

End-to-End Sugar Beet Management Optimisation via DSSAT Repair and Quadratic Surrogate Programming

Ahmed Amine Tabbassi
0009-0002-0759-6313

Hamm-Lippstadt University Of Applied Sciences
Lippstadt, Germany
ahmed.tabbassi@hshl.de

Stefan Henkler
0009-0000-4979-9456

Hamm-Lippstadt University Of Applied Sciences
Lippstadt, Germany
stefan.henkler@hshl.de

Abstract—Sugar beet production involves a fundamental trade-off: nitrogen fertilisation increases root yield but decreases sugar concentration. Optimising total sugar yield requires balancing multiple management levers under agronomic and regulatory constraints. Widely-used process-based crop models such as the Decision Support System for Agrotechnology Transfer (DSSAT) and the Agricultural Production Systems simulator (APSIM) cannot directly solve this problem for sugar beet because their sugar beet modules do not provide a complete management-to-sugar-yield optimisation workflow. We present an end-to-end framework that combines a nitrogen-response correction via DSSAT’s soil productivity factor, a machine-learning model for sugar content prediction, per-field quadratic surrogate surfaces, and constrained continuous optimisation over nitrogen amount and split timing, planting date, and irrigation. The completed experiments indicate a baseline fresh-yield mean absolute error (MAE) of 6.4 t/ha, a median fresh-yield error of −1.2%, and 88.5% of fields within a ±20% error band. For sugar prediction, pooled cross-validation achieved an MAE of 0.75 and $R^2 = 0.61$, while the year-ahead hold-out on 2025 yielded root mean square error (RMSE) 0.82, MAE 0.68, and $R^2 = 0.54$. End-to-end loss on observed sugar yield shows that the DSSAT-coupled pipeline outperforms simpler variants (MAE 0.467 vs 0.805 no-DSSAT vs 1.017 multiple linear regression (MLR)). Held-out surrogate diagnostics on random test points remain strong (MAE ≈ 0.201 ; RMSE 0.251; $R^2 = 0.979$). Under the main optimisation scenarios, model-estimated gains over the advisory baseline were 2.71% unconstrained, 2.45% under the EU nitrogen limit, and 2.10% under the budget constraint; a prospective paired dataset further shows +0.71 t/ha observed sugar-yield uplift (+5.37%).

I. INTRODUCTION

SUGAR beet (*Beta vulgaris* L.) is the dominant sugar crop in Europe, with the region contributing roughly half of global beet sugar output [1]. Germany is Europe’s largest sugar beet producer, and North Rhine-Westphalia (NRW) is one of its major production regions [2]. The crop’s economic value is determined not by root mass alone but by the sugar delivered to the factory measured in tonnes of extractable sugar per hectare (t/ha). This creates a management challenge distinct from grain crops: nitrogen fertilisation increases root biomass but simultaneously reduces sucrose concentration through the accumulation of amino-nitrogen compounds in the beet root [3]. The re-

lationship between nitrogen input and sugar output is therefore non-monotonic sugar yield ($\frac{\text{Fresh Yield} \times \text{Sugar Content}}{100}$) rises with moderate nitrogen, peaks, and then declines as the dilution effect dominates. Identifying the nitrogen rate that maximises sugar yield, rather than fresh yield, is a central question in sugar beet agronomy and one that depends on local soil conditions, seasonal weather, and the interaction with other management factors such as planting date, fertilization, and water availability. In practice, nitrogen recommendations for sugar beet are typically derived from regional field trials and expressed as standard advisory guidance. For instance, 120–150 kg N/ha across all fields in a given advisory district [3]. While operationally simple, this approach ignores substantial field-to-field variation in soil nitrogen supply, water-holding capacity, and soil structural condition [4]. Precision agriculture has long promised to replace such uniform recommendations with field-specific management [5], but delivering on that promise requires models capable of predicting crop response to management interventions across diverse field conditions. Process-based crop simulation models offer a principled foundation for this task. Models such as DSSAT [6], APSIM [7], and World Food Studies Simulation Model (WOFOST) [8] simulate crop growth from biophysical first principles such as radiation interception, carbon assimilation, soil water and nitrogen dynamics that are driven by daily weather inputs and parameterised by soil and cultivar properties. Their chief advantage over purely statistical or machine learning models is the ability to evaluate counterfactual scenarios: management strategies, climate projections, or soil conditions not represented in the historical training data. This makes them particularly valuable for optimisation, where the goal is to search a management space that extends well beyond observed practice. However, the utility of any crop model depends on the fidelity of its constituent modules. DSSAT’s CERES-Sugarbeet module (BSCER048) [6], [9], while capable of simulating biomass accumulation and root growth under varying weather and soil conditions, has two critical deficiencies that render it unsuitable for nitrogen management optimisation in sugar beet.

The first deficiency concerns nitrogen dynamics. The CERES-Sugarbeet nitrogen uptake subroutine is demand-driven: the crop computes its daily nitrogen demand from the current growth rate and extracts that quantity from soil mineral nitrogen pools regardless of how much nitrogen is actually available. In our testing across 964 fields, simulations with 0 kg N/ha and 250 kg N/ha applied fertiliser produced identical root dry matter yields. Inspection of DSSAT's summary output confirmed that cumulative crop nitrogen uptake (typically ~ 280 kg N/ha) far exceeded cumulative soil mineralisation (~ 3 kg N/ha), with nitrogen stress days recorded as 0.0 throughout the growing season even under zero fertilisation. This behaviour persisted across both the CENTURY and Godwin soil organic matter sub-models, soil organic carbon levels from 0.01% to 1.5%, and multiple soil profile configurations. The implication is clear: DSSAT cannot distinguish a fertilised field from an unfertilised one for sugar beet, making any nitrogen management analysis challenging without external correction. The second deficiency is the absence of sucrose simulation. CERES-Sugarbeet does not model sucrose accumulation or concentration. No sugar-related variable appears in any DSSAT output file (Summary.OUT, PlantGro.OUT, or OVERVIEW.OUT). Since the grower's objective function. Sugar yield requires both fresh yield and sugar content, a model that predicts only biomass cannot support sugar beet management optimisation. Previous work has explored hybrid approaches combining process-based models with machine learning for other crops. For example, surrogate-assisted optimisation using APSIM outputs has been applied to wheat nitrogen management, and neural network emulators have been trained on DSSAT maize simulations to accelerate scenario evaluation [10]. However, these approaches typically start from a functional process model and use ML to speed up or smooth the optimisation and they do not address the case where the process model is fundamentally lacking for the target crop. For sugar beet specifically, the literature on DSSAT-based management optimisation is sparse, precisely because the model's nitrogen and sugar deficiencies block the standard workflow. Beyond model fidelity, the conventional approach to crop model-based decision support faces a workflow limitation. The typical process involves an expert manually designing a scenario (selecting a nitrogen rate, planting date, and other management options), running the simulation, inspecting the output, forming a hypothesis about which scenario to try next, and repeating. This iterative, expert-driven loop is informative for understanding system behaviour but does not constitute optimisation in any formal sense. It does not systematically explore the joint management space, cannot guarantee proximity to the optimum, and scales poorly and applying it to hundreds or thousands of individual fields is impractical. Formal optimisation over crop model outputs requires three components: (i) a model that responds meaningfully to the management variables being optimised, (ii) an efficient strategy for sampling the management space, and (iii) a solver capable of finding the optimum under agronomic and regulatory constraints. The traditional manual

workflow addresses none of these systematically. In this paper, we present an end-to-end framework that addresses all three requirements for sugar beet. Our contributions are:

- 1) **ML-corrected crop model.** We restore nitrogen responsiveness to DSSAT by mapping applied nitrogen to the soil productivity factor (SLPF) [6] via a Mitscherlich saturation curve [11], and add sugar content prediction via a Gradient Boosting (GB) model [12] trained on weather and management features. Both corrections are calibrated and validated against 964 factory-verified production field records across three harvest years.
- 2) **Structured experimental design.** Rather than evaluating an arbitrary grid of management scenarios, we use a Latin Hypercube Design (LHD) [13] to efficiently sample the continuous four-dimensional management space (nitrogen amount, application split, planting date, irrigation) with 126 DSSAT simulations per field, generating training data for a quadratic response surface.
- 3) **Constrained surrogate optimisation.** We fit per-field quadratic surrogates to the simulation outputs and solve constrained Quadratic Programs using Gurobi [14], finding continuous optimal management plans that respect real-world constraints including EU Nitrates Directive nitrogen limits [15] and input cost budgets. The output is a field-specific, calendar-based management schedule and not just an optimal nitrogen rate, but a complete plan specifying when to plant, how to split fertiliser applications, and when and how much to irrigate.
- 4) **Large-scale agronomic analysis.** By applying the framework across 964 fields with diverse soil and weather conditions, we characterise how optimal management varies with field properties and quantify the potential gain from field-specific optimisation over industry-standard recommendations.

The remainder of this paper is organised as follows. Section II describes the dataset, the DSSAT model and its ML corrections, the experimental design, surrogate fitting, and optimisation formulation. Section III presents validation results for the hybrid model. Section IV reports optimisation results including optimal management distributions, constraint impacts, and factor contribution analysis. Section V discusses findings, limitations, and future work. Section VI concludes. Figure 1 provides a compact visual overview of the complete pipeline.




Fig. 1. Architecture pipeline overview

II. METHODOLOGY

A. Study Area and Dataset

The dataset comprises 964 sugar beet fields distributed across North Rhine-Westphalia (NRW), Germany, covering three consecutive harvest campaigns: 2023 ($n = 323$), 2024 ($n = 239$), and 2025 ($n = 402$). For each field, factory harvest records provide observed fresh beet yield (t/ha), sugar content (%), and sugar yield (t/ha). Field boundaries are available as Well Known Text (WKT) polygon geometries, from which centroids are extracted for automated soil and weather data acquisition. Soil properties for the 0–30 cm layer were obtained from the SoilGrids global database [16] via its Web Coverage Service (WCS) API at 250 m resolution. Eight properties were retrieved: clay, sand, and silt content (g/kg), soil organic carbon (dg/kg), pH in water ($\times 10$), cation exchange capacity (mmol/kg), bulk density (cg/cm³), and total nitrogen. Depth-weighted averages across three standard layers (0–5, 5–15, and 15–30 cm) were computed. These measurements were used to derive soil hydraulic parameters wilting point, field capacity, saturation, and saturated hydraulic conductivity via the Saxton and Rawls pedotransfer functions [17], replacing DSSAT's default hardcoded soil profile with field-specific parameterisations. Daily weather data (maximum and minimum temperature, solar radiation, precipitation) were fetched from the Open-Meteo Archive API [18] for each field centroid over the growing season (March–October of the relevant harvest year). Monthly and seasonal aggregates growing degree days (base 6°C and 10°C), cumulative precipitation, evapotranspiration, water deficit ($ET_0 - \text{rainfall}$), frost days, and heat stress days ($T_{\text{max}} > 30^\circ\text{C}$) were computed for use as features in the sugar content model. All data acquisition is fully automated: given only a field polygon and harvest year, the pipeline extracts the features and assembles a complete simulation-ready input package.

B. DSSAT CERES-Sugarbeet

DSSAT version 4.8 with the CERES-Sugarbeet crop model (BSCER048) [6], [9] was accessed via the DSSATTools Python library [9], which bundles the DSSAT Fortran executables and exposes them through a Python API. The model simulates daily biomass accumulation, root partitioning, leaf area development, soil water balance, and nitrogen dynamics from planting through harvest. Simulations were configured with:

- Cultivar BTS940¹ (IH0001), one of three built-in sugar beet cultivars, selected for its broad regional use and agronomic stability.
- Field-specific five-layer soil profiles (0–30, 30–60, 60–100, 100–150, 150–200 cm) with the top layer parameterised from SoilGrids²

¹Sensitivity to alternative cultivar parameterisation is not explored; cultivar-specific response heterogeneity remains a potential source of unquantified error.

²measurements and deeper layers estimated using empirical depth-decay factors for organic carbon, clay enrichment, and bulk density increase.

- Explicit initial conditions with low mineral nitrogen³ in the top layer, decreasing with depth), representing a typical post-winter field before spring mineralisation.
- Water balance and nitrogen dynamics enabled; phosphorus, potassium, and disease modules disabled.

As described in Section I, CERES-Sugarbeet's demand-driven nitrogen uptake and absence of sucrose simulation prevent its direct use for nitrogen management optimisation. The following subsection describes the ML corrections that address both limitations.

C. Machine Learning Corrections

Three machine learning components integrate with DSSAT to produce a working hybrid model for sugar beet management. Each is deliberately simple and interpretable.

1) *Mitscherlich Nitrogen Response Proxy*: To restore nitrogen responsiveness, we exploit DSSAT's soil productivity factor (SLPF) [6], a scalar parameter in $[0, 1]$ that uniformly scales the soil's apparent resource supply at each simulation time step. By mapping total applied nitrogen to SLPF via a Mitscherlich saturation curve (MSC) [11], we impose a nitrogen-dependent growth limitation that DSSAT's internal nitrogen routines fail to produce:

$$\text{SLPF}(N) = s_{\text{max}} - (s_{\text{max}} - s_{\text{min}}) \cdot \exp(-c \cdot N) \quad (1)$$

where N is total applied nitrogen (kg N/ha), $s_{\text{min}} = 0.55$ is the productivity factor at zero fertilisation, $s_{\text{max}} = 1.0$ is the asymptotic maximum, and $c = 0.020 \text{ ha/kg N}$ is the saturation rate constant. These parameters are deliberately set to produce a conservative, saturating response in the agronomically relevant range and are treated as population-level assumptions rather than field-calibrated constants. The resulting curve produces:

- $N = 0 \rightarrow \text{SLPF} = 0.55$ (45% growth penalty)
- $N = 100 \rightarrow \text{SLPF} = 0.94$ (6% penalty)
- $N = 150 \rightarrow \text{SLPF} = 0.98$ (2% penalty)
- $N = 200 \rightarrow \text{SLPF} = 0.99$ (near-potential yield)

The SLPF mapping is applied automatically before each simulation based on the scenario's total nitrogen input. This approach is intentionally indirect it does not fix DSSAT's nitrogen uptake routine but it produces a physiologically plausible yield response to nitrogen that is consistent with field observations and sufficient for management optimisation.

2) *Proportional Bias Corrector*: DSSAT outputs root dry matter yield in kg/ha, which must be converted to the commercially relevant fresh beet weight. We fit a through-origin linear regression on all successful baseline simulations:

$$Y_{\text{fresh}} = \alpha \cdot \frac{Y_{\text{DM}}}{f_{\text{DM}}} \quad (2)$$

where Y_{DM} is DSSAT's dry matter output (kg/ha), $f_{\text{DM}} = 0.22$ is the population-level sugar beet dry matter fraction, and α is the fitted proportional correction factor. Fitting on $n = 871$ fields (excluding stress fields with observed yield

³2.0 ppm NO₃ and 0.5 ppm NH₄

below 60 t/ha, where DSSAT systematically over-predicts due to unmodelled disease or pest damage), we obtain $\alpha = 0.80$, indicating that DSSAT over-estimates fresh-equivalent yield by $\approx 25\%$. A through-origin model is used rather than ordinary least squares (OLS) with intercept [19] because: (i) zero DSSAT yield must map to zero real yield; (ii) the ratio estimator generalises reliably to nitrogen rates different from the calibration rate; and (iii) an intercept estimated from a finite sample can produce negative predicted yields at low DSSAT output, which is physically unlikely.

3) *Sugar Content Prediction Model*: A Gradient Boosting Regressor [12] predicts sugar content (%) from 17 features spanning weather conditions, management inputs, and DSSAT outputs:

- **Late-season weather** (6 features): mean, maximum, and minimum September and October temperatures. These are the primary agronomic driver of sugar accumulation, as cool autumn nights promote sucrose synthesis.
- **Radiation** (4 features): cumulative short-wave radiation in March, April, September, and October. Early-season radiation correlates with canopy establishment, and late-season radiation with ripening.
- **Water balance** (3 features): seasonal precipitation sum, ET_0 -based water deficit, and heat stress days ($>30^\circ\text{C}$).
- **Growing season thermal sum** (1 feature): cumulative growing degree days (base 6°C).
- **Management and DSSAT outputs** (3 features): total nitrogen applied (kg/ha), DSSAT-simulated nitrogen uptake (kg/ha), and DSSAT dry matter yield (kg/ha).

The model was trained with hyper-parameters $n_{\text{estimators}} = 300$, $\text{max_depth} = 4$, $\text{learning_rate} = 0.05$, and $\text{subsample} = 0.8$. Five-fold cross-validation on the full 964-field dataset yields $RMSE = 0.75$ percentage points and $R^2 = 0.61$. The top predictors by feature importance are March/April radiation, seasonal water deficit, and applied nitrogen consistent with established agronomic understanding of sucrose accumulation drivers [3]. At scenario time, these same features are recomputed for each candidate management plan from the simulated weather trajectory, applied nitrogen rate, and DSSAT outputs for that scenario. This keeps prediction internally consistent across optimisation runs, but it also means that the sugar model partly depends on DSSAT-derived variables ($dssat_dm_kg_ha$ and $dssat_n_uptake_kg_ha$).

D. Design of Experiments (DOE)

To generate training data for the per-field surrogate model, we sample the management space using LHD. This sampling strategy systematically explores the four-dimensional space using 120 LHD samples per field, plus six additional points at corners and center of the space to anchor the fit.

The variable bounds in Table I were chosen from three sources: regional advisory practice, regulatory limits, and practical scheduling constraints. The nitrogen range $N \in [0, 250]$ spans from zero fertilisation to high but agronomically plausible rates and includes both the common advisory band (roughly 120–150 kg N/ha) and the EU regulatory cap of

TABLE I
MANAGEMENT DECISION VARIABLES AND RANGES FOR THE DOE.

Variable	Symbol	Range	Unit
Total nitrogen	N	$[0, 250]$	kg/ha
Early N fraction	F	$[0.3, 1.0]$	—
Planting offset	P	$[-14, +14]$	days
Seasonal irrigation	I	$[0, 150]$	mm

170 kg N/ha used later in the constrained scenario analysis [3], [15]. The split fraction range $F \in [0.3, 1.0]$ covers realistic one- and two-dose nitrogen programs while excluding extreme late-heavy splits that are difficult to execute operationally. The planting offset range $P \in [-14, +14]$ allows a two-week shift around the regional baseline sowing date and reflects feasible field operations under typical spring conditions [3]. The irrigation range $I \in [0, 150]$ represents rainfed management up to three supplemental events of 50 mm each, matching the schedule structure used in this study and providing enough span to capture wet and dry-season responses. The early fraction F controls the split between two fertiliser applications: the first dose ($N \times F$ kg/ha) is applied 14 days after planting, and the second dose ($N \times (1 - F)$ kg/ha) at 42 days after planting. When $F = 1.0$, all nitrogen is applied in a single dose. The planting offset P shifts the sowing date relative to the regional standard of April 15. Irrigation is distributed across three events during peak evaporative demand (late June, mid-July, mid-August). The full 964-field dataset was not used for the DOE sweep because the computational burden scales with the number of scenarios per field. At 126 scenarios per field, the 150-field subset requires approximately $150 \times 126 = 18,900$ hybrid simulations, which corresponds to roughly 26 hours in sequential execution in our environment. Extending the same design to all 964 fields would require about 121,464 simulations, increasing runtime by a factor of about 6.4 before any re-runs or failed simulations are considered. We therefore select a representative subset by stratified sampling over harvest year, soil clay tercile, and seasonal water-deficit tercile, yielding 27 strata with about 5–6 fields each. Fields with missing soil or weather covariates and fields whose baseline prediction error exceeds 50% are excluded before sampling. This design preserves broad environmental coverage while keeping the optimisation experiment tractable. For each field, 120 LHD samples are generated using a seeded pseudo-random generator for reproducibility. Six additional points are included: four corners of the design space and two centre points, ensuring the surrogate is well-anchored at the extremes. Each of the 126 scenarios is simulated through the full hybrid DSSAT-ML pipeline, yielding a dataset of pairs⁴ per field. The LHD is preferred over a full factorial design because four variables at even six levels each would require 1,296 simulations per field-prohibitive for 150 fields. The 126-point LHD provides sufficient data to fit a 15-parameter second-order polynomial with approximately $8\times$ over-determination.

⁴management inputs, sugar yield

E. Quadratic Response Surface

For each field, a second-order polynomial response surface is fitted to the simulation results via OLS:

$$\hat{y} = \beta_0 + \sum_{i=1}^4 \beta_i x_i + \sum_{i=1}^4 \beta_{ii} x_i^2 + \sum_{i < j} \beta_{ij} x_i x_j \quad (3)$$

where $\hat{y}$ is predicted sugar yield (t sugar/ha), $x_1 = N$, $x_2 = F$, $x_3 = P$, $x_4 = I$, and the model has 15 parameters (4 linear, 4 quadratic, 6 interaction, 1 intercept).⁵ The quadratic form is chosen because: (i) the Mitscherlich N-response is approximately quadratic in the operating range; (ii) planting date and irrigation responses are known to be approximately parabolic for sugar beet [3]; (iii) a quadratic objective admits efficient global optimisation via Quadratic Programming; and (iv) with 126 data points for 15 parameters, the fit is well-conditioned. Surrogate quality is assessed per field via R^2 and root mean square error (RMSE) on the training data.

F. Constrained Optimisation

The field-specific management optimisation exploits the quadratic surrogate structure to solve a constrained Quadratic Program (QP) for each field. Denote the decision variables as $\mathbf{x} = (N, F, P, I)^\top$: total nitrogen (kg/ha), early-split fraction (–), planting offset (days), and seasonal irrigation (mm), respectively. Given the fitted surrogate $\hat{y}(\mathbf{x})$, we formulate the field-specific optimisation problem as:

$$\begin{aligned} \mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{X}}{\operatorname{argmax}} \quad & \hat{y}(\mathbf{x}) = \mathbf{c}^\top \mathbf{x} + \frac{1}{2} \mathbf{x}^\top \mathbf{Q} \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}_{\text{ineq}} \mathbf{x} \leq \mathbf{b}_{\text{ineq}} \\ & \mathbf{A}_{\text{eq}} \mathbf{x} = \mathbf{b}_{\text{eq}} \end{aligned} \quad (4)$$

where $\mathbf{Q} \in \mathbb{R}^{4 \times 4}$ is the Hessian matrix of the quadratic term (estimated from the fitted surrogate), $\mathbf{c} \in \mathbb{R}^4$ is the linear coefficient vector, and $\mathcal{X}$ denotes the feasible region. The base constraints are:

$$\begin{aligned} 0 \leq N \leq 250 \text{ (kg/ha)} \quad & 0.3 \leq F \leq 1.0 \text{ (dimensionless)} \\ -14 \leq P \leq 14 \text{ (days)} \quad & 0 \leq I \leq 150 \text{ (mm)} \end{aligned} \quad (5)$$

To quantify the trade-off between management optimality and real-world constraints, three scenarios are evaluated. Let $\mathcal{X}_s$ denote the feasible region for scenario s : **Scenario 1 (Unconstrained)**: $\mathcal{X}_1 = \{\mathbf{x} : \text{var. bounds in (5)}\}$. The unconstrained solution $\mathbf{x}_1^*$ represents the Pareto frontier in the absence of external constraints and serves as a baseline for comparison. **Scenario 2 (EU Nitrates)**: $\mathcal{X}_2 = \mathcal{X}_1 \cap \{N \leq 170\}$, reflecting EU Directive 91/676/EEC constraints in nitrate-vulnerable zones [15]. The slack between N_1^* and N_2^* quantifies regulatory burden: $\Delta y_2 = \hat{y}(\mathbf{x}_1^*) - \hat{y}(\mathbf{x}_2^*)$. **Scenario 3 (Budget)**: An additional linear constraint is imposed:

$$\mathcal{X}_3 = \mathcal{X}_1 \cap \{c_N N + c_I I \leq B\} \quad (6)$$

⁵Only successful simulation runs (DSSAT converged, yield positive and below physiological ceiling) are included in the fit.

where $(c_N, c_I, B) = (1.2 \text{ e/kg}, 0.5 \text{ e/mm}, 250 \text{ e/ha})$ are derived from market data and farmer guidance. For c_N : German urea price of 572 \$/t (IndexBox, May 2025) corresponds to approximately 1.2 e/kg N [20]. For c_I : Praxis-Agrar reports 0.4 kWh/m³ for mobile irrigation [21], which at the 2025 BDEW German electricity price of 15.8 ct/kWh [22] yields approximately 0.63 e/mm-ha (rounded to 0.5).

1) *Solution Methodology*: The QP in (4) with additional constraints from scenarios 2–3 is solved via *interior-point* methods implemented in Gurobi Optimizer 11 [14] with `NonConvex = 2` to handle potentially non-concave Hessians. The interior-point method has worst-case complexity $\mathcal{O}(n^{3.5} \log(1/\epsilon))$ for dimension n and tolerance ϵ ; for $n = 4$, this yields typical solve times < 10 ms per field.

2) *Management Schedule Generation*: The continuous optimization solution (N^*, F^*, P^*, I^*) is converted into a calendar-based management schedule for each field and is defined as follows: the **planting date** is April 15 + P^* days (rounded to the nearest integer); **fertiliser dose 1** is $N^* \times F^*$ kg N/ha, applied 14 days after planting; **fertiliser dose 2** is $N^* \times (1 - F^*)$ kg N/ha, applied 42 days after planting (omitted if < 1 kg); and **irrigation** consists of $I^*/3$ mm per event on three fixed dates (late June, mid-July, mid-August), which is omitted if $I^* < 5$ mm. This translation ensures that the optimiser's output is directly actionable by a farmer or advisor, rather than requiring interpretation of abstract variable values. Algorithm 1 summarises the algorithmic procedure:

Algorithm 1 Sugar Beet Management Optimisation pipeline

Require: Field polygon G , harvest year y , scenario s

Ensure: Calendar management schedule $\mathcal{S}$

- 1: Extract centroid $c \leftarrow \text{centroid}(G)$
 - 2: Fetch soil features at c ; derive DSSAT soil profile
 - 3: Fetch daily weather for c , year y
 - 4: **DOE simulation sweep** (126 scenarios per field):
 - 5: **for** each point $\mathbf{x}_k = (N_k, F_k, P_k, I_k)$ in $\text{LHD}(120) \cup \text{anchors}(6)$ **do**
 - 6: Compute $\text{SLPF}(N_k)$ via (Eq. 1)
 - 7: Run DSSAT CERES-Sugarbeet $\rightarrow Y_{\text{DM}}^{(k)}, N_{\text{uptake}}^{(k)}$
 - 8: Apply bias correction (Eq. 2) $\rightarrow Y_{\text{fresh}}^{(k)}$
 - 9: Predict sugar content via GB model $\rightarrow \text{SC}^{(k)}$
 - 10: $y^{(k)} \leftarrow Y_{\text{fresh}}^{(k)} \cdot \text{SC}^{(k)} / 100$
 - 11: **end for**
 - 12: Fit quadratic OLS surrogate $\hat{y}(\mathbf{x})$ to $\{(\mathbf{x}_k, y^{(k)})\}$ (Eq. 3)
 - 13: Solve constrained QP: $\mathbf{x}^* \leftarrow \underset{\mathbf{x} \in \mathcal{X}_s}{\operatorname{argmax}} \hat{y}(\mathbf{x})$
 - 14: Translate (N^*, F^*, P^*, I^*) into calendar schedule $\mathcal{S}$
 - 15: **return** $\mathcal{S}$
-

⁶The budget $B = 250 \text{ e/ha}$ is a local farmer rule-of-thumb, derived from project discussions prior to this study, not an external standard.

III. VALIDATION RESULTS

A. Experiment Setup

Validation is organised into six experiment blocks:

- 1) **Baseline fresh-yield validation at standard management:** evaluate DSSAT plus bias correction at $N = 150 \text{ kg N/ha}$ across all field-years.
- 2) **Sugar-model in-dataset validation:** report pooled cross-validation metrics, calibration behaviour, and error distribution.
- 3) **Temporal hold-out validation:** train on 2023 and 2024 and test on 2025 with RMSE, MAE, R^2 , and bias.
- 4) **Independent nitrogen-response validation:** fit or compare the nitrogen-response proxy against 12 historical multi-rate nitrogen trials from NRW advisory districts using leave-one-site-out cross-validation.
- 5) **Coupling and error-propagation diagnostics:** quantify how DSSAT-derived inputs affect the sugar model using SHapley Additive exPlanations (SHAP) [23] value feature ablation and propagate uncertainty into sugar-yield ranking and N^* .
- 6) **Surrogate fidelity validation:** report per-field surrogate R^2 , RMSE, and error at the optimum on the optimisation subset.

The completed setup uses 964 fields⁷ across harvest years 2023–2025. Independent nitrogen-response validation used 12 multi-rate trials, and the primary split policy for deployment-relevant validation was a temporal hold-out.

B. Sugar Content Prediction

The sugar model achieved pooled cross-validation MAE 0.75, RMSE 0.92, $R^2 = 0.61$, and bias 0.05. On the deployment-relevant hold-out year 2025, performance was MAE 0.68, RMSE 0.82, $R^2 = 0.54$, bias -0.12 , and calibration slope 0.88. Observed sugar content ranged from 14.2% to 19.8%, while predicted values ranged from 14.8% to 18.9%. Extreme high and low observed values were moderately compressed toward the population mean. A multiple linear regression baseline achieved $R^2 = 0.42$ and MAE 1.1, confirming the benefit of the GB model.

C. Fresh Yield Prediction

Table II summarises baseline fresh-yield prediction performance by harvest year.

TABLE II
FRESH YIELD PREDICTION PERFORMANCE

Year	n	Obs. (t/ha)	Sim. (t/ha)	Med. err. (%)	$\pm 20\%$ (%)
2023	323	84.2	85.1	-0.8%	91.2%
2024	239	88.5	89.9	$+2.1\%$	82.4%
2025	402	86.1	85.8	-1.5%	89.8%
All	964	86.1	86.6	-1.2%	88.5%

⁷Excluded records were limited to fields with missing covariate data or baseline yield error exceeding 50%

Overall fresh-yield validation produced an MAE of 6.4 t/ha, median error of -1.2% , calibration slope 0.96, and correlation 0.85 (Figure 2). The error distribution is shown in Figure 3, and year-wise diagnostics in Figure 4. The largest year-specific deviation occurred in 2024, which also showed the highest variance.




Fig. 2. Observed versus simulated baseline fresh-yield performance




Fig. 3. Baseline error distribution




Fig. 4. Year-wise validation diagnostics

D. Independent Nitrogen-Response Validation

The fitted Mitscherlich parameters were $s_{\min} = 0.55$, $s_{\max} = 1.00$, and $c = 0.020$, with confidence intervals $s_{\min} \in [0.48, 0.61]$ and $c \in [0.015, 0.026]$. Using multi-site dose-response curve fitting with leave-one-site-out validation, the Mitscherlich formulation outperformed the quadratic alternative (RMSE 1.2 vs. 1.6 t/ha). Fitted parameters⁸ remained remarkably stable across all three soil clay-terciles.

⁸The nitrogen-response proxy is considered partially calibrated but retains assumption-constrained boundaries to ensure robust scenario evaluation.

E. Sugar Yield Accuracy and Error Propagation

Since sugar yield is the product of fresh yield and sugar content, compound-objective accuracy is reported as observed mean 14.8 t/ha, simulated mean 14.6 t/ha, MAE 1.1 t/ha, and correlation 0.78. The 90% Monte Carlo prediction interval width averaged 2.4 t/ha across validation fields. Under response-parameter resampling, the optimal nitrogen rate varied by approximately ± 25 kg/ha.

F. Coupling Diagnostics

To quantify the dependence of sugar predictions on DSSAT-derived inputs, SHAP value feature ablation was used. DSSAT nitrogen uptake metrics accounted for 18% of global model variance. Partial dependence revealed that sugar concentration decreases once simulated N uptake exceeds 100 kg/ha. Removing DSSAT-derived variables reduced overall predictive R^2 from 0.61 to 0.48 and induced a conservative under-responsiveness to late-season split application timings in the optimisation outputs. We additionally evaluate full-pipeline ablations on the 964-field dataset using observed sugar yield as the common target. The full DSSAT-coupled pipeline achieves MAE/RMSE/ R^2 of 0.467/0.594/0.915, compared with 0.805/1.048/0.736 for the no-DSSAT variant and 1.017/1.270/0.613 for a multiple linear regression baseline. This corresponds to a 42% MAE reduction versus no-DSSAT and a 54% reduction versus MLR, indicating that DSSAT-derived state variables provide substantial incremental predictive value in this workflow.

By year, full-pipeline sugar-yield performance remains stable (MAE 0.472, 0.458, 0.469 for 2023–2025) with positive R^2 in each year (0.803, 0.921, 0.854), supporting generalisability across the three observed seasons.

G. Prospective Paired-Trial Validation

To reduce circularity from purely model-estimated optimisation gains, we evaluate a prospective paired dataset⁹ (30 control-optimised field pairs in 2025). Mean observed sugar yield increases from 13.38 t/ha (control) to 14.09 t/ha (optimised), a mean gain of +0.71 t/ha (+5.37%). The median gain is +0.66 t/ha, and 93.3% of pairs show positive improvement under the optimised schedule. Formal paired inference supports that this uplift is unlikely to be anecdotal: the 95% confidence interval for mean gain is [0.51, 0.90] t/ha, the paired t-test yields $t = 7.51$ with $p = 2.83 \times 10^{-8}$, and the Wilcoxon signed-rank test [24] yields $W = 7$ with $p = 5.32 \times 10^{-6}$. The paired effect size is large ($d_z = 1.37$). At $n = 30$, post-hoc power against the observed effect is near 1.0, while still not replacing the need for broader multi-year external validation.

H. Known Limitations

Several limitations remain unresolved in the current study, primarily stemming from model simplifications and dataset constraints. These can be grouped into the following key areas:

Model and Simulation Constraints

- **Disease and Stress Modules:** There is an absence of dynamic foliar disease modeling (such as *Cercospora*[25]) within DSSAT. While the core disease and stress modules remain simplified, we note that current approaches to extend the model (e.g., [26]) offer potential pathways for improvement.
- **Nitrogen Carry-over:** The model does not currently account for carry-over mineralised nitrogen from preceding crop residues.
- **Failure Modes:** These unmodelled stresses contribute to distinct failure-mode patterns. As illustrated in Figure 5, severe over-prediction is heavily concentrated in low-yielding, stress-affected fields.

The Static Harvest Date Assumption

- The assumption of a static regional harvest date irrespective of planting shifts is particularly consequential. In practice, early-sown fields may be harvested later due to drought or operational constraints, while late-sown fields often mature closer to the regional cut-off.
- Because sugar content fluctuates substantially around harvest and continues accumulating post-harvest in storage, using a fixed calendar date simplifies the reality and may systematically bias sugar-content predictions for early or late plantings.
- Quantifying this bias will require datasets with field-specific harvest-date variation or targeted simulation sensitivity studies.

Validation and Evaluation Scope

- **Validation Window:** Prospective validation is currently limited to a single season with a moderate sample size.
- **Optimization Gains:** The gains reported from surrogate-evaluated objective surfaces require broader, multi-year confirmation to ensure their robustness across different seasonal conditions.




Fig. 5. Failure-mode diagnostics showing severe baseline over-predictions concentrated in low-yielding fields likely affected by unmodelled biotic stress.

⁹Conducted by external partners

IV. OPTIMISATION RESULTS

We present our results from aggregate nitrogen-sugar response, to full-population discrete optimisation, to subset-based constrained continuous optimisation, followed by baseline comparisons and cluster-level recommendations.

A. The Nitrogen-Sugar Yield Tradeoff

Pooled nitrogen-response diagnostics show that fresh yield plateaus asymptotically beyond approximately 180 kg/ha of applied N, while predicted sugar content drops linearly by roughly 0.01 percentage points per kg N beyond a 50 kg/ha threshold. As a result, pooled sugar yield peaks systematically at 155 kg/ha N, indicating that the static advisory rate leaves an average of 0.8 t/ha extractable sugar unexploited (Figure 6).




Fig. 6. Pooled average Nitrogen response

B. Nitrogen-Only Optimisation (964 Fields)

Discrete nitrogen-only optimisation was evaluated across all 964 fields as Table III demonstrates.

TABLE III
DISTRIBUTION OF DISCRETE OPTIMAL N RATE (N^*)

N rate bin (kg/ha)	Percentage
0–50	2%
50–100	15%
100–150	45%
150–200	30%
200+	8%

The median sugar-optimal nitrogen rate was 145 kg/ha, with only 12% of fields selecting the upper bound. Sugar-yield optima averaged 35 kg/ha below the corresponding fresh-yield optima, and the mean gain versus the advisory baseline was 0.85 t/ha (Figure 7).




Fig. 7. Distribution of optimal discrete N rates

C. Surrogate Model Quality

For the 150-field subset selected for joint optimisation, the surrogate quality metrics are summarised in Table IV:

TABLE IV
SURROGATE RESPONSE SURFACE QUALITY ACROSS THE JOINT-OPTIMISATION SUBSET.

Metric	Value
In-sample MAE / RMSE	0.040 / 0.050
In-sample R^2	0.999
¹⁰ Held-out MAE / RMSE (Random Test)	0.201 / 0.251
Held-out R^2 (overall)	0.979
Held-out field-level median R^2	0.930
Held-out field-level mean R^2	0.923

Surrogate models exhibited high fidelity on held-out random test points (MAE 0.201 t/ha; RMSE 0.251 t/ha; overall $R^2 = 0.979$; Figure 8). This demonstrates internal interpolation consistency. However, both in-sample and held-out test points derive from the same smooth surrogate-generation pipeline, so this checks within-model coherence rather than external generalization to truly independent experimental or field data.




Fig. 8. Surrogate fit quality with both in-sample and held-out random-point

D. Joint Optimisation Results

Table V reports mean optimal management variables and expected sugar-yield gain under each of the three optimisation scenarios.

TABLE V
MEAN OPTIMAL MANAGEMENT AND SUGAR YIELD ACROSS THE OPTIMISATION SUBSET UNDER EACH SCENARIO.

	Unconstr.	EU N-limit	Budget
Mean N^* (kg/ha)	162	148	135
Mean F^* (frac.)	0.65	0.68	0.72
Mean P^* (days)	+2	+2	+1
Mean I^* (mm)	45	48	20
Mean sugar yield (t/ha)	16.8	16.5	15.9
Mean gain vs advisory rate (%)	2.71	2.45	2.10

While the EU Nitrate limit marginally reduces high-clay optima, budget constraints heavily restrict irrigation across all soils. To resolve calibration inconsistencies, we reran the discrete optimisation (964 fields) using the cluster-specific Mitscherlich calibration (Section III-D). This unconstrained rerun achieved a mean gain of +2.60 t/ha (median +2.20 t/ha;

¹⁰ in-sample $R^2 = 0.999$ reflects the inherent smoothness of the pipeline.

15.0%) over the baseline. Optimisation remains beneficial even in a strict rainfed-only scenario ($I=0$), yielding a $+0.98\text{ t/ha}$ mean gain (median $+0.84\text{ t/ha}$; 5.78%). The budget scenario (Table V) closely tracks the EU-cap case, demonstrating that qualitative scenario rankings are robust to the cost triplet in Eq. 6, though comprehensive economic sensitivity analysis remains necessary. Ultimately, these rainfed results confirm the method's practical relevance for NRW farms lacking routine irrigation (Figure 9).




Fig. 9. Comparison of management decision variables and expected yield across scenarios

E. Comparison with Simpler Baselines

- 1) advisory practice: static 150 kg/ha nitrogen application with regional standard planting dates;
- 2) nitrogen-only optimisation: discrete N-rate optimisation holding planting and irrigation to regional baselines;
- 3) best DOE point: the highest-yielding simulated observation from the 126 Latin Hypercube samples;
- 4) additional heuristic baseline: lookup-table heuristics based purely on soil clay terciles.

Joint continuous optimisation outperformed the best observed DOE sample by an average of 0.15 t/ha, confirming the value of the quadratic interpolation step.

F. Cluster-Level Recommendations

Cluster-level recommendations are summarised as median management profiles under the primary scenario (Table VI).

TABLE VI
MEDIAN OPTIMAL MANAGEMENT BY SOIL OR MANAGEMENT CLUSTER

Cluster	N^* (kg/ha)	P^* (days)	I^* (mm)	Sugar yield (t/ha)
High-Clay Rainfed	140	+4	0	15.2
Sandy Irrigated	175	-2	120	14.8
Loam Optimal	165	0	50	17.5
Advisory baseline	150	0	0	14.1

G. Factor Contribution Analysis

Sequential attribution assigned 45% of the gain to nitrogen, 20% to planting-date adjustment, and 35% to irrigation across the full scenario-unconstrained population (Figure 10). This factorisation is scenario- and soil-type dependent: irrigation optimisation dominates particularly in low-clay soil stratifications with high water deficit risk, where supplemental water

is both available and agronomically effective. For rainfed-only usage (see $I = 0$ scenario), irrigation attribution becomes irrelevant, and gains are driven by nitrogen and planting-date optimisation alone.




Fig. 10. Factor contributions to sugar-yield gain

H. Continuous vs. Discrete Optimisation

Continuous optimisation added 0.15 t/ha relative to the best discrete DOE point, corresponding to an 8% relative improvement in the optimisation margin, with runtime below 10 ms per field. This increment remains smaller than the held-out surrogate MAE scale (0.201 t/ha), so it is interpreted as a modest model-level ranking benefit rather than a standalone agronomic effect size. Discrete search was generally sufficient only for purely unconstrained single-variable N optimisation.

I. Optimal Management Schedules

Figure 11 illustrates representative field-level management schedules from the completed optimisation runs.




Fig. 11. Representative management schedules by cluster and scenario

V. DISCUSSION

A. Practical Implications

The completed experiments indicate that a static advisory strategy leaves measurable sugar-yield potential unused. Relative to the regional baseline (150 kg/ha N, standard planting date, no supplemental irrigation), the cluster-calibrated rerun still identifies consistent field-level upside: in full-population discrete optimisation, the unconstrained scenario yields +2.60 t/ha mean gain (median +2.20 t/ha; median +15.0%). Operationally, the main value is the end-to-end integration: publicly available soil and weather sources are translated into field-level DSSAT simulations, corrected with lightweight ML components, and converted into calendar-ready management actions. This allows the framework to function as a practical decision-support system rather than only an exploratory modelling workflow. The scenario comparison is also decision-relevant for advisory use. It quantifies not only the best achievable outcome but also how much performance is surrendered under policy and economic constraints, providing a transparent trade-off between agronomic intensity and feasibility.

B. Role of Each Framework Component

Factor attribution in this dataset assigns 45% of improvement to nitrogen optimisation, 20% to planting-date shifts, and 35% to irrigation adjustment. This ranking is empirical and scenario-dependent, but it highlights that gains come from coordinated management rather than nitrogen-only tuning. Irrigation recommendations should be interpreted conditionally in NRW. While irrigation contributes strongly in low-clay and deficit-prone fields, a strict rainfed-only scenario ($I = 0$) still retains positive benefit (+0.98 t/ha mean; +0.84 t/ha median; median +5.78%), which improves practical usability for non-irrigating farms. The surrogate layer is critical for practical optimisation. It provides smooth interpolation between DOE points, yields an additional 0.15 t/ha over the best sampled DOE point, and supports fast constrained solving (< 10 ms per field), which is essential for large-scale scenario analysis.

C. Model Limitations and Failure Modes

Several limitations of the current framework should be acknowledged.

For interpretability and deployment relevance, we treat temporal hold-out and structural robustness checks as primary evidence, and pooled cross-validation as secondary in-dataset evidence.

Mitscherlich parameterisation. The calibration evidence and cluster-specific rerun values are reported in Section III-D. Those results are strong enough for comparative scenario analysis, but not strong enough to remove uncertainty in curve shape, so optimal-N recommendations should still be interpreted as parameter-sensitive.

Population-level versus local calibration. The main pipeline still relies on a single global Mitscherlich curve for tractability, whereas the multi-rate evidence indicates that calibration granularity matters. In particular, cluster-specific fits

reduce error relative to a one-size-fits-all curve and therefore support cluster-level rather than population-level parameterisation as the more credible deployment path (Figure 12).

Nitrogen split timing. Split scheduling remains simplified through fixed application windows and aggregate uptake effects. Timing sensitivity is present but secondary to total-N and water constraints in this implementation.

Year-specific bias. Error variance peaks in 2024, consistent with late-season drought and stress heterogeneity. The current correction remains population-level, so residual year effects are not fully removed.

Fixed dry-matter fraction. The fresh-yield conversion uses a single $f_{DM} = 0.22$ value, although dry matter fraction varies with cultivar, harvest timing, and growing conditions. That simplification adds unmodelled noise to the $Y_{DM} \rightarrow Y_{fresh}$ step and is one reason the proportional bias corrector should be viewed as a pragmatic conversion layer rather than a mechanistic submodel.

Stress and disease. Unmodelled biotic stress (e.g., foliar disease[27]) remains a limitation; severe baseline over-predictions appear concentrated in low-yielding 2024 fields, indicating stress dominates failure modes in adverse seasons.

Sugar model uncertainty. Although pooled performance is acceptable ($R^2 = 0.61$), hold-out performance remains lower ($R^2 = 0.54$), and value-range compression persists at extremes. This uncertainty directly affects sugar-yield ranking.

Temporal generalisation of the sugar model. Hold-out validation on 2025 (MAE 0.68, $R^2 = 0.54$) remains positive but weaker than pooled metrics (MAE 0.75, $R^2 = 0.61$), supporting a cautious interpretation: the framework is currently stronger for near-term scenario support than fully out-of-distribution forecasting.

Subset-based optimisation and external validity. Joint optimisation is performed on a stratified 150-field subset and reports model-estimated gains. Results are therefore indicative of comparative potential, not yet confirmed farm-level realised gains.

Nitrogen sensitivity attribution in GB model. Coupling diagnostics show dependence on DSSAT-derived variables: removing DSSAT features drops sugar-model R^2 from 0.61 to 0.48. This is operationally useful but limits fully independent mechanistic attribution.

DSSAT contribution versus simpler pipelines. The present work demonstrates that DSSAT-derived states are useful in the coupled predictor, but it does not yet establish how much of total performance is uniquely attributable to DSSAT versus simpler crop-model surrogates or purely statistical pipelines. A full end-to-end ablation against non-DSSAT baselines is therefore a priority before making stronger claims about DSSAT's incremental value. Using the remote ablation dataset, we can now quantify this increment: the full DSSAT-coupled pipeline substantially outperforms the no-DSSAT and MLR alternatives for sugar-yield prediction (MAE 0.467 vs 0.805 vs 1.017; R^2 0.915 vs 0.736 vs 0.613). This supports a non-trivial DSSAT contribution in the present




Fig. 12. Sensitivity of optimisation outputs. Left: fields assigned to $N^* = 250$ across the $(s_{\min}, c)$ grid. Right: mean sugar-yield gain versus $N = 150$

architecture, while still leaving open whether other process-model alternatives could achieve similar gains.

Uncertainty and circularity checks. The 90% Monte Carlo interval width (2.4 t/ha) and ± 25 kg/ha variation in optimal N under parameter resampling indicate that point-optimal recommendations should be interpreted with confidence bands.

Economic scenario simplification. The budget constraint now makes its provenance explicit: the N coefficient is anchored to a German urea benchmark, the irrigation coefficient is a rounded operating-cost proxy derived from published German field-irrigation energy demand and electricity prices, and the 250 e/ha cap itself reflects local farmer rule-of-thumb rather than audited bookkeeping records. In the present dataset the budget-constrained outcome stays close to the EU-cap scenario, which suggests the headline ranking is not dominated by one exact price triplet, but a formal price-sensitivity analysis is still needed before treating the budget scenario as an economic recommendation.

Why many fields select high N in some settings. In the present calibrated configuration, endpoint concentration is reduced (12% upper-bound share), and lower-N selections appear systematically in low water-holding soils, supporting a more heterogeneous recommendation structure.

Validation of optimised plans. Prospective field-trial data (reported in Section III, 30 control-optimised pairs in 2025) confirm positive mean sugar-yield improvement (+0.71 t/ha, +5.37%) with formal statistical support (95% CI, paired t-test $p < 10^{-7}$, effect size $d_z = 1.37$). While encouraging and reducing anecdotal risk, this single-season validation is not a substitute for broader multi-year independent confirmation.

Circularity in optimisation evaluation. Because optimisation is solved and evaluated on the same surrogate objective, reported gains are model-estimated and partially circular by construction. This is acceptable for scenario screening, but not equivalent to independent confirmation of realised on-farm gains.

Surrogate model generalisation. The quadratic surrogate achieves high in-sample R^2 (0.999) and held-out R^2 (0.979), but both derive from the same smooth DSSAT-ML pipeline. This indicates strong internal consistency and interpolation

stability, but does not constitute external validation against independent observations or alternative crop models. The high fidelity is partly an artefact of the pipeline’s inherent smoothness (Mitscherlich saturation plus bias correction), which imposes functional structure. Independent field-level validation of surrogate predictions, ideally from blind multi-year trials on new farms, remains necessary before deploying surrogates as operational decision tools.

D. Comparison with Related Work

Compared with prior surrogate-based optimisation studies in cereals, this work addresses a different bottleneck: sugar beet requires explicit sucrose prediction and nitrogen-response repair before optimisation is agronomically meaningful. The contribution is therefore not a new optimiser class, but a deployable integration of DSSAT repair, ML augmentation, DOE-based surrogate fitting, and constrained optimisation at field scale. In contrast to regional forecasting platforms such as CRAFT [28], which prioritise macro-scale yield monitoring, the present framework targets field-level management scheduling and can be viewed as complementary rather than competing.

VI. CONCLUSION

This work couples DSSAT with external correction and optimisation layers to form a field-level sugar beet decision pipeline: a Mitscherlich-based nitrogen-response mapping, a data-driven sugar-content predictor, DOE-based surrogate construction, and constrained quadratic optimisation. The output is an actionable management schedule rather than a single recommendation variable. Validation on 964 production fields across 2023–2025 shows robust baseline behaviour (fresh-yield MAE 6.4 t/ha; median error -1.2% ; 88.5% within $\pm 20\%$) and moderate but useful sugar-model transfer (pooled MAE 0.75, $R^2 = 0.61$; 2025 hold-out MAE 0.68, RMSE 0.82, $R^2 = 0.54$). End-to-end ablation on observed sugar yield shows clear incremental value of the DSSAT-coupled pipeline over simpler alternatives (MAE 0.467 vs 0.805 no-DSSAT vs 1.017 MLR). Optimisation results indicate consistent model-estimated gains over advisory management: 2.71% unconstrained, 2.45% under EU-N constraints, and 2.10% under

budget constraints. On the full population, median optimal N is 145 kg/ha with limited upper-bound concentration, while subset optimisation identifies differentiated strategies by soil-water context. A prospective paired-trial dataset provides first external support for realised uplift (+0.71 t/ha sugar yield; +5.37%; 93.3% positive pairs), while the marginal continuous-over-discrete advantage remains small relative to surrogate error scale. The next priorities are expanded multi-year prospective validation, explicit disease/stress integration, stronger temporal robustness of sugar prediction, and cluster-level (or field-level) calibration of nitrogen-response parameters. Until then, the framework is best used as a transparent decision-support and scenario-analysis tool rather than a fully autonomous prescription engine.

ACKNOWLEDGMENT

Supported by the Ministry of Economic Affairs, Industry, Climate Action and Energy of the State of North Rhine-Westphalia. Co-funded by the European Union. Grant number EFRE-20800498.

REFERENCES

- [1] CEFS, "The european sugar sector," 2026, accessed: 2026. [Online]. Available: <https://cefs.org/the-european-sugar-sector/>
- [2] Eurostat, "Sugar beet production statistics in the european union," 2026, accessed: 2026. [Online]. Available: <https://ec.europa.eu/eurostat/>
- [3] A. P. Draycott, *Sugar Beet*. Oxford, UK: Blackwell Publishing, 2006.
- [4] H.-J. Koch, S. Heuer, O. Tomanova, and B. Marlander, "Cumulative effect of annually repeated passes of heavy agricultural machinery on soil structural properties and sugar beet yield under two tillage systems," *Soil and Tillage Research*, vol. 101, no. 1–2, pp. 69–77, 2009.
- [5] R. Gebbers and V. I. Adamchuk, "Precision agriculture and food security," *Science*, vol. 327, no. 5967, pp. 828–831, 2010.
- [6] J. W. Jones, G. Hoogenboom, C. H. Porter, K. J. Boote, W. D. Batchelor, L. A. Hunt, P. W. Wilkens, U. Singh, A. J. Gijsman, and J. T. Ritchie, "The DSSAT cropping system model," *European Journal of Agronomy*, vol. 18, no. 3–4, pp. 235–265, 2003.
- [7] D. P. Holzworth, N. I. Huth, P. G. de Voil, E. J. Zurcher, N. I. Herrmann, G. McLean, K. Chenu, E. J. van Oosterom, V. Snow, C. Murphy, A. D. Moore, H. Brown, J. P. M. Whish, S. Verrall, J. Fainges, L. W. Bell, A. S. Peake, P. L. Poulton, Z. Hochman, P. J. Thorburn, D. S. Gaydon, N. P. Dalgliesh, D. Rodriguez, H. Cox, S. Chapman, A. Doherty, E. Teixeira, J. Sharp, R. Cichota, I. Vogeler, F. Y. Li, E. Wang, G. L. Hammer, M. J. Robertson, J. P. Dimes, A. M. Whitbread, J. Hunt, H. van Rees, T. McClelland, P. S. Carberry, J. N. G. Hargreaves, N. MacLeod, C. McDonald, J. Harsdorf, S. Wedgwood, and B. A. Keating, "APSIM – evolution towards a new generation of agricultural systems simulation," *Environmental Modelling & Software*, vol. 62, pp. 327–350, 2014.
- [8] A. de Wit, H. Boogaard, D. Fumagalli, S. Janssen, R. Knapen, D. van Kraalingen, I. Supit, R. van der Wijngaart, and K. van Diepen, "25 years of the WOFOST cropping systems model," *Agricultural Systems*, vol. 168, pp. 154–167, 2019.
- [9] D. Romero, "DSSATTools: A Python library for DSSAT crop model simulation," 2023, python Package Index. [Online]. Available: <https://pypi.org/project/DSSATTools/>
- [10] M. R. Anwar, D. L. Liu, R. Farquharson, I. Macadam, A. Abadi, J. Finlayson, B. Wang, and T. Ramilan, "Climate change impacts on phenology and yields of five broadacre crops at four climatologically distinct farm locations in Australia," *Agricultural Systems*, vol. 132, pp. 133–144, 2015.
- [11] E. A. Mitscherlich, "Das gesetz des minimums und das gesetz des abnehmenden bodenertrages," *Landwirtschaftliche Jahrbucher*, vol. 38, pp. 537–552, 1909.
- [12] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
- [13] M. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 21, no. 2, pp. 239–245, 1979.
- [14] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, 2026. [Online]. Available: <https://www.gurobi.com>
- [15] Council of the European Union, "Council directive 91/676/EEC concerning the protection of waters against pollution caused by nitrates from agricultural sources," 1991, official Journal L 375, 31/12/1991, pp. 1–8.
- [16] L. Poggio, L. M. de Sousa, N. H. Batjes, G. B. M. Heuvelink, B. Kempen, E. Ribeiro, and D. Rossiter, "SoilGrids 2.0: Producing soil information for the globe with quantified spatial uncertainty," *SOIL*, vol. 7, no. 1, pp. 217–240, 2021.
- [17] K. E. Saxton and W. J. Rawls, "Soil water characteristic estimates by texture and organic matter for hydrologic solutions," *Soil Science Society of America Journal*, vol. 70, no. 5, pp. 1569–1578, 2006.
- [18] Open-Meteo, "Open-meteo historical weather API," 2026, accessed: 2026. [Online]. Available: <https://open-meteo.com/en/docs/historical-weather-api>
- [19] G. A. F. Seber and A. J. Lee, *Linear Regression Analysis*, 2nd ed. Hoboken, NJ: Wiley, 2003.
- [20] IndexBox, "Urea price in germany - 2025," 2026, reports a German average urea FOB export price of 572 USD/t in May 2025; accessed: 2026. [Online]. Available: <https://www.indexbox.io/search/urea-price-germany/>
- [21] Praxis-Agrar and de Witte, Thomas, "Wirtschaftlichkeit der feldbewässerung," 2026, summarises German field-irrigation economics, including 0.4 kWh/m³ for mobile irrigation systems and full costs of 3.17–4.00 EUR per mm and hectare; accessed: 2026. [Online]. Available: <https://www.praxis-agrar.de/pflanze/bewaesserung/wirtschaftlichkeit-der-feldbewaesserung/>
- [22] BDEW, "Bdew-strompreisanalyse januar 2026," 2026, reports a 2025 electricity price of 15.8 ct/kWh for medium-sized industry customers in Germany; accessed: 2026. [Online]. Available: <https://www.bdew.de/service/daten-und-grafiken/bdew-strompreisanalyse/>
- [23] X. Huang, D. Kroening, W. Ruan, J. Sharp, Y. Sun, E. Thamo, M. Wu, and X. Yi, "A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability," *Computer Science Review*, vol. 37, p. 100270, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1574013719302527>
- [24] R. F. Woolson, "Wilcoxon signed-rank test," in *Wiley Encyclopedia of Clinical Trials*, 2008, pp. 1–3. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9780471462422.eoct979>
- [25] W. Tan, K. Li, D. Liu, and W. Xing, "Cercospora leaf spot disease of sugar beet," *Plant Signaling & Behavior*, vol. 18, no. 1, p. 2214765, 2023, PMID: 37209061.
- [26] E. Memic, S. Graeff-Hönniger, O. Hensel, and W. D. Batchelor, "Extending the csm-ceres-beet model to simulate impact of observed leaf disease damage on sugar beet yield," *Agronomy*, vol. 10, no. 12, pp. 1–20, 2020.
- [27] V. Misra, S. Srivastava, A. K. Mall, and S. Srivastava, *Foliar Sugar Beet Diseases and Their Management Approaches in India*. Singapore: Springer Nature Singapore, 2022, pp. 541–564. [Online]. Available: https://doi.org/10.1007/978-981-19-2730-0_26
- [28] S. Vakhtang, J. Hansen, V. Sharda, C. Porter, P. Aggarwal, C. J. Wilkerson, and G. Hoogenboom, "A multi-scale and multi-model gridded framework for forecasting crop production, risk analysis, and climate change impact studies," *Environmental Modelling and Software*, vol. 115, pp. 144–154, 2019.