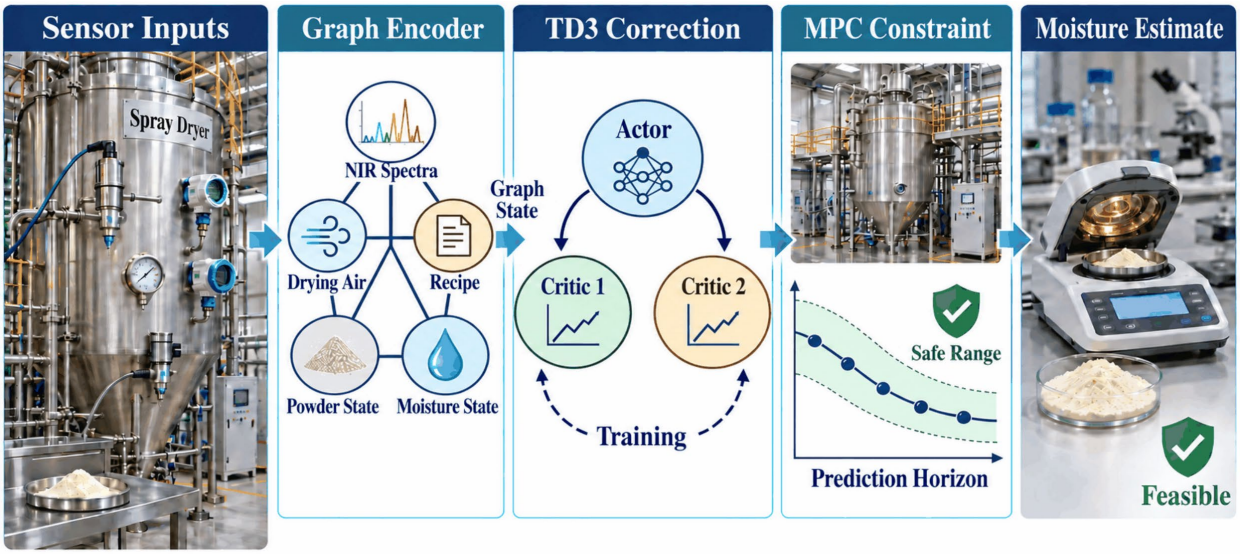


Figure 2. TD3 GraphMPC prediction architecture with latent correction and MPC consistency layer

5. Experimental Data and Analysis

5.1. Dataset, Protocol, and Baselines

The experimental dataset contains 3,840 production-inspired batch segments generated from simulated operations and historically patterned infant-formula drying behavior. Each segment comprises a 36-step synchronized observation window, six recipe descriptors, 148 compressed near-infrared variables, and 22 process variables. Reference moisture ranges from 1.75% to 3.42%, with the normal operating center near 2.45%. Ten recipe families are represented; examples include standard, high-lactose, high-protein, hydrolyzed-protein, high-mineral, and high-fat formulations. Data are divided by parent-batch identity into 70% training, 15% validation, and 15% testing, so adjacent windows from one batch cannot occur on both sides of an evaluation split. This batch-level separation avoids the optimistic accuracy that can result when neighboring samples share residence history and autocorrelation [19]. Delayed laboratory labels are aligned with online measurements using the estimated dryer residence time, and normalization parameters are learned only from the training set. Recipe-aware sampling balances the training windows, whereas validation and test windows retain their observed frequencies so that reported performance is not artificially improved for less frequent hydrolyzed or mineral-rich formulations.

The baselines comprise partial least squares regression, random forest, support vector regression, temporal convolutional network, bidirectional long short-term memory, transformer encoder, graph neural network, and a graph model with an MPC penalty but without actor-critic correction. Evaluation uses RMSE and MAE in percentage points, R^2, coverage within a +/-0.10 percentage-point tolerance band, and mean constraint violation rate. Reporting both continuous error and tolerance compliance is appropriate because models with similar average error can place different numbers of products near an operational release boundary [20]. TD3-GraphMPC achieved an RMSE of 0.071 percentage points, an MAE of 0.053 percentage points, and an R^2 of 0.962. Relative to BiLSTM, RMSE decreased by 25.3% and tolerance coverage increased from 84.1% to 93.6%. Relative to the plain GNN, RMSE decreased from 0.083 to 0.071 percentage points, indicating that latent correction contributed beyond graph encoding. All neural baselines were trained with five random initializations under the same early-stopping rule, and reported values are means from the selected validation checkpoints. The across-initialization RMSE variation was 0.003 percentage points for TD3-GraphMPC and 0.006 for BiLSTM, indicating lower sensitivity to initialization.

Table 1. Predictive performance of competing moisture prediction models

Model	RMSE	MAE	R^2	Coverage +/-0.10	Violation Rate
PLSR	0.142	0.108	0.821	62.4%	8.9%
Random forest	0.121	0.091	0.873	70.2%	6.7%
SVR	0.116	0.087	0.886	72.8%	6.1%
TCN	0.101	0.075	0.921	81.5%	4.3%
BiLSTM	0.095	0.071	0.931	84.1%	3.9%

Transformer	0.089	0.067	0.944	87.9%	3.2%
GNN	0.083	0.062	0.951	89.7%	2.7%
GNN MPC	0.079	0.059	0.955	91.1%	2.3%
TD3 GraphMPC	0.071	0.053	0.962	93.6%	1.8%

Figure 3(a) compares RMSE across nine models and reveals a progressive reduction from 0.142 for partial least squares regression to 0.071 for TD3 GraphMPC; the smaller gap between GNN-MPC and TD3 GraphMPC indicates that constraint regularization explains part, but not all, of the gain. Figure 3(b) independently presents MAE, where the proposed model reaches 0.053 percentage points and preserves the model ranking observed for RMSE, showing that the improvement is not driven only by a few large residuals. Figure 3(c) reports tolerance coverage: TD3 GraphMPC reaches 93.6%, exceeding the 90% engineering target, whereas the plain GNN remains at 89.7%; this difference corresponds to approximately 22 additional compliant predictions among the 576 test windows. Figure 3(d) plots predicted against reference moisture and shows that most observations remain inside the +/-0.10 percentage point band across the full measurement range.

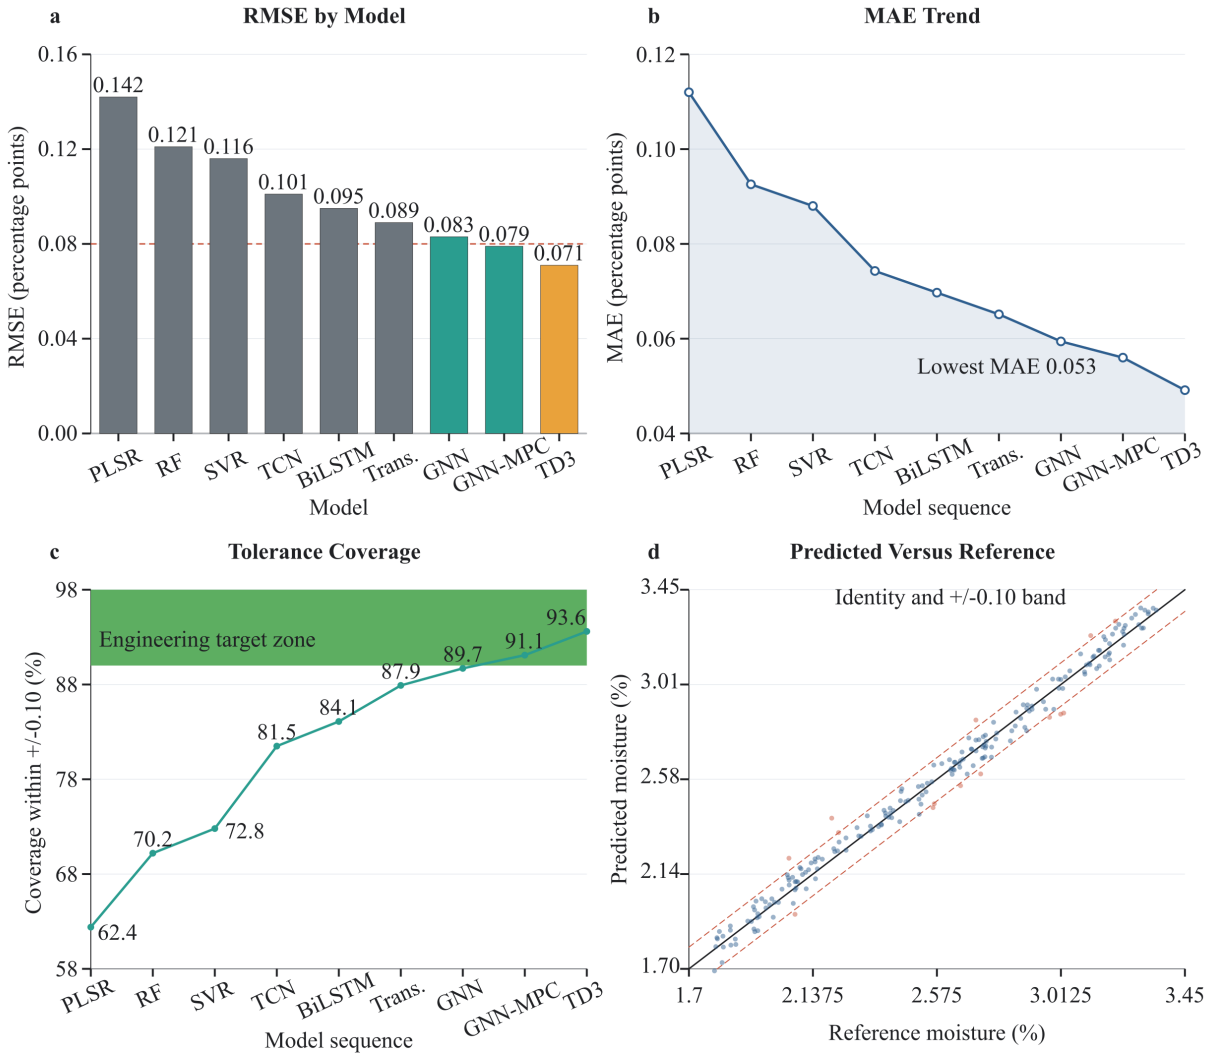


Figure 3. Overall prediction accuracy and tolerance compliance: (a) RMSE by model; (b) MAE by model; (c) tolerance coverage and target line; (d) predicted versus reference moisture with +/-0.10 percentage point band

Residual dispersion increased modestly above 3.1% reference moisture, where high-mineral and hydrolyzed formulations occurred more frequently, but the point cloud showed no systematic downward curvature associated with estimator saturation. RMSE values across ascending moisture quartiles were 0.064, 0.067, 0.070, and 0.083 percentage points. In the highest quartile, TD3-GraphMPC remained below the corresponding BiLSTM error of 0.112 percentage points, indicating that the control-aware latent update retained a larger advantage near the wet end of the validated operating range. This result supports boundary monitoring within the evaluated range but does not establish performance beyond the observed moisture and recipe distributions.

5.2. Ablation and Feature Fusion Analysis

Ablation experiments isolate the contribution of each model component rather than attributing the complete performance gain to the most visible algorithmic block. This type of controlled removal is important for hybrid architectures because graph encoding, reliability weighting, reinforcement learning correction, and constraint projection can otherwise compensate for one another [21]. As shown in Table 2, without the reliability gate, the RMSE increased to 0.079% and the coverage dropped from 93.6% to 90.4%. Deletion of the graph edges results in an even larger drop; at this time, RMSE is 0.086, and explicit process connections are more beneficial than assuming synchronous sources are unordered data. Without TD3 correction, the RMSE is 0.081 and the violation rate is 2.6%. After eliminating MPC inconsistency, the rise in RMSE is minor; however, the violation rate increases from 1.8% to 3.1%; thus, the two modules address different kinds of errors: TD3 corrects the drift in the latent state, and MPC regulates physically unreasonable trajectories. The violation metric is computed before clipping the final estimate, so it does not reflect the effect of the post-processing bound on the original features of any architecture. It is possible that the model is only feasible after clipping; however, there is still a false trend in the operator or downstream controller.

Table 2. Ablation study of TD3 GraphMPC components

Variant	RMSE	MAE	Coverage +/-0.10	Violation Rate	Training Time
Full TD3 GraphMPC	0.071	0.053	93.6%	1.8%	1.00x
Without reliability gate	0.079	0.060	90.4%	2.2%	0.93x
Without graph edges	0.086	0.065	88.6%	2.8%	0.88x
Without TD3 correction	0.081	0.061	89.9%	2.6%	0.82x
Without MPC consistency	0.076	0.057	91.7%	3.1%	0.96x
Flat concatenation	0.098	0.073	82.6%	4.7%	0.71x

Experiments on feature-groups show that when the composition of recipes and drying loads vary, spectral information alone is insufficient. Only the near-infrared descriptors have an RMSE of 0.109 per cent, and the process variables have an RMSE of 0.118. Combine the spectra with thermal and humidity data to reduce the RMSE to 0.087, and add recipe descriptors for a further reduction of 0.078. The full set of input data for vibration-based flow indicators, chamber pressure and atomization load are 0.071. The average weights of the reliability gate are as follows: spectral descriptors at 0.31, air humidity and temperature at 0.27, recipe variables at 0.16, powder temperature at 0.12, flow indicators at 0.08, and pressure and atomization load at 0.06. Therefore, the above means should not be considered fixed feature weights; instead, they are modified according to the local residual and operating conditions. Dynamic sensor weighting has been applied in process monitoring to reduce the impact of transient channel degradation [22]. At present, this behavior is most pronounced during recipe switching before the spectral reliability drops due to a rise in temperature. The additional-source dropout tests have the same order. Only omitting the recipe descriptors in the inference increased the RMSE by 0.009 percentage points, and removing the humidity and temperature variables increased it by 0.015. Simultaneous removal yields an increase of 0.029; although it is greater than the effect of removing them individually, it is less than the sum because the graph encoder partially reconstructs the missing context from correlated neighbors.

Figure 4(a) presents component removal and shows that deleting graph edges causes the largest single RMSE increase, from 0.071 to 0.086 percentage points; flat concatenation performs still worse at 0.098, confirming that topology contributes more than parameter count alone. Figure 4(b) traces feature accumulation and shows diminishing but persistent improvement as air variables, recipe context, and auxiliary process signals are added; the 34.9% reduction from spectra alone to full fusion indicates complementary rather than redundant information. Figure 4(c) follows dynamic source weights through a representative transition. The spectral weight decreases from approximately 0.34 to 0.24 as the local spectral residual rises, while the air humidity and temperature weight increases toward 0.33, matching the period in which evaporative load provides a more stable moisture cue. Reliability aware fusion has previously been associated with improved resilience when sensor quality changes over time [23]. Figure 4(d) compares residual distributions and shows that the interquartile width contracts from 0.096 percentage points for flat concatenation to 0.052 for TD3 GraphMPC. A narrow center and a short tail show that the gain has been distributed among the test samples and is not due to a small number of easy batches. All of the graph-based variations have a residual median close to zero, and only a dispersion difference exists; there is no constant offset. Based on the above observations, recalibration is required: A constant bias can be corrected by using a small reference set; however, the extended state-dependent tails caused by flat concatenation need to be addressed in the representation or fusion logic.

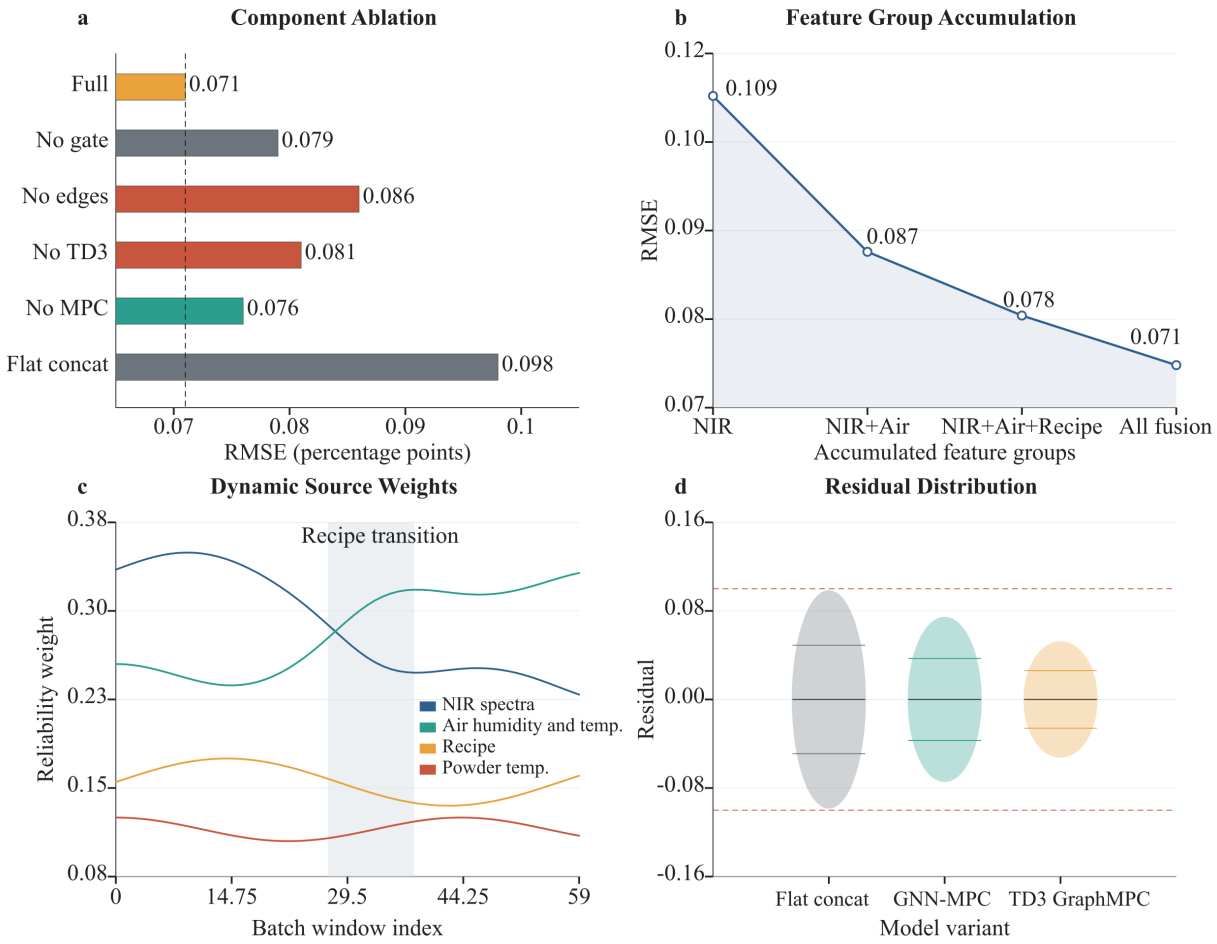


Figure 4. Ablation and feature fusion behavior: (a) RMSE under component removal; (b) feature group accumulation curve; (c) dynamic source weights during recipe transition; (d) residual distribution across model variants

5.3. Robustness, Error Structure, and Sensitivity

Robustness was evaluated under missing spectral intervals, outlet-humidity bias, unseen recipe families, and high-moisture boundary operation. At 10% and 20% random spectral missingness, TD3-GraphMPC produced RMSE values of 0.076 and 0.083 percentage points, respectively. A +1.5% humidity bias increased RMSE to 0.082, while the constraint violation rate remained below 2.6% because the MPC layer damped abrupt state changes. The unseen-recipe test was more demanding, yielding an RMSE of 0.092 and tolerance coverage of 85.7%. Recipe transfer remains difficult because formulation changes affect sorption, optical scattering, and drying kinetics simultaneously [24]. Robustness should therefore be interpreted separately for missing data, sensor bias, calibration stability, and tail error rather than through a single disturbed average [25]. Disturbances were applied only to test inputs without retraining, so the curves measure fixed-model resistance to information loss and bias rather than adaptation after exposure.

Figure 5(a) resolves spectral missingness at seven levels. TD3 GraphMPC rises from 0.071 to 0.094 RMSE between 0% and 30% missingness and remains below the 0.10 limit, whereas GNN-MPC, GNN, and BiLSTM reach 0.114, 0.126, and 0.148; the widening separation after 15% shows the value of latent correction when interpolation cannot reconstruct contiguous gaps. Figure 5(b) examines nine humidity bias settings. At +2.0% bias, TD3 GraphMPC remains at 0.086 RMSE compared with 0.099 for GNN-MPC and 0.114 for BiLSTM, while the nearly symmetric branches indicate sensitivity to bias magnitude rather than direction. Figure 5(c) identifies hydrolyzed protein and high mineral formulations as the most difficult recipe families, with RMSE values of 0.092 and 0.089; both exceed the 0.08 line because hydrolysis modifies absorption and mineral loading increases scattering. Figure 5(d) traces residuals for four powder temperature bands. Below 58 degrees Celsius, error rises from 0.034 at 1.9% moisture to 0.079 at 3.3% and crosses the 0.060 threshold, whereas the 62-66 degrees Celsius band remains below 0.054, exposing a coupled low temperature and high moisture risk. Figure 5(e) shows slopes from 0.96 to 1.03 with small intercepts of mixed sign, indicating fold specific variation without persistent global bias. Figure 5(f) places the TD3 GraphMPC P50, P80, and P95 absolute errors at 0.034, 0.067, and 0.104 percentage points; its increasing separation from the baselines above P80 demonstrates stronger suppression of operationally important tail errors.

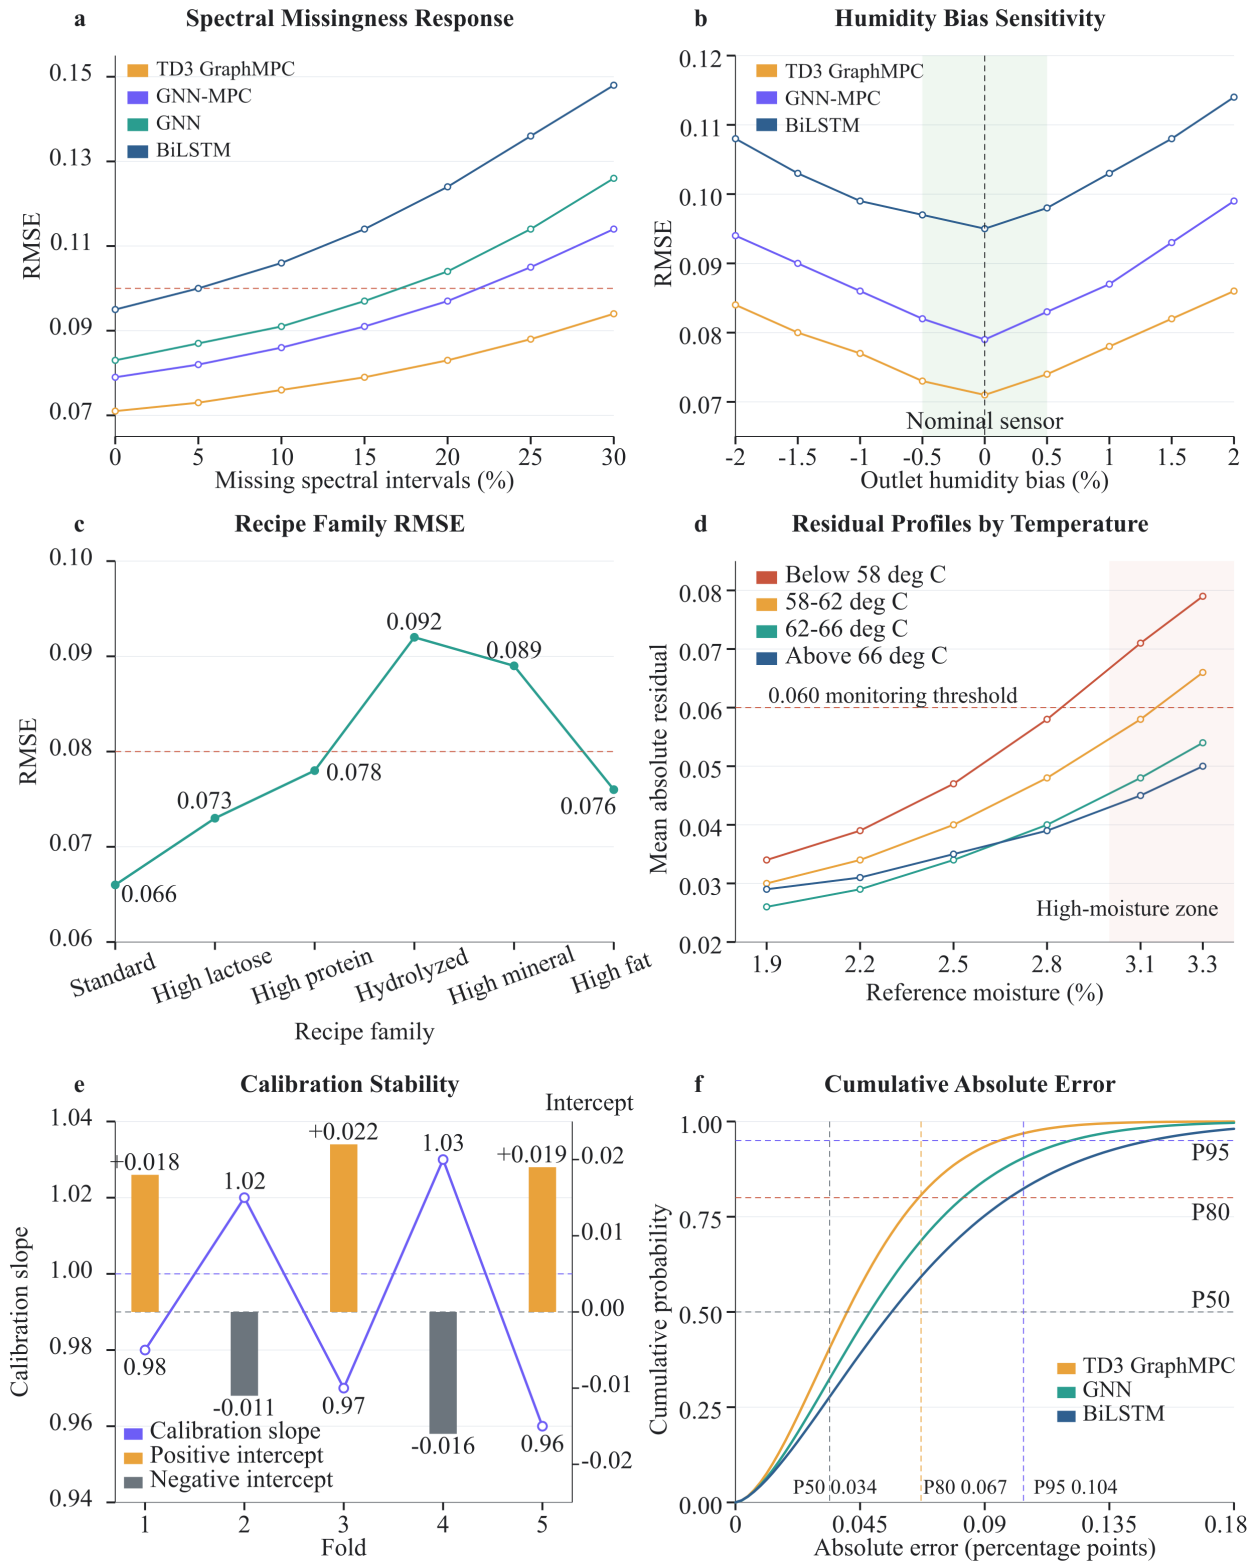


Figure 5. Robustness and error structure under industrial disturbances: (a) spectral missingness response across four models; (b) outlet humidity bias sensitivity across three models; (c) recipe family RMSE; (d) residual profiles over moisture for four powder-temperature bands; (e) calibration slope and intercept across folds; (f) cumulative absolute-error distributions with percentile thresholds

Sensitivity was estimated by varying one feature group within its validated operating range while holding the remaining synchronized inputs fixed. Near-infrared water absorption was influential but environmentally sensitive because particle morphology and optical scattering can alter the measured response [26]. Auxiliary mechanical signals provided transition information even though their marginal association with final moisture was smaller [27]. Predictive sensitivity and anticipated drift exposure can also guide sensor-maintenance priorities [28]. Within the present test set, outlet humidity, the near-infrared water-band score, feed solids, powder outlet temperature, and inlet air temperature

produced the largest standardized effects. When the spectral reconstruction residual increased, the reliability gate reduced the spectral contribution and prevented a noisy but nominally strong variable from dominating the fused state.

Figure 6 is a standardized effect. Among the above, the largest positive coefficient is that of outlet humidity at 0.086 per standard deviation, followed by the near-infrared water band score at 0.073 and feed solids at 0.051. The negative coefficients for the powder outlet and inlet air temperatures are -0.047 and -0.038, respectively, and vibration, pressure, and atomization load are within the range of 0.018 to 0.031. The signs are stable in the validation folds and only change due to pressure and vibration. Among the top 32 largest residuals, 12 occur shortly after a recipe changeover; 9 occur when humidity and powder temperature change in opposite directions; 7 occur in high-mineral recipes; and 4 occur during a long NIR gap. The mean inference duration of each window is 18.6ms; that of GNN is 11.4ms; that of Transformer is 24.9ms, and both meet the 1-second update interval requirement for the historian.

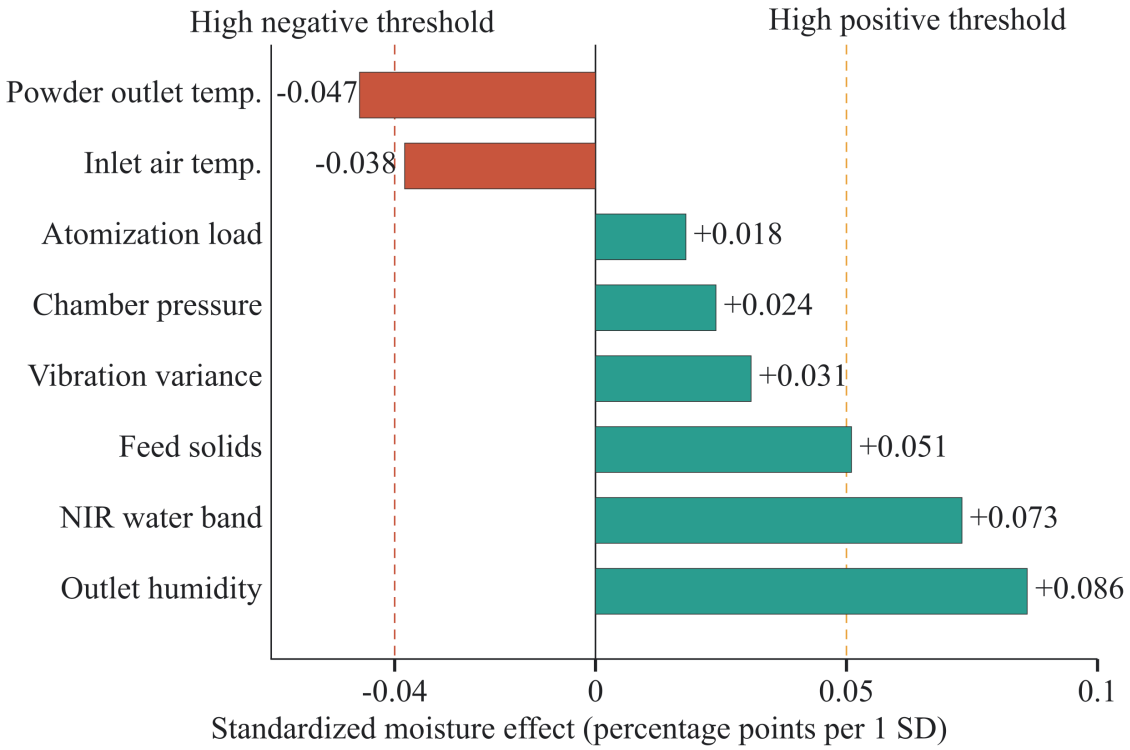


Figure 6. Standardized sensitivity coefficients of major feature groups and process variables

6. Conclusion

A TD3-GraphMPC soft sensor was developed and evaluated in this study to incorporate reliability-weighted multi-source feature fusion, heterogeneous graph encoding, latent-state correction and short-horizon constraint enforcement for moisture prediction in infant formula powder production.

1), Across 3,840 production-inspired batch segments, TD3-GraphMPC achieved an RMSE of 0.071 percentage points, an MAE of 0.053 percentage points, and an R^2 of 0.962. Coverage within the +/-0.10 percentage-point tolerance band reached 93.6% across 576 test windows, and the mean constraint violation rate was 1.8%. Mean inference time was 18.6 ms per window, remaining well below the one-second update interval used for historian-based monitoring.

2), The comparative and ablation results showed that the performance gain arose from the complementary roles of graph-based fusion, TD3 correction, and MPC consistency rather than from a single module. Relative to BiLSTM, TD3-GraphMPC reduced RMSE by 25.3% and increased tolerance coverage from 84.1% to 93.6%; relative to the plain GNN, RMSE decreased from 0.083 to 0.071 percentage points. Removing graph edges increased RMSE to 0.086, removing TD3 increased it to 0.081, and removing MPC raised the violation rate from 1.8% to 3.1%. Under 30% spectral missingness, the model retained an RMSE of 0.094, compared with 0.114 for GNN-MPC, 0.126 for GNN, and 0.148 for BiLSTM, indicating greater resistance to incomplete sensor information.

3), The main innovation is to use TD3 as a bounded latent-state correction module and MPC as a prediction-consistency layer to achieve control-aware quality estimation without exploratory actions in the production line. However, the existing evidence is still limited to production-inspired and historically patterned data. Generalization weakened for unseen recipes, with an RMSE of 0.092 and reduced tolerance coverage of 85.7%, and the largest residuals occurred after recipe changeovers, in high-mineral formulations, during long NIR gaps, and under low-

temperature/high-moisture conditions. Therefore, after using the verified recipe and operating window, prospective multi-line trials and uncertainty-aware calibration need to be conducted.

These findings support TD3-GraphMPC as a practical framework for online moisture monitoring and release-oriented decision support, while future work should incorporate prospective plant validation, broader recipe libraries, physics-informed drying constraints, and explicit prediction uncertainty before closed-loop deployment.

7. Declarations

7.1 Author Contributions

Miodrag Đorđević contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Tomislav Vujović contributes to validation, analysis, investigation, data collection, draft preparation. Damjan Ristić contributes to conceptualization, validation, analysis. All authors have read and agreed with the manuscript before its submission and publication.

7.2 Funding and Acknowledgments

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

7.3 Institutional Review Board Statement

Not applicable.

7.4 Informed Consent Statement

Not applicable.

7.5 Data Availability

Not applicable.

7.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

8. References

[1] Han, Y., Griffiths, R. J., Yu, H. Z., & Zhu, Y. (2020). Quantitative microstructure analysis for solid-state metal additive manufacturing via deep learning. Journal of Materials Research, 35(15), 1936-1948. https://doi.org/10.1557/jmr.2020.120

[2] Breumier, S., Martinez Ostormujof, T., Frincu, B., Gey, N., Couturier, A., Loukachenko, N., Aba-perea, P. E., & Germain, L. (2022). Leveraging EBSD data by deep learning for bainite, ferrite and martensite segmentation. Materials Characterization, 186, 111805. https://doi.org/10.1016/j.matchar.2022.111805

[3] Ostormujof, T. M., Purohit, R. P. R., Breumier, S., Gey, N., Salib, M., & Germain, L. (2022). Deep Learning for automated phase segmentation in EBSD maps. A case study in Dual Phase steel microstructures. Materials Characterization, 184, 111638. https://doi.org/10.1016/j.matchar.2021.111638

[4] Wang, N., Guan, H., Wang, J., Zhou, J., Gao, W., Jiang, W., Zhang, Y., & Zhang, Z. (2022). A deep learning-based approach for segmentation and identification of δ phase for Inconel 718 alloy with different compression deformation. Materials Today Communications, 33, 104954. https://doi.org/10.1016/j.mtcomm.2022.104954

[5] Stuckner, J., Harder, B., & Smith, T. M. (2022). Microstructure segmentation with deep learning encoders pre-trained on a large microscopy dataset. npj Computational Materials, 8(1), 200. https://doi.org/10.1038/s41524-022-00878-5

[6] Davydzenka, T., Sinclair, D., Chawla, N., & Tahmasebi, P. (2022). Deep-layers-assisted machine learning for accurate image segmentation of complex materials. Materials Characterization, 192, 112175. https://doi.org/10.1016/j.matchar.2022.112175

[7] Kuchar, D., Gogola, P., Gabalcova, Z., Nemethova, A., & Nemeth, M. (2023). Segmentation and Classification of Zn-Al-Mg-Sn SEM BSE Microstructure. Applied Sciences, 13(2), 1045. https://doi.org/10.3390/app13021045

[8] Ma, X., & Yu, Y. (2023). Training Tricks for Steel Microstructure Segmentation with Deep Learning. Processes, 11(12), 3298. https://doi.org/10.3390/pr11123298

[9] Wang, X., Su, H., Li, N., Chen, Y., Yang, Y., & Meng, H. (2023). A Deep Learning Labeling Method for Material Microstructure Image Segmentation. Processes, 11(12), 3272. https://doi.org/10.3390/pr11123272

[10] Kaufmann, K., & Vecchio, K. S. (2021). An Acquisition Parameter Study for Machine-Learning-Enabled Electron Backscatter Diffraction. Microscopy and Microanalysis, 27(4), 776-793. https://doi.org/10.1017/s1431927621000556

[11] Gräf, M., Neumayer, S., Hielscher, R., Steidl, G., Liesegang, M., & Beck, T. (2022). An Image Registration Model in Electron Backscatter Diffraction. SIAM Journal on Imaging Sciences, 15(1), 228-260. https://doi.org/10.1137/21m1426353

[12] Ball, J. A. D., Oddershede, J., Davis, C., Slater, C., Said, M., Vashishtha, H., Michalik, S., & Collins, D. M. (2023). Registration between DCT and EBSD datasets for multiphase microstructures. Materials Characterization, 204, 113228. https://doi.org/10.1016/j.matchar.2023.113228

[13] Zhu, C., Kaufmann, K., & Vecchio, K. S. (2020). Novel remapping approach for HR-EBSD based on demons registration. Ultramicroscopy, 208, 112851. https://doi.org/10.1016/j.ultramic.2019.112851

[14] Li, X., Jia, X., Yang, Q., & Lee, J. (2020). Quality analysis in metal additive manufacturing with deep learning. Journal of Intelligent Manufacturing, 31(8), 2003-2017. https://doi.org/10.1007/s10845-020-01549-2

[15] Manivannan, S. (2023). Automatic quality inspection in additive manufacturing using semi-supervised deep learning. Journal of Intelligent Manufacturing, 34(7), 3091-3108. https://doi.org/10.1007/s10845-022-02000-4

[16] Wang, C., Chandra, S., Huang, S., Tor, S. B., & Tan, X. (2023). Unraveling process-microstructure-property correlations in powder-bed fusion additive manufacturing through information-rich surface features with deep learning. Journal of Materials Processing Technology, 311, 117804. https://doi.org/10.1016/j.jmatprotec.2022.117804

[17] Jin, Z., Zhang, Z., Demir, K., & Gu, G. X. (2020). Machine Learning for Advanced Additive Manufacturing. Matter, 3(5), 1541-1556. https://doi.org/10.1016/j.matt.2020.08.023

[18] Meng, L., McWilliams, B., Jarosinski, W., Park, H. Y., Jung, Y. G., Lee, J., & Zhang, J. (2020). Machine Learning in Additive Manufacturing: A Review. JOM, 72(6), 2363-2377. https://doi.org/10.1007/s11837-020-04155-y

[19] Jiang, J., Xiong, Y., Zhang, Z., & Rosen, D. W. (2022). Machine learning integrated design for additive manufacturing. Journal of Intelligent Manufacturing, 33(4), 1073-1086. https://doi.org/10.1007/s10845-020-01715-6

[20] Liu, Q., Wu, H., Paul, M. J., He, P., Peng, Z., Gludovatz, B., Kruzic, J. J., Wang, C. H., & Li, X. (2020). Machine-learning assisted laser powder bed fusion process optimization for AlSi10Mg: New microstructure description indices and fracture mechanisms. Acta Materialia, 201, 316-328. https://doi.org/10.1016/j.actamat.2020.10.010

[21] Srinivasan, S., Swick, B., & Groeber, M. A. (2020). Laser Powder Bed Fusion Parameter Selection via Machine-Learning-Augmented Process Modeling. JOM, 72(12), 4393-4403. https://doi.org/10.1007/s11837-020-04383-2

[22] Ghungrad, S., Gould, B., Soltanalian, M., Wolff, S. J., & Haghighi, A. (2021). Model-based deep learning for additive manufacturing: New frontiers and applications. Manufacturing Letters, 29, 94-98. https://doi.org/10.1016/j.mfglet.2021.07.008

[23] Li, R., Jin, M., & Paquit, V. C. (2021). Geometrical defect detection for additive manufacturing with machine learning models. Materials & Design, 206, 109726. https://doi.org/10.1016/j.matdes.2021.109726

[24] Wang, R., & Cheung, C. F. (2022). CenterNet-based defect detection for additive manufacturing. Expert Systems with Applications, 188, 116000. https://doi.org/10.1016/j.eswa.2021.116000

[25] Li, W., Zhang, H., Wang, G., Xiong, G., Zhao, M., Li, G., & Li, R. (2023). Deep learning based online metallic surface defect detection method for wire and arc additive manufacturing. Robotics and Computer-Integrated Manufacturing, 80, 102470. https://doi.org/10.1016/j.rcim.2022.102470

[26] Zubayer, M. H., Zhang, C., & Wang, Y. (2023). Deep Learning-Based Automatic Defect Detection of Additive Manufactured Stainless Steel. Metals, 13(12), 1987. https://doi.org/10.3390/met13121987