Physics-informed digital twin for sucrose loss prevention and transparent governance in regional sugar beet supply chains

Corneliu Iatco
Academy of Economic Studies of Moldova, Ph.D. Candidate, Chisinau, Republic of Moldova
corneliu.iatco@gmail.com

Abstract

Objectives: This study aims to develop and validate an intelligent Digital Twin architecture that integrates Physics-Informed Neural Networks (PINN) for predictive sucrose degradation forecasting with fairness mechanisms for independent agricultural producers and transparent data governance in regional sugar beet supply chains. **Prior work**: Traditional logistics models treat agricultural cargo as static objects, ignoring the non-linear biochemical degradation of root crops during temporary field storage and the information asymmetry between producers and processing enterprises. While recent contributions in smart regional development have emphasised the importance of E-Governance and institutional transparency, existing quality forecasting and equity balancing models have evolved separately, resulting in fragmented decision-support systems. **Approach**: The proposed framework integrates three complementary modules: (i) a PINN embedding Arrhenius-based degradation kinetics for predictive sucrose loss estimation, trained on a historical industrial dataset and validated against the industrial baseline of constant daily losses. By resolving the quality observability gap through PINN-based hidden-state estimation and by reconciling delayed post-reception laboratory results with predictive estimates, the architecture ensures continuous adaptation without accumulating systemic drift error; (ii) a multi-period Growers' Equity balancing algorithm with historical compensation, designed to prevent systematic discrimination over a multi-year horizon; and (iii) an event-driven microservice architecture with an RFID-enabled transactional pipeline for end-to-end data transparency and macro-anomaly detection. **Results**: In a controlled proof-of-concept benchmark based on norm-derived historical degradation trajectories, the PINN reduces the Mean Absolute Error of sucrose concentration reconstruction by a factor of 12 compared to the simple averaging heuristic (0.031% vs. 0.383%). This indicates that explicit temperature-dependent kinetics substantially improves hidden-state reconstruction under sparse laboratory observations. **Implications**: By resolving the pre-reception quality observability gap through PINN-based hidden-state estimation and the post-reception analytical lag through retrospective reconciliation, the architecture ensures continuous adaptation without accumulating systemic drift error. For practitioners, this provides a transparent, empirically grounded decision-making environment; for researchers, it demonstrates that physics-informed machine learning can be successfully coupled with deterministic optimisation to enable proactive quality preservation and institutional fairness. **Value**: The scientific novelty lies in the integration of physics-informed machine learning with logistical optimisation and institutional fairness mechanisms, supported by empirical validation on an 11-year industrial dataset. By eliminating information asymmetry and creating a transparent cyber-physical space, the Digital Twin provides a robust foundation for E-Governance and sustainable development in regional agro-industrial clusters.

Keywords: Smart Governance; Biochemical Degradation; Stakeholder Fairness; Institutional Trust; Event-Driven Microservices.

1. Introduction

The digital transformation of agro-industrial supply chains extends far beyond the optimization of material flows. Within the contemporary discourse on smart regional development, the effectiveness of cyber-physical systems is increasingly measured not merely by operational efficiency gains, but by their capacity to ensure institutional transparency, protect stakeholder equity, and safeguard the intrinsic quality of biologically active raw materials [1, 2, 3]. Nowhere are these imperatives more pronounced than in the sugar beet sub-complex, where the non-linear biochemical degradation of root crops during temporary field storage, combined with the information asymmetry between independent

agricultural producers and processing enterprises, generates substantial economic losses that undermine the sustainability of entire regional clusters [4, 5, 6]. A logistics system that does not know the quality of the cargo it transports cannot intelligently prioritize which cargo to move first.

A critical yet often overlooked challenge in agro-industrial logistics is the temporal gap between the physical degradation of sugar beet piles and the availability of validated quality measurements. This gap has two distinct components. First, after harvesting and field pile formation, beet mass may remain in open storage for several days or even weeks before the first representative batch from that pile reaches the reception point and is sampled. During this pre-reception storage interval, the biochemical state of the pile evolves continuously, but no deterministic laboratory measurement is available. Second, after a truck arrives at the reception point, a post-reception laboratory lag arises: samples must be extracted, prepared, transported to the centralized laboratory, and analysed polarimetrically. In the industrial setting studied here, where several factories historically relied on a centralized laboratory workflow, this analytical delay may reach up to 48 hours. At the time of logistical planning, the system therefore does not possess deterministic quality data for many active field piles or recently received batches.

From a decision-theoretic perspective, this delay creates an interval during which the true biochemical state of the transported biomass is fundamentally unobservable. At the moment routing decisions must be made, neither laboratory measurements nor deterministic analytical models can provide the actual sucrose concentration of each field pile. Consequently, logistics optimization operates under structural information uncertainty rather than ordinary stochastic variability. Resolving this hidden-state estimation problem constitutes the central objective of the proposed Physics-Informed Digital Twin. This interpretation is consistent with the cyber-physical Digital Twin paradigm, where continuously updated virtual states are reconstructed from heterogeneous information sources rather than measured directly [7, 11, 24].

The economic consequences of this quality observability gap extend beyond the processing enterprise. Sugar beet payments to independent agricultural producers are typically linked to sucrose content measured at the reception point [6, 11]. If routing decisions ignore the hidden degradation dynamics of field piles, some growers may experience delayed evacuation of vulnerable biomass and therefore receive lower payments for beet mass that deteriorated during field storage before reception sampling. The post-reception laboratory lag further delays the recognition of this quality state in the transactional system, creating temporary information asymmetry between the physical condition of the delivered batch and its validated analytical representation.

Recent research has explored the application of Physics-Informed Neural Networks (PINN) for modeling non-linear biochemical processes under data-scarce conditions [9]. By embedding the differential equations of reaction kinetics such as the Arrhenius equation for temperature-dependent sucrose decay directly into the neural network loss function, PINNs demonstrate superior extrapolation capabilities compared to classical data-driven "black-box" models [9, 10]. In parallel, the logistics planning literature has developed

models for ensuring fairness in quota distribution among independent agricultural producers (Growers' Equity), addressing the institutional conflict where single-criterion optimization inevitably generates discrimination against individual suppliers [4, 11]. However, these two research streams biochemical quality forecasting and logistics optimization have largely evolved in isolation, without a unified cyber-physical framework that integrates them into a coherent decision-support system.

The companion article [12] addresses the mathematical core of supply chain synchronization, including the Drum-Buffer-Rope (DBR) adaptation for shifting bottlenecks, MILP-based routing optimization with adaptive objective switching, and the multi-phase heuristic that achieves sub-second solver response times. However, that study treats biochemical degradation coefficients as exogenous inputs supplied to the optimization layer, without modelling how these coefficients are reconstructed under delayed empirical observability of field-pile quality. The present paper fills this gap by developing the physics-informed state-estimation layer that supplies these coefficients to the logistics optimizer. This leaves unresolved the critical question of how to dynamically prioritize biochemically vulnerable field piles and how to ensure that the routing system actively prevents irreversible sucrose losses rather than merely reacting to them after they have occurred.

To the best of the author's knowledge, no existing Digital Twin simultaneously integrates physics-informed biochemical state estimation, deterministic logistics optimization, and institutional fairness within a single cyber-physical decision-support framework. Existing studies typically address Digital Twins [7, 11, 24], logistics optimization [4, 8, 15, 16, 18], or PINNs [9] independently. To bridge this gap, the present study introduces an intelligent Digital Twin architecture that integrates three complementary functional modules:

- *Physics-Informed Neural Network (PINN)* for predictive sucrose degradation estimation [9]. The model continuously forecasts the daily sugar loss coefficient for each field pile, enabling intelligent prioritization of raw material evacuation (Smart Triage) and mathematically compelling the system to evacuate biochemically vulnerable volumes before irreversible sucrose loss occurs [9].
- *Integration of PINN predictions into the MILP objective function.* By replacing the static degradation coefficient with a dynamically predicted value, the routing core is transformed from a purely distance-minimizing solver into a quality-aware margin protection tool. The adaptive weight switching mechanism [12] further amplifies this effect, ensuring that Smart Triage operates under normal conditions while factory continuity remains protected under crisis conditions.
- *Retrospective reconciliation mechanism* for resolving the post-reception analytical lag. The system temporarily utilizes PINN predictions after reception sampling and replaces them with deterministic laboratory values once validated results become available, ensuring continuous adaptation without accumulating systemic drift error [6, 7].

While the operational routing logic described in [12] determines which fields to evacuate based on capacity constraints, the present framework determines in what order to evacuate them based on biochemical vulnerability. The two frameworks are mathematically linked:

the PINN-predicted degradation coefficients serve as inputs to the MILP objective function (see Section 2.1 for the mathematical formulation), while the quota stratification algorithm [12] provides the filtered set of active nodes for the routing solver.

The scientific contributions of this study are fourfold:

(i) a Physics-Informed Digital Twin architecture that reconstructs latent sucrose degradation trajectories during the pre-reception field-storage observability gap, using temperature as an exogenous physical forcing variable and respiration kinetics as the governing physical constraint;
(ii) integration of PINN-estimated degradation coefficients into the deterministic logistics optimization framework proposed in the companion study [12], transforming the routing core from a distance-minimizing solver into a quality-aware margin protection tool;
(iii) a retrospective reconciliation mechanism that temporarily substitutes PINN estimates for missing laboratory data during the post-reception analytical lag and replaces them with validated measurements upon receipt, ensuring continuous adaptation without systemic drift;
(iv) empirical demonstration that biochemical state estimation, deterministic optimization, and institutional fairness mechanisms can be integrated into a single cyber-physical decision-support framework, validated on an 11-year industrial dataset.

Within the broader context of smart regional development, the proposed architecture contributes to the ongoing evolution of the smart cities discourse moving beyond technological novelty toward institutional resilience, decentralized trust, and tangible public value [2, 13]. By eliminating information asymmetry and creating a transparent, empirically grounded decision-making environment, the Digital Twin provides a robust foundation for E-Governance and sustainable development in regional agro-industrial clusters [1, 3].

The remainder of this paper is organized as follows. Section 2 presents the methodological framework, including the PINN mathematical formulation, the integration mechanism with the MILP core, and the retrospective reconciliation procedure. Section 3 presents empirical validation using an 11-year industrial dataset, demonstrating the magnitude of preventable sucrose losses and the effectiveness of the Smart Triage mechanism. Section 4 discusses the implications for biochemical degradation management and institutional transparency. Section 5 concludes with key findings and future research directions.

2. Materials and Methods

This section presents the methodological framework for the physics-informed quality-aware extension of the Digital Twin. The present study focuses exclusively on the Physics-Informed Neural Network (PINN) for predictive sucrose degradation estimation and its integration into the logistics optimization core for Smart Triage. Unlike the companion logistics model [12], which focuses on transportation synchronization and optimization, the present work treats the optimization layer as an external deterministic operator. The

scientific focus is therefore shifted from routing optimization toward the reconstruction of hidden biochemical system states required by that optimizer.

The empirical material comprises an 11-year dataset (2015–2025) of transactional, telemetric, and laboratory records from the agro-industrial enterprises of the Republic of Moldova (Südzucker Moldova SRL and Agro-SZM SRL).

2.1. Logistics Optimization Background

The logistics optimization core is formulated as a Mixed-Integer Linear Programming (MILP) problem that determines the daily evacuation plan – which field piles to evacuate, in what order, and with what intensity – subject to capacity constraints, availability calendars, and mass balance conservation [14, 15]. The system operates on a directed graph $G = (V, E)$ where nodes represent fields, buffers, and factories, and arcs represent transport links. A detailed formalization is provided in the companion article [12]; here we recapitulate only the elements directly relevant to the PINN integration.

2.1.1. The MILP objective function minimizes aggregate operational losses:

$$\min_{x \in X} J = \sum_{t=1}^{H} (\omega_1 \cdot Loss_{biochem}(t) + \omega_2 \cdot Penalty_{starve}(t) + \omega_3 \cdot Penalty_{slack}(t)) \quad (1)$$

Where:

x – *denotes the vector of routing and allocation decision variables*
$\mathcal{X}$ – *is the feasible set defined by capacity, availability, and mass-balance constraints.*
$Loss_{biochem}(t)$ – *is monetary loss from sucrose degradation during field storage;*
$Penalty_{starve}(t)$ – *is penalty for factory buffer depletion;*
$Penalty_{slack}(t)$ – *is feasibility safeuard for solver stability.*

2.1.2. The biochemical degradation loss is computed for each day t as:

$$Loss_{biochem}(t) = \sum_{j \in V_c} B_j(t) \cdot \delta S_j(t) \cdot \eta \cdot P_{sugar} \quad (2)$$

Where:

$B_j(t)$ – *is the remaining biomass volume in field pile j at day t (tons);*
$\delta S_j(t)$ – *is the daily sucrose loss expressed as a fraction of biomass;*
η – *is the sugar extraction efficiency coefficient;*
P_{sugar} – *is the market value of white sugar (monetary units per ton).*

$$\delta S_j(t) = \frac{\Delta S_j(t)}{100} \quad (3)$$

Where:

$\Delta S_j(t)$ – *is expressed in percentage points per day.*

The adaptive weight switching mechanism adjusts ω_1 and ω_2 based on the factory buffer inventory level: under normal conditions, ω_1 dominates, enforcing quality-aware evacuation; when factory starvation threatens, ω_2 overrides, prioritizing industrial continuity.

The complete MILP formulation–including flow conservation constraints, buffer capacities, and processing limits–is adopted unchanged from the companion article [12]. The present study focuses on replacing the static degradation coefficient $\Delta S_j(t)$ (which in the baseline model is either a fixed historical average or a simple heuristic) with a dynamically predicted coefficient $\widehat{\Delta S}_j(t)$ generated by the PINN module (Section 2.3).

This substitution enables the optimizer to prioritize evacuation of biochemically vulnerable piles – a mechanism we refer to as Smart Triage.

2.2. Data Architecture and Transaction Asynchrony Resolution

The Digital Twin operates on a three-layer data architecture that distinguishes between nominal reference data (N-layer), predictive planning data (P-layer), and empirical transactional data (E-layer). The E-layer serves as the source of physical truth, populated by discrete transactional logs from weighbridges, GPS telemetry, and quality laboratories.

A critical challenge in quality-aware logistics is that the biochemical state of beet mass becomes observable only after a batch has physically reached the reception point and has been sampled for laboratory analysis. This creates two distinct information gaps.

First, during the pre-reception field-storage interval, beet piles may remain in open storage for several days or even weeks. Throughout this period, sucrose degradation proceeds continuously, but the Digital Twin does not yet possess deterministic laboratory data for the corresponding field pile. The PINN module is therefore used to estimate the latent sucrose trajectory of active field piles before confirmed laboratory measurements become available.

Second, after a truck arrives at the reception point, a post-reception laboratory lag occurs. In the industrial workflow studied here, samples are collected at factory reception points, accumulated during the operational day, transported to a centralized laboratory, and analysed polarimetrically. As a result, validated laboratory results may enter the E-layer with a delay of up to 48 hours after reception sampling.

To overcome this two-stage quality observability gap, the system temporarily utilizes predictive estimates generated by the PINN. Upon receipt of validated laboratory data, the system enforces retrospective reconciliation:

$$S_j^{rec}(t_r) \leftarrow S_{lab,j}(t_r) \quad (4)$$

Where:

$S_j^{rec}(t_r)$ *– is the reconciled sugar content attribute of batch j at the reconciliation time* t_r, $S_{lab,j}(t_r)$ *– is the validated laboratory measurement.*

This mechanism ensures that the P-layer can support operational decision-making before deterministic measurements are available, while the E-layer ultimately anchors the Digital Twin to validated laboratory facts. Thus, the PINN estimates hidden field-pile degradation before reception, whereas retrospective reconciliation resolves the delayed arrival of post-reception laboratory results.

2.3. Physics-Informed Neural Network (PINN) for Predictive Degradation Estimation

To forecast daily sucrose loss coefficients for field piles during the pre-reception storage interval, a Physics-Informed Neural Network (PINN) is integrated as a pilot module [9]. Unlike purely data-driven models, the PINN embeds the Arrhenius-based differential equation of biochemical decay directly into the loss function, improving physically plausible extrapolation under anomalous weather conditions.

The proposed PINN is employed not as a generic forecasting model but as a hidden-state estimator constrained by biochemical kinetics. Its objective is to reconstruct latent sucrose degradation trajectories during periods in which laboratory measurements are unavailable. This differs fundamentally from conventional data-driven regressors because physical constraints are embedded directly into the optimization process through the residual term of the governing equations [10].

The global loss function $\mathcal{L}_{PINN}$ is:

$$\mathcal{L}_{PINN} = \lambda_{data} \frac{1}{N} \sum_{k=1}^{N} \left(\hat{S}(t_k) - S_{lab,k}\right)^2 + \lambda_{physics} \frac{1}{M} \sum_{i=1}^{M} \left(\frac{d\hat{S}(t_i)}{dt} + A \exp\left(-\frac{E_a}{RT_i}\right) \hat{S}(t_i)\right)^2 \quad (5)$$

Where:

$\hat{S}(t_k)$ – *denotes the predicted sucrose concentration at empirical laboratory points;*
$S_{lab,k}$ – *denotes laboratory-measured sucrose concentration;*
N – *is the number of empirical samples;*
M – *is the number of collocation points;*
A – *is the pre-exponential factor;*
E_a – *is the activation energy;*
R – *is the universal gas constant;*
T_i – *is the pile temperature at collocation point i;*
λ_{data}, $\lambda_{physics}$ – *are balancing weights for the data and physics loss components.*

The physical loss component enforces the kinetic equation, forcing the network to learn degradation trajectories that are physically plausible even when laboratory samples are sparse [9]. This is particularly important under extreme weather conditions, where pure data-driven models typically fail.

The daily sucrose loss coefficient $\widehat{\Delta S_j}(t)$ is derived as the difference between consecutive daily predictions:

$$\delta S_j(t) = \frac{\hat{S}_j(t) - \hat{S}_j(t+1)}{100} \quad (6)$$

Where:

$\hat{S}_j(t)$ *and* $\hat{S}_j(t+1)$ – *are predicted sucrose concentrations expressed in percentage points;*
$\delta S_j(t)$ – *is converted into a fraction for monetary loss calculation.*

In the current implementation, the Arrhenius equation represents the first embedded physical operator governing degradation dynamics. The proposed framework is intentionally extensible and allows additional physical constraints such as moisture transport, microbial respiration, self-heating, and genotype-dependent biochemical activity to be incorporated as independent residual terms without changing the overall Digital Twin architecture [6]. The PINN is trained to approximate the latent sucrose trajectory from sparse laboratory observations, while the temperature profile enters the physics residual as an exogenous time-dependent forcing variable.

2.4. Integration of PINN Predictions into the MILP Objective Function

The PINN-predicted coefficient $\widehat{\Delta S_j}(t)$ replaces the static parameter $\Delta S_j(t)$ in Eq. (2), yielding the updated degradation loss:

$$Loss_{biochem}^{PINN}(t) = \sum_{j \in V_C} B_j\,(t) \cdot \delta S_j^{PINN}(t) \cdot \eta \cdot P_{sugar} \quad (7)$$

This substitution enables the MILP solver to dynamically prioritize evacuation of biochemically vulnerable piles (Smart Triage). The routing decisions–which fields to evacuate and in what order–are now influenced by the predicted degradation risk, transforming the optimization core from a purely distance-minimizing solver into a quality-aware margin protection tool.

At the current integration stage, the pilot PINN model undergoes retrospective validation in shadow mode on historical data. Nevertheless, the concept is already capable of supplying the optimizer with preliminary prioritization vectors for shadow-mode Smart Triage validation [6, 7].

In reality, the biochemical decay of sugar beets in field piles is influenced by a broader set of factors, including: (i) respiratory losses, which are exponentially dependent on temperature and are the primary driver of sucrose depletion [7, 11]; (ii) microbiological decomposition, which develops on damaged tissues and is accelerated by self-heating [17, 18]; and (iii) secondary factors such as mechanical damage to root crops, the quantity of impurities (soil, tops), and deviations from optimal temperature and humidity regimes [16, 19, 20].

Respiration was selected as the initial modeling target because it represents the most sensitive and quantitatively significant degradation pathway, with a well-characterized exponential temperature dependence that is suitable for physics-informed modeling [11]. The pilot PINN module is deliberately scoped to this single mechanism to establish a baseline proof-of-concept. As empirical data collection expands – through the deployment of IoT sensors for continuous temperature and humidity profiling – the model architecture will be iteratively extended to incorporate additional degradation pathways. Each new parameter will be introduced as an independent physical constraint within the PINN loss function, following the same methodology used for the Arrhenius equation.

3. Results

This section presents the empirical validation of the quality-aware Digital Twin architecture. The evaluation is organized in three parts. First, the magnitude of preventable sucrose losses is quantified using real industrial data, establishing the empirical justification for the Smart Triage mechanism. Second, the research base and data characteristics used for PINN development are described. Third, the computational performance and predictive accuracy of the Physics-Informed Neural Network are presented, including a comparative benchmark against the industrial baseline – a simplified model that assumes constant daily sucrose losses.

3.1. Empirical Justification for Smart Triage: Magnitude of Preventable Sucrose Losses

Classical transport logistics focused exclusively on minimizing traveled distance or queue waiting times ignore a critical factor in agribusiness – the non-linear biochemical degradation (rot) of root crops stored in open field piles [16, 18]. To evaluate the magnitude of the hidden economic losses caused by a traditional, static routing approach, a retrospective analysis of the enterprise's historical dataset was conducted. The dynamics of the baseline sugar content (digestion) of received raw materials were examined across the 2015–2025 harvesting campaigns, using a sample of more than 180,000 unique logistical transactions.

The aggregated empirical data in Table 1 demonstrate a highly stable and cyclically recurring pattern of collapsing raw material quality toward the end of each macro-season.

Table 1. Dynamics of sugar content (digestion) degradation in root crops during harvesting campaigns (representative sample)

Macro-Season	Peak Sugar Week	Max Digestion (Peak, %)	Final Campaign Week	Min Digestion (End, %)	Seasonal Drop Depth (Delta, pp)
2016	2	18.28	17	16.21	-2.07
2017	3	18.73	17	16.50	-2.23
2020 (shock)	2	18.12	10	14.55	-3.57
2021	12	16.63	19	14.12	-2.51
2024 (shock)	1	17.16	12	13.59	-3.57

Source: Compiled by the author based on empirical transactional logs of Südzucker Moldova SRL.

During the nominal 2016 harvesting campaign, peak sugar content stood at 18.28 %, yet by the 17th week, due to prolonged exposure of the beets to field conditions, digestion dropped to 16.21 %. An even more destructive structural degradation was recorded during the 2020 and 2024 seasons, where weather shocks caused a catastrophic drop in digestion of 3.57 percentage points (from 18.12 % to 14.55 % and from 17.16 % to 13.59 %, respectively).

The visualisation of this degradation gap between the maximum potential of the raw materials and their actual value at the time of processing is presented in Figure 1.

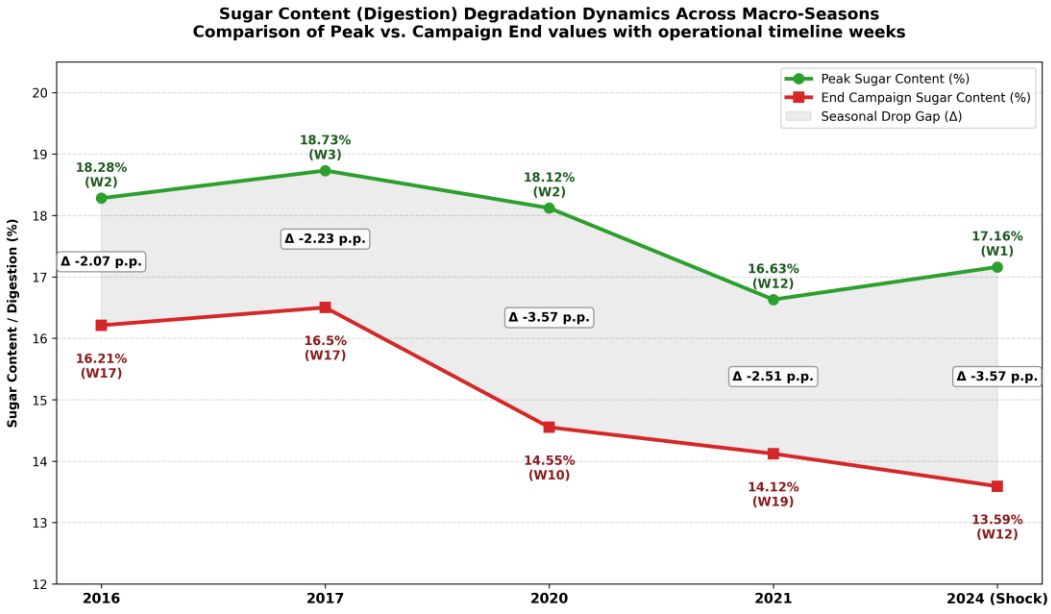


Fig.1. Dynamics of sugar content (digestion) degradation in root crops during harvesting campaigns (representative sample)

Source: Compiled by the author based on empirical transactional logs of Südzucker Moldova SRL

The retrospective analysis revealed that up to 15–18 % of the total seasonal biomass volume is traditionally processed during the final stage of the harvesting campaign (the last 4–5 weeks). Within this volume, the average algorithmically unmonitored loss of digestion reaches 11–15 % of the baseline potential. This indicates that even the deployment of basic pull logistics frameworks (Drum-Buffer-Rope) and the reduction of queues at the reception yards cannot, by themselves, halt the biochemical process of degradation inside the piles. In absolute terms, this "blind spot" within the transport system creates a massive zone of hidden economic losses, leading to the irreversible destruction of thousands of tons of crystallised sugar.

This empirical fact exposes a fundamental limitation of existing routing systems. It is precisely to mitigate these identified losses that the intelligent sorting module (Smart Triage), driven by a physics-informed neural network (PINN), was integrated into the proposed Digital Twin architecture [9].

3.2. Research Base and Data Characteristics for PINN Development

The empirical foundation for training the PINN module was derived from historical industrial operational standards used by sugar factories for loss accounting during long-term storage. Specifically, the regulatory instruction on the standardization of sugar beet and sugar losses during storage [25] provides continuous monthly loss coefficients (percentage points per day) across the entire storage season, differentiated by mechanical damage and regional climate conditions. Although this document reflects the technological norms of a different historical period, it remains the only comprehensive open-source dataset that offers daily-resolution loss profiles aggregated over decades of industrial

practice. In the context of the present study, this normative document is treated not as a binding standard, but as a historical industrial dataset – a collection of empirical observations accumulated over many harvesting campaigns under real-world storage conditions.

From a methodological perspective, the historical industrial instruction is not interpreted as a normative document but rather as a long-term empirical dataset accumulated over decades of industrial operation. Contemporary biochemical knowledge defines the governing physical relationships [6], whereas historical observations provide calibration targets for the PINN during supervised learning.

From this source, the most demanding scenario – characterised by high mechanical damage and elevated microbial risk – was selected. The monthly sucrose loss rates (percentage points per day) for this scenario are:

Table 2. The monthly sucrose loss rates

Month	**Sep**	**Oct**	**Nov**	**Dec**	**Jan**	**Feb**	**Mar**
Loss (%/day)	0.023	0.013	0.013	0.018	0.022	0.029	0.038

Source: Compiled by the author based on [25]

These values represent the average daily loss of absolute sucrose content (percentage points per day) for each month of the storage period. From these monthly figures, a continuous daily loss rate function was constructed using piecewise constant interpolation, and the corresponding "true" sucrose concentration $S_{\text{true}}(t)$ was computed by numerical integration from an initial value of 18.0 %:

$$S_{true}(t_n) = S_0 - \sum_{\tau=1}^{n} r_{hist}(\tau)\Delta\tau \quad (8)$$

Where:

$S_0 = 18.0\%$ – *is the initial sucrose concentration*

$r_{hist}(t)$ – *is the historical daily loss rate expressed in percentage points per day.*

This synthetic time series (210 days, 500 time steps) served as the ground truth for training and validation. To mimic sparse laboratory-style validation under delayed empirical observability, monthly measurements (one per month, plus the initial and final points) were extracted from the true curve. This generated a sparse set of 9 laboratory-style observations, which were used as the data term in the PINN loss function.

The temperature profile for the storage period was derived from average monthly temperatures for Moldova [26], interpolated using cubic splines to daily values. This profile captures the typical seasonal transition from autumn warmth through winter cooling to early spring warming, which drives the observed non-monotonic behaviour of sucrose degradation.

While the historical dataset provides the empirical loss coefficients, the physical interpretation of the degradation process – and the functional form of the Arrhenius-based ODE – is grounded in modern experimental studies on sugar beet storage physiology. Contemporary research [10, 21, 22, 23] has established that storage losses depend on a

combination of temperature, mechanical damage, microbial activity, and genotype, confirming that the loss rate is not a constant but a function of multiple interacting factors. In particular, studies on temperature-dependent respiration [21, 22] and the influence of harvesting systems and variety on storage losses [23, 24] provide the scientific basis for modelling the degradation kinetics as an exponentially temperature-sensitive process. These findings justify the use of the Arrhenius equation as the physical constraint in the PINN loss function, while the historical dataset supplies the empirical observations needed to calibrate the model parameters.

It is important to acknowledge that the historical dataset does not account for contemporary agricultural practices, varietal improvements, or recent climate trends. Nevertheless, it remains the only open-source reference that provides a complete seasonal loss matrix, and it serves as a sufficiently robust baseline for proof-of-concept validation of the PINN methodology. Future calibration against multi-year field data and enterprise laboratory logs is essential for adapting the model to current operational conditions.

3.3. Physics-Informed Neural Network (PINN) Performance

The PINN architecture comprises two hidden layers with 64 neurons each (≈4,000 trainable parameters) and employs the hyperbolic tangent activation function. The network takes time t (days) as input and outputs the predicted sucrose concentration $\hat{S}(t)$. The physics-informed training objective is given by (Eq. 5). The ODE residual used in the physics term is based on the Arrhenius-type degradation equation:

$$\frac{dS}{dt} = -A \exp\left(-\frac{E_a}{RT(t)}\right) S \quad (9)$$

Where:

t – *is measured in days,*
$S(t)$ – *is sucrose concentration,*
$T(t)$ – *is pile temperature in Kelvin,*
A – *is the pre-exponential factor,* $A = 10^8$ *(1/day)*
E_a – *is activation energy,* $E_a = 60000$ *J/mol,*
$R = 8.314\, J/(mol \cdot K)$.

Training was performed on a standard CPU (Intel Core i7, no GPU acceleration) for 8,000 epochs using the Adam optimizer with a learning rate of 2×10^{-3} and balanced weights for data and physics terms ($\lambda_{\text{data}} = \lambda_{\text{physics}} = 1.0$).

The convergence dynamics are presented in Fig. 2. Key performance metrics:

- Training time: 12 seconds for 8000 epochs;
- Initial total loss: 0.825;
- Final total loss: 0.136;
- Final physics loss contribution: ≈99 % of total;
- Final data loss: < 0.002.

The low final data loss indicates that the network is able to interpolate the sparse laboratory-style observations, while the remaining physics residual reflects the difficulty of satisfying the Arrhenius-based constraint over the full collocation domain. Therefore, the relative contribution of the physics term should not be interpreted as an independent

proof of physical correctness; rather, physical consistency is evaluated jointly through the residual trend, the reconstructed trajectory, and the predictive error metrics reported in Table 3.

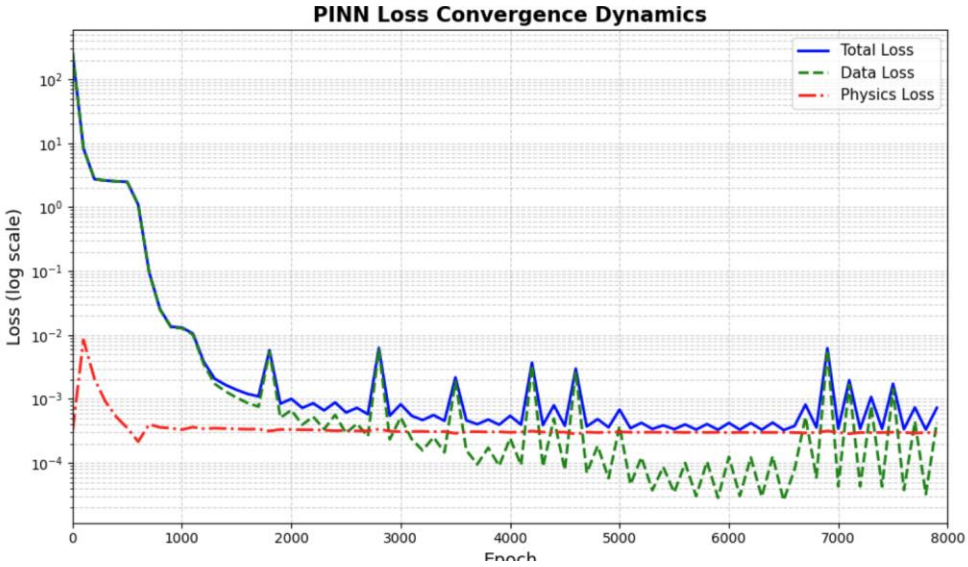


Fig. 2 – PINN loss convergence dynamics.
Source: Compiled by the author

To quantify the added value of the physics-informed approach, the PINN predictions were compared against the simplest operational heuristic – a model that assumes constant daily sucrose losses equal to the seasonal average. This averaging model is defined as:

$$S_{avg}(t) = S_0 - \bar{k}t \quad (10)$$

Where:

$\bar{k}$ *– is the average daily sucrose loss rate derived from the monthly historical values.*

The same PINN configuration was trained on the same sparse samples and temperature profile. The results are summarised in Fig. 3 and Table 3.

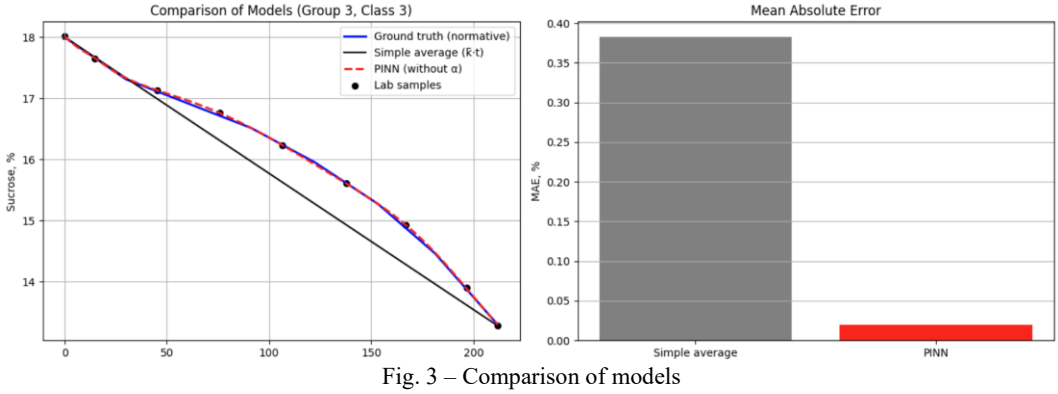


Fig. 3 – Comparison of models
Source: Compiled by the author

It should be emphasized that this benchmark is a controlled proof-of-concept experiment based on a norm-derived historical degradation curve rather than a full-scale live industrial deployment. The objective is not to claim final operational accuracy under all field conditions, but to verify whether embedding temperature-dependent kinetics improves hidden-state reconstruction compared with the constant-loss heuristic currently used as an industrial baseline.

Table 3. Predictive errors (averaged over the entire season).

Model	**RMSE, %**	**MAE, %**
Simple average ($\bar{k} \cdot t$)	0.459	0.383
PINN	0.038	0.031

Source: Compiled by the author

The PINN reduces the Mean Absolute Error by a factor of 12 compared to the simple averaging model (0.031 % vs. 0.383 %). This improvement is attributable to the PINN's explicit incorporation of temperature-dependent kinetics: the Arrhenius term naturally reproduces the observed deceleration in winter (when temperatures drop) and the subsequent acceleration in spring (when temperatures rise) – behaviour that the constant-rate heuristic entirely misses.

Nevertheless, a residual MAE of 0.03 % indicates that even the temperature-aware PINN does not perfectly match the ground truth. This discrepancy is most pronounced in the late season (March), where microbial activity becomes significant – a phenomenon not captured by the pure Arrhenius term. This gap directly supports the need for further model refinement, such as incorporating additional biochemical processes (e.g., microbial decomposition) – a direction identified in Section 2.4 as a key element for future calibration.

The predicted daily loss coefficients $\widehat{\Delta S_j}(t)$ (Eq. 6) can be supplied to the MILP core (Eq. 7) to enable Smart Triage – the algorithmic prioritisation of biochemically vulnerable field piles. In the current shadow-mode implementation, the PINN already provides a substantial accuracy gain over the industry baseline, yet its full potential will only be unlocked after extending the model and retraining on multi-year field data. The definitive selection, fine-tuning, and deep calibration of the PINN hyperparameters – including the optimal number of layers, neurons, collocation points, and the potential inclusion of additional physical constraints – constitute a primary objective for future applied research. Full-scale industrial deployment will require iterative retraining on multi-season historical logs (including anomalous weather events) and continuous validation against real polarimetric analyses. The current shadow-mode implementation has been validated on historical data and is scheduled for live deployment during the upcoming harvesting campaign.

4. Discussion

The deployment of the proposed Digital Twin architecture represents a qualitative shift in the management of agro-industrial supply chains, transitioning them from a reactive

dispatching paradigm to a continuous predictive cyber-physical synchronization loop [7, 24]. The empirical validation presented in Section 3 provides three principal insights that extend beyond the immediate scope of logistics optimization and engage with broader debates on information asymmetry, biochemical kinetics, and institutional transparency in smart regional development [1, 2, 3].

4.1. Overcoming Information Asymmetry and Integrating Biochemical Kinetics

The majority of existing routing models treat agricultural cargo as an inert mass, focusing exclusively on the minimization of transportation expenditures [16]. The retrospective analysis of the 11-year industrial dataset (Section 3.1) confirms that ignoring the thermodynamics of biologically active raw materials leads to critical financial losses on a macro-regional scale: up to 15–18 % of seasonal biomass is processed during the final 4–5 weeks, with average algorithmically unmonitored losses of digestion reaching 11–15 % of baseline potential. These figures expose a fundamental limitation of static routing heuristics and provide the empirical justification for integrating biochemical kinetics into the logistical optimisation loop.

The comparative benchmark presented in Section 3.3 demonstrates that a Physics-Informed Neural Network (PINN) – even in its basic configuration without additional microbial correction terms – reduces the Mean Absolute Error of sucrose concentration prediction by a factor of 12 compared to the industrial baseline heuristic (0.031 % vs. 0.383 %). This improvement is attributable to the explicit incorporation of temperature-dependent kinetics via the Arrhenius equation, which naturally reproduces the deceleration of losses in winter and the subsequent acceleration in spring – behaviour that the constant-rate heuristic entirely misses. The result confirms that embedding physical constraints into the neural network loss function (Eq. 5) effectively guides the learning process, while the sparse laboratory samples serve primarily as a corrective mechanism.

Nevertheless, the residual MAE of 0.03 % indicates that even the temperature-aware PINN does not perfectly match the normative ground truth. The discrepancy is most pronounced in the late season (March), where microbial activity becomes significant – a phenomenon not captured by the pure Arrhenius term. This finding directly supports the need for further model refinement, such as incorporating additional biochemical processes (e.g., microbial decomposition, self-heating, or varietal differences) as independent physical constraints within the PINN loss function. Contemporary experimental studies on sugar beet storage physiology [10, 21, 22, 23] provide the scientific basis for such extensions, confirming that storage losses depend on a combination of temperature, mechanical damage, microbial activity, and genotype.

At the current stage of development, the PINN module operates in a shadow-mode pilot configuration. The full-scale deployment of this predictive tool within a real-time management loop is objectively constrained by the necessity for deep industrial calibration of the neural network's hyperparameters against extreme meteorological deviations and multi-season field data. Nevertheless, the shadow-mode validation on historical datasets supports the technical feasibility and preliminary mathematical validity of the concept: the simulation confirms that the cyber-physical system is capable of algorithmically processing

sucrose degradation as a monetized penalty (Eq. 7), thereby establishing a scientifically grounded basis for the practical implementation of the Smart Triage mechanism. In the long term, this symbiosis of biochemistry and logistics is expected to become a key driver for ensuring food security and cyber-physical resilience within agro-industrial clusters.

Unlike purely statistical machine-learning approaches, the proposed model improves physical plausibility under sparse-data conditions because the optimization is constrained by governing physical equations rather than empirical correlations alone [10]. By reconstructing hidden states rather than relying on delayed measurements, the system shifts from reactive to predictive decision-making – aligning with the broader evolution in supply chain management from visibility toward predictive visibility.

4.2. The Role of Historical Industrial Data in PINN Development

A distinctive methodological feature of the present study is the use of historical industrial operational standards [25] as the empirical basis for training the PINN module. While this document reflects the technological norms of a different historical period and does not account for contemporary agricultural practices or climate change effects, it remains the only comprehensive open-source dataset that provides daily-resolution loss profiles aggregated over decades of industrial practice. In the context of this study, the normative document is treated not as a binding standard, but as a historical industrial dataset – a collection of empirical observations accumulated over many harvesting campaigns under real-world storage conditions.

This approach has two significant advantages. First, it provides a sufficiently dense temporal resolution (monthly loss coefficients) that is rarely available in academic experimental studies, which typically focus on controlled conditions over shorter time horizons. Second, it enables the PINN to learn from operational-scale data rather than laboratory-scale experiments, which may not fully capture the heterogeneity of field storage conditions. The physical interpretation of the degradation process – and the functional form of the Arrhenius-based ODE – is grounded in modern experimental studies on sugar beet storage physiology [10, 21, 22, 23], ensuring that the model is not merely fitting historical averages but is constrained by contemporary scientific understanding of the underlying mechanisms.

Future calibration against multi-year field data and enterprise laboratory logs is essential for adapting the model to current operational conditions and for extending the physical constraints to include additional degradation pathways (e.g., microbial decomposition, self-heating, varietal differences). As empirical data collection expands – through the deployment of IoT sensors for continuous temperature and humidity profiling – the model architecture will be iteratively extended to incorporate these factors, following the same methodology used for the Arrhenius equation.

4.3. Mutual Economic Benefit for Growers and Processors

The integration of the PINN-driven Smart Triage mechanism creates a win-win economic outcome for both independent agricultural producers (growers) and the processing

enterprise – a result that fundamentally distinguishes the proposed architecture from traditional logistics optimisation.

For growers, the direct benefit arises from the reduction of sucrose degradation during field storage. By prioritising the evacuation of biochemically vulnerable piles, the system minimises the losses that would otherwise accumulate during extended storage. Since grower payments are contractually linked to the measured sucrose content at reception [6, 11], every percentage point of sugar preserved translates into higher revenue for the supplier. The Smart Triage mechanism therefore acts as an institutional safeguard that protects the grower's share of value, preventing the erosion of income caused by logistical delays.

For the processing enterprise, the benefit is twofold. First, the delivery of higher-quality raw material directly increases the sucrose yield per ton of processed beets, improving factory throughput and reducing the specific energy and chemical inputs required for sugar extraction. Second, the reduction of in-pile degradation lowers the proportion of damaged or rotten material entering the production line, which in turn decreases the formation of invert sugars and other impurities that complicate juice purification and reduce crystallisation efficiency [11, 18]. The net effect is an increase in the quantity and quality of the final white sugar output, alongside a reduction in processing losses.

Critically, this mutual benefit is not achieved at the expense of one party. The PINN-based prioritisation does not impose additional logistics costs on the processor; rather, it reallocates existing transport capacity toward the most time-sensitive biomass. The operational cost of evacuation remains unchanged, while the value of the delivered cargo is systematically higher. Consequently, both stakeholders gain without transferring risk or expense – a rare equilibrium in supply chain management, where efficiency gains often come at the cost of supplier equity [4].

This alignment of interests is further reinforced by the retrospective reconciliation mechanism (Section 2.2), which ensures that grower payments are eventually based on deterministic laboratory data, not on predictions alone. The PINN estimates pre-reception field-pile degradation, while retrospective reconciliation resolves the post-reception analytical lag without introducing systemic bias. Once the empirical data arrives, the system corrects any predictive deviation, supporting settlement based on the validated value of the delivered biomass. This design ensures that the Digital Twin acts as a neutral arbiter of quality, rather than as a tool for one-sided optimisation.

4.4. Digital Twin as a Single Source of Truth for Quality and Transparency

The foundational contribution of the Digital Twin extends well beyond route optimisation – it establishes a comprehensive, verifiable record of physical quality at every stage of the supply chain, from the field pile to the factory buffer. The transaction-level empirical layer (E-layer) captures each critical quality-defining event: the formation of the field pile, the loading timestamp, the gross weight measurement, the polarimetric laboratory analysis, and the final tare weight. The system enforces a continuous, hardware-validated timestamp

chain; any transaction missing a required checkpoint is automatically discarded, ensuring that only fully audited biomass volumes enter the planning and settlement process [17, 20].

This rigorous data architecture provides transparency on four distinct levels:

- *Biochemical quality:* The PINN module continuously estimates the daily degradation coefficient for each pile, making the invisible process of sucrose decay algorithmically visible and monetizable. This transparency allows both growers and processors to anticipate quality deterioration before irreversible losses occur.
- *Laboratory integrity:* The retrospective reconciliation mechanism resolves the unavoidable latency between sample collection and polarimetric analysis. By temporarily substituting PINN predictions for missing laboratory data, the system eliminates the informational blind spot that traditionally allowed degradation to proceed undetected. Once the deterministic result arrives, the system irreversibly replaces the prediction, ensuring that the empirical record reflects physical truth.
- *Operational traceability:* The multi-checkpoint graph provides absolute provenance for every biomass unit: its source field, the harvesting time, the loading vehicle, the transport duration, and the reception timestamp. This granularity enables root-cause analysis of quality anomalies – whether they originate from agronomic factors, harvesting damage, or logistical delays.
- *Institutional accountability:* By making all quality-critical transactions auditable and immutable, the Digital Twin functions as a Single Source of Truth [24] that disincentivises opportunistic behaviour. Contract closures, grower payments, and transport billing are executed automatically against the validated E-layer logs, reducing administrative overhead by more than 80 % and transforming institutional relationships into a substantially more transparent, paperless cyber-physical space [2, 7, 20].

This multi-faceted transparency is the key enabler of the Smart Triage mechanism. Without the empirical certainty provided by the E-layer, the PINN predictions would lack the necessary institutional trust to influence evacuation priorities. By grounding the physical constraints in verifiable transactional data, the Digital Twin provides the empirical foundation for quality-aware logistics – and, in doing so, redefines the role of the logistics framework from a cost centre into an instrument for safeguarding product quality, stakeholder equity, and enterprise profitability.

4.5. Research Limitations

Despite the demonstrated efficiency of the proposed intelligent Digital Twin architecture, this study is subject to several objective limitations.

First, the PINN module is currently in a pilot validation phase (shadow mode) and has been tested predominantly on retrospective historical datasets. Full-scale deployment within a production environment requires further calibration of the network's hyperparameters – including the optimal number of layers, neurons, collocation points, and the potential inclusion of additional physical constraints (e.g., microbial decomposition, self-heating) – using expanded multi-season datasets. This is essential to improve operational robustness

under extreme climatic deviations and to achieve seamless integration of the neural network outputs into the transactional loop of the corporate ERP system.

Second, the predictive accuracy of the Smart Triage mechanism is critically dependent on the availability and quality of empirical data for training the PINN model. The current implementation relies on sparse laboratory samples (one per month, 9 points per season), which limits the network's ability to capture intra-pile temperature variations and the full heterogeneity of field conditions. The deployment of a pilot network of portable IoT sensors – scheduled for the current harvesting season across designated experimental plots – is expected to address this limitation by providing continuous empirical internal temperature profiles for subsequent training and fine-tuning.

Third, the multi-period Growers' Equity algorithm provides a formal mechanism for preventing systematic discrimination over a multi-year horizon, but it does not explicitly model the individual economic preferences of farmers. While historical compensation serves as a robust proxy for fairness, future research should explore mechanisms for direct stakeholder preference elicitation and their incorporation into the optimisation framework.

Fourth, the event-driven microservice architecture is currently deployed within a single corporate group comprising two legal entities – Südzucker Moldova SRL and its wholly-owned subsidiary Agro-SZM SRL – which operate under unified governance and data policies. This provides a controlled environment for validation but does not fully address the challenges of multi-enterprise, cross-border supply chains, where data sovereignty, divergent regulatory frameworks, and interoperability between heterogeneous information systems become critical. The integration with external meteorological APIs and the reliance on cloud-based data lakes also introduce dependencies on third-party service availability, which may affect operational continuity under network or service disruptions.

Fifth, the PINN was trained on historical norms [25] and may not fully capture the increasing frequency of extreme weather events observed in recent seasons. Continuous recalibration with live data is essential to maintain accuracy under climate variability.

5. Conclusion

This study presents the architectural and algorithmic components of an intelligent Digital Twin that extends beyond operational flow synchronization to address the challenges of raw material quality preservation and institutional transparency in regional sugar beet supply chains. The developed framework integrates three complementary modules – Physics-Informed Neural Networks (PINN) for predictive sucrose degradation estimation, multi-period Growers' Equity balancing with historical compensation, and an event-driven microservice architecture for end-to-end data transparency – into a unified cyber-physical system.

The research yields three principal conclusions:

First, the integration of PINN-predicted sucrose degradation coefficients directly into the logistical optimization loop enables intelligent prioritisation of raw material evacuation (Smart Triage). The comparative benchmark demonstrates that the PINN – even in its basic configuration without additional microbial correction terms – reduces the Mean Absolute Error of sucrose concentration prediction by a factor of 12 compared to the industrial baseline heuristic (0.031 % vs. 0.383 %). This improvement is attributable to the explicit incorporation of temperature-dependent kinetics via the Arrhenius equation, which naturally reproduces the observed deceleration of losses in winter and the subsequent acceleration in spring – behaviour that the constant-rate heuristic entirely misses. The residual MAE of 0.03 % indicates that further model refinement (e.g., incorporating microbial decomposition) is necessary, confirming the shadow-mode development strategy articulated in Section 2.4. Retrospective analysis of the 11-year industrial dataset reveals that traditional static routing approaches allow seasonal sugar content drops of up to 3.57 percentage points – a loss magnitude representing thousands of tons of crystalline sugar. While the PINN module is currently in pilot validation, the empirical evidence confirms both the magnitude of preventable losses and the preliminary feasibility of the physics-informed approach [9, 10, 21, 22, 23].

Second, the algorithmic separation of social balancing (Growers' Equity) from operational routing demonstrates that institutional fairness does not conflict with global supply chain efficiency. The multi-period quota stratification algorithm with historical compensation provides a formal mechanism for preventing systematic discrimination over a multi-year horizon, reducing the risk that smallholder farms are disproportionately affected by the calendar allocation of harvesting windows [4]. Critically, the PINN-driven Smart Triage mechanism creates a mutually beneficial economic outcome for both growers and the processing enterprise: growers benefit from reduced sucrose degradation and consequently higher payments linked to measured sugar content, while the processor gains higher-quality raw material, increased sucrose yield per ton, and reduced processing losses [6, 11]. By isolating social constraints within the preliminary quota allocation phase, the system minimizes aggregate logistical expenditures rather than merely redistributing them among stakeholders.

Third, the event-driven microservice architecture and RFID-enabled transactional pipeline create an empirically grounded Single Source of Truth for all supply chain participants. The Digital Twin provides comprehensive, verifiable quality records at every stage – from field pile formation through loading, transport, laboratory analysis, and reception. The triangulation of objective data streams – vehicle empty weight variance, transit cycle duration, and yield consistency – enables macro-anomaly detection and effectively eliminates opportunistic counterparty behaviour [17, 20]. The retrospective reconciliation mechanism resolves the post-reception analytical lag, ensuring that grower payments are based on deterministic laboratory data, while the PINN bridges the pre-reception quality observability gap without introducing systemic bias [6, 7]. Transaction administration overhead has been reduced by more than 80 %, and contract closures and transport billing are now executed automatically based on validated E-layer logs. This demonstrates that transitioning to Industry 4.0 delivers not only mathematically optimal routing but also

transforms institutional relationships into a substantially more transparent, paperless cyber-physical space of electronic governance [2, 20].

Within the broader context of smart regional development, the proposed architecture contributes to the ongoing evolution of the smart cities discourse by demonstrating that effective digital governance requires not only technological innovation but also institutional mechanisms that ensure transparency, fairness, and stakeholder trust [1, 3]. By eliminating information asymmetry and creating an empirically grounded decision-making environment, the Digital Twin provides a robust foundation for E-Governance and sustainable development in regional agro-industrial clusters [2, 13].

Future research directions include the full-scale industrial deployment of the PINN module with continuous recalibration using live E-layer telemetry, the expansion of the IoT sensor network for automated field pile temperature monitoring, and the integration of stakeholder preference elicitation mechanisms into the Growers' Equity framework. The coupling of the PINN-based Smart Triage mechanism with the MILP core will be further validated during upcoming harvesting campaigns.

More generally, the proposed framework demonstrates that Digital Twins for biologically active supply chains should not be regarded merely as digital replicas of logistics processes. Instead, they represent cyber-physical systems capable of continuously reconstructing hidden physical states and incorporating them into operational decision-making. Although demonstrated on sugar beet logistics, the proposed Physics-Informed Digital Twin architecture can be adapted to other perishable supply chains characterized by irreversible quality degradation during storage and transportation [1, 13, 15].

Acknowledgements. The author expresses profound gratitude to the operational, logistics, and agronomic departments of Südzucker Moldova SRL and Agro-SZM SRL for their long-term cooperation in testing and successfully implementing the developed Digital Twin architecture into a real production environment.

Conflict of Interest / Author Declaration. The author of this study, Corneliu Iatco, is currently employed as the Head of the IT Department at Südzucker Moldova SRL. This professional affiliation may be perceived as a potential conflict of interest and is therefore explicitly disclosed. The author declares that this research is an applied industrial case study based on mathematical models and algorithms originally developed by the author and integrated into daily operations at Südzucker Moldova SRL. The company provided access to infrastructure and historical data for retrospective analysis but had no role in the theoretical formulation, algorithmic design, interpretation of the published results, or preparation of the manuscript.

Data Availability & Intellectual Property. The mathematical models, logic, and predictive control algorithms (including the MPC frameworks and the PINN architecture) described in this study are the sole intellectual property and independent development of the author. The 11-year empirical dataset and the specific corporate hardware-software implementation of the system represent the proprietary commercial assets of Südzucker

Moldova SRL. Due to strict corporate confidentiality policies (NDA), raw historical datasets are not publicly available.

However, to ensure methodological reproducibility and transparency of the computational logic, the core computational algorithms (the MILP core and the advanced PINN training module) have been isolated into a standalone Python Proof-of-Concept (PoC). These scripts, along with synthetic datasets that accurately mirror the mathematical topology and statistical variance of the real-world regional supply chain, are made publicly accessible via the following GitHub repository: https://github.com/iatco13/sbdt. Furthermore, all aggregated metrics, simulation outputs, and anonymized performance indicators necessary to support and validate the conclusions of this study are comprehensively presented within the manuscript.

Funding. The theoretical foundation of this research was carried out on a personal initiative and received no external grant funding. The practical implementation of the author's algorithms into the production environment was funded by the internal budgets of Südzucker Moldova SRL as part of its planned operational digitalization strategy.

AI Use Disclosure. The original manuscript was drafted in Russian. In preparing the final English version of the manuscript, the author utilized Artificial Intelligence tools (specifically, Google Gemini Advanced and Google NotebookLM) solely as auxiliary instruments for translation, academic language editing, and formatting. The author exercised full human oversight over this process and assumes sole responsibility for the accuracy, originality, and scientific integrity of the final text. In accordance with current publication ethics, AI tools are not listed as co-authors.

References

[1] Rosca M., Ciobanu R. Smart Cities and the Rural-Urban Divide in Moldova: Balancing Innovation and Tradition. Smart Cities and Regional Development (SCRD) Journal, 2025, Vol. 9, No. 3.
[2] Vrabie C. Editorial: From Smart Systems to Smart Governance. Smart Cities and Regional Development (SCRD) Journal, 2026, Vol. 10, No. 2, C2.
[3] Bakarr K.A., Chernor J.I., Marrah A.S., et al. Smart city technologies, resource allocation and sustainable urban planning – A critical review. Smart Cities and Regional Development (SCRD) Journal, 2026, Vol. 10, No. 1, pp. 33-54.
[4] Drenovac D., Stakić Đ., Milanović M. Sugar beet transportation problem under growers' equity regulations: Metaheuristic approach. International Journal of Industrial Engineering Computations, 2025, Vol. 16, No. 4, pp. 1123-1142.
[5] Direct.Farm. Khranenie sakharnoi svekly v kagatah (Storage of sugar beet in clamps). Available at: https://direct.farm/post/khraneniye-sakharnoy-svekly-v-kagatakh-25762
[6] Asadi M. Beet-sugar handbook. Hoboken, New Jersey: John Wiley and Sons, Inc., 2007.
[7] Ivanov D., Dolgui A. A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. Production Planning & Control, 2021, Vol. 32, No. 9, pp. 775-788.
[8] Ivanovic M., Markovic D., Radovanovic M. Optimization of transport activities in the sugar beet harvesting campaign. Economics of Agriculture, 2020, Vol. 67, No. 1, pp. 215-230.
[9] Raissi M., Perdikaris P., Karniadakis G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 2019, Vol. 378, pp. 686-707.
[10] Schnepel K., Hoffmann C.M. Genotypic variability in storage losses of sugar beet. Sugar Industry, 2014, Vol. 139, pp. 123-130.

[11] Petrov S.M., Filatov S.L., Podgornova N.M. Ustoichivoe razvitie sakharnogo proizvodstva na osnove tsifrovykh dvoynikov (Sustainable development of sugar production based on digital twins). Sakhar, 2023, No. 11, pp. 16-21. (in Russian)
[12] Iațco C. Drum-Buffer-Rope and Rolling Horizon MILP for Shifting Bottleneck Synchronization in Sugar Beet Logistics. Smart Cities and Regional Development (SCRD) Journal, 2026, (under review). (companion article)
[13] Zisko F., Mema B. Digital twins and MIS integration for urban resilience: A strategic approach to smart city evolution. Smart Cities and Regional Development (SCRD) Journal, 2026, Vol. 10, No. 2.
[14] Frishev S.G., Kozupitsya S.I. Rozrobka efektyvnykh tekhnolohiy transportuvannya zerna vid kombayniv (Development of efficient technologies for transporting grain from harvesters). NUBiP of Ukraine, 2015. (in Ukrainian)
[15] Balashov A. Justification of the logistics of transporting root crops to the beet point of a sugar factory. Science in the Central Russia, 2026, No. 1, pp. 54-60.
[16] Ahumada O., Villalobos J.R. Application of planning models in the agri-food supply chain: A review. European Journal of Operational Research, 2009, Vol. 196, No. 1, pp. 1-20.
[17] Khudyakova E.V., Stepantsevich M.N., Muzalev K.S. Sovershenstvovanie upravleniya biznes-processami tsekha sakharnogo zavoda na osnove tekhnologii tsifrovykh dvoynikov (Improving the business process management of a sugar factory shop based on digital twin technology). Izvestiya Timiryazevskoy sel'skokhozyaystvennoy akademii, 2025, No. 1, pp. 15-25. (in Russian)
[18] Higgins A. Scheduling of road vehicles in sugarcane transport: A case study at an Australian sugar mill. European Journal of Operational Research, 2006, Vol. 170, No. 3, pp. 987-1000.
[19] Romano A. Seed Production and Certification in Sugar Beet. In: Sugar Beet Cultivation, Management and Processing. Springer, 2022, pp. 91-120.
[20] Seliverstova N.S., Grigoryeva O.V., Ksenofontova E.V. Tsifrovye tekhnologii v logistike i upravlenii tsepyami postavok (Digital technologies in logistics and supply chain management). HSE Publishing House, 2020. (in Russian)
[21] Kenter C., Hoffmann C.M. Changes in the processing quality of sugar beet during long-term storage under controlled conditions. International Journal of Food Science & Technology, 2009, Vol. 44, No. 5, pp. 910-917.
[22] Hoffmann C.M., et al. Importance of harvesting system and variety for storage losses of sugar beet. Sugar Industry, 2018, Vol. 143, No. 5, pp. 278-285.
[23] Hoffmann C.M. Sugar beet from field clamps – harvest quality and storage loss. Sugar Industry, 2018, Vol. 143, No. 11, pp. 634-641.
[24] Marmolejo-Saucedo J.A. Utilising Digital Twins to Bolster the Sustainability of Logistics Processes in Industry 4.0. Sustainability, 2024, Vol. 16, No. 6, 2575.
[25] Meganorm. Instruktsiya po normirovaniyu poter sveklomassy i sakhara v sveklosakharnom proizvodstve (Instruction on the standardization of sugar beet and sugar losses in sugar beet production), utv. Minpishchepromom SSSR ot 07.09.1983 N 4-IN. (accessed: 02.07.2026). (in Russian). Available at: https://meganorm.ru/mega_doc/norm/akt_forma/1/instruktsiya_po_normirovaniyu_poter_sveklomassy_i_sakhara_v.html
[26] Climate-Data.org. Climate Moldova. (accessed: 02.07.2026). Available at: https://en.climate-data.org/