

Article

Data-Driven Manufacturing: Reduced Models and Optimization in Glass Fiber Processes

Ashreet Mishra , Alex Lohse, Steve D. C. Pham , Ravi Arora and Bruno A. Purnode *

Owens Corning Science & Technology Centre, Granville, OH 43023, USA;

alexander.lohse@owenscorning.com (A.L.); mpedcpham@gmail.com (S.D.C.P.); arorar_2000@yahoo.com (R.A.)

* Correspondence: ashreet.mishra@owenscorning.com (A.M.); bruno.purnode@owenscorning.com (B.A.P.)

Abstract

Joule-heated electric glass melters are highly energy-intensive and operate under stringent throughput, quality, and refractory-protection constraints, yet existing optimization approaches either rely on offline high-fidelity CFD—which is too slow for day-to-day operating decisions—or on purely data-driven energy models that do not enforce physics-based glass-quality constraints. This work presents a simulation-informed supervisory optimization framework that closes this gap by integrating (i) a high-fidelity CFD design of experiments on a cold-top electric glass melter, (ii) interpretable linear reduced-order models (ROMs) for furnace temperatures and glass-process indices (residence time, sand dissolution, mixing, and melting indices), and (iii) a convex linear-program supervisory optimizer that minimizes electrical power subject to throughput, glass-quality, and refractory-protection constraints. Sixteen CFD operating points spanning power inputs of 900–1671 kW, pull rates of 0.12–0.36 kg/s, and cullet fractions of 0–30% are used to fit and cross-validate the ROMs. The novelty of this contribution is the unification of CFD-derived multi-KPI quality constraints with a fast, interpretable supervisory optimizer suitable for advisory deployment on industrial electric melters. For a representative 1350 kW baseline, the optimizer recommends a reduced operating power of 1256 kW (~7% reduction) while satisfying all quality, throughput, and refractory constraints—demonstrating a practical ROM-enabled digital-twin pathway for energy-efficient electric glass melting.

Keywords: electric glass melter; computational fluid dynamics; reduced-order model; joule heating; supervisory optimization; digital twin



Academic Editor: Raul D. S. G. Campilho

Received: 29 May 2026

Revised: 8 July 2026

Accepted: 17 July 2026

Published: 20 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Glass fiber manufacturing is an energy-intensive process under mounting pressure to improve efficiency, reduce emissions, and enhance operational robustness—all while maintaining strict product quality and maximizing furnace longevity [1–3]. Melting furnaces, which consume on the order of several GJ per ton of glass, are central to both energy costs and CO₂ emissions, making them prime targets for process optimization and decarbonization efforts [1–4]. Indeed, recent assessments indicate that many furnaces operate above theoretical best efficiency, suggesting substantial scope for improvement in practice [4]. To meet these challenges, manufacturers have begun embracing advanced operational strategies that combine physics-based modeling and data-driven methods, aligning with broader Industry 4.0 initiatives to digitize and optimize manufacturing processes [2]. The present work contributes to this movement by proposing a simulation-informed, data-

driven framework for electric glass melters aimed at systematically reducing energy usage while upholding throughput and quality requirements.

High-fidelity computational fluid dynamics (CFD) has emerged as a powerful tool for analyzing glass melting furnaces, capturing the complex multiphysics environment of industrial melters [3,5,6]. In an electric cold-top melter, volumetric resistive heating (Joule heating) generates strong buoyancy-driven convection in the molten glass, inherently coupling electrical, thermal, and flow fields within the furnace [7,8]. Modern CFD models can resolve these interactions in detail: they account for electric potential distributions in the melt, heat transfer by conduction, convection, and radiation, as well as fluid flow patterns that are influenced by temperature-dependent viscosity and density variations [3,7]. Numerous studies illustrate the insights available from such high-fidelity simulations. For example, Chang et al. [9] incorporated multiphase effects (batch blanket and gas bubbles) in a glass melter model, showing how these phenomena alter melt flow structures and temperature fields. Likewise, Matsuno et al. [10] coupled electric field calculations with CFD to investigate how electrode placement and power input affect temperature uniformity, finding that even small changes in electrode configuration can significantly impact under-glass temperatures and melting rates. These and other analyses confirm that furnace performance metrics—such as thermal homogeneity, melt residence times, and defect formation tendencies—are highly sensitive to design and operating parameters [9,10]. Foundational studies (e.g., Stanek [8]) have long recognized this sensitivity in electric melting, and subsequent modeling advancements have continued to increase fidelity. Choudhary et al. [11], for instance, developed detailed mathematical models encompassing flow and heat transfer from melting through fiber forming, highlighting the complex couplings and long-time scales involved in achieving steady production. Additional refinements, such as incorporating cold-cap (batch blanket) behavior and gas release models, have also been explored to better predict melting rates and foam formation [12]. Beyond understanding current operations, CFD has been pivotal in evaluating new furnace technologies and decarbonization strategies. Simulations have assessed the effects of electric boosting, hybrid heating, and full electrification on furnace throughput and energy efficiency, informing the tradeoffs between reduced emissions and product quality under various scenarios [1,13]. Together, these efforts underscore that high-fidelity modeling provides deep process insight—from fundamental melting physics to the strategic evaluation of energy and carbon reduction pathways.

However, direct use of full CFD in day-to-day furnace operation remains impractical, primarily due to computational cost and complexity. Industrial glass furnaces are large, with melt volumes and flow structures that require millions of computational cells to simulate accurately. As a result, a single steady-state CFD run can take on the order of many hours (if not days) on high-performance computing resources [14]. For example, Abbassi and Khoshmanesh [15] modeled a full-scale glass melter and noted that even with careful problem partitioning, the simulation required substantial computing time and careful numerical strategies to converge. This high latency means that plant engineers cannot feasibly run a new CFD simulation for each potential adjustment or disturbance—they must rely instead on static operating recipes or heuristic adjustments. The gap between detailed modeling and agile operations leaves potential efficiency gains untapped. There is a clear need for surrogate modeling techniques that preserve the essential physics-based sensitivities of CFD but can be evaluated almost instantaneously, enabling rapid “what-if” analysis and data-driven optimization in an industrial setting.

Reduced-order models (ROMs) provide a promising solution by distilling complex CFD results into fast, lower-dimensional approximations [16,17]. These can take various forms—e.g., linear regression models, projection-based reduced-basis models, or machine

learning surrogates—but the common goal is to capture dominant input–output relationships of the process with drastically less computation. ROMs have been successfully used in related process systems to enable near-real-time predictions. For instance, Aversano et al. [18] built a ROM-driven model of a flameless combustion furnace, accurately replicating CFD-predicted temperature fields at a fraction of the cost, thereby enabling rapid scenario evaluations. More generally, frameworks for industrial digital twins emphasize combining simplified physics models with live data to inform operations in real time [2,19,20]. A case in point is Molinaro et al. [21], who integrated CFD simulations with data analytics in a digital twin context, demonstrating how first-principles models can be coupled with plant data to continuously monitor and optimize a chemical process. These examples show that embedding physics knowledge into agile, data-supported models can allow advanced decision support and control, without running a full simulation for each new scenario. In the context of glass furnaces, such surrogate-assisted approaches hold particular promise for bridging the gap between high-fidelity analysis and mill operations.

In practice, however, surrogate modeling for glass melting has so far been limited in scope, and further integration is needed to fully harness its benefits. Some recent efforts have tackled specific sub-problems: for example, Zybala et al. [22] optimized an electric furnace's electrode configuration by using a surrogate model derived from an electromagnetic field simulation, showing the efficiency of surrogate-based design but focusing on a narrow aspect of furnace physics. On the data-driven side, industry-oriented studies have explored empirical modeling to cut energy consumption. Choudhary [6] highlighted opportunities for furnace energy reduction via data analytics in manufacturing operations, and Kodak et al. [23] developed linear statistical models from plant data to optimize a furnace's specific energy consumption under production constraints. These data-centric approaches demonstrate measurable energy savings through careful tuning of operating variables, yet they typically lack direct enforcement of physics-based quality or safety constraints. Meanwhile, other high-fidelity studies have begun linking operating conditions to glass-quality metrics: for instance, Raič et al. [24] used CFD to analyze how adding electric boosting in an oxy-fuel furnace influences indicators like minimum residence time and melting index, finding that increased power input improves melting completeness via stronger melt convection (while naturally incurring higher energy usage). However, such insights were not embedded in a real-time optimization framework; they remain off-line analyses that operators must interpret qualitatively.

Recent developments since 2022 have further accelerated the integration of data-driven methods into high-temperature process modeling. Machine-learning surrogates and physics-informed neural networks have been increasingly applied to combustion furnaces and melting systems for rapid output prediction and design exploration, while digital-twin frameworks combining CFD with real-time plant data have been demonstrated for chemical and metallurgical processes [21]. In the glass sector specifically, Miserocchi et al. [4] benchmarked the energy performance of industrial glass furnaces and identified substantial efficiency gaps relative to theoretical minima; Atzori et al. [25] reviewed CO₂-reduction technologies applicable to glass furnaces, including electrification, hydrogen, and oxy-fuel combustion; and Kodak et al. [23] demonstrated statistical surrogate-based optimization of furnace-specific energy consumption using plant data. Despite these advances, no published framework to date has unified CFD-derived multi-KPI quality and refractory constraints with a fast, interpretable supervisory optimizer for an electric cold-top melter—which is the specific gap addressed by the present work.

Compared with recent ML-based furnace-surrogate studies, which focus primarily on rapid prediction of output fields or KPIs without embedded constraint enforcement, the present framework is distinguished by its explicit integration of CFD-derived multi-KPI

quality and refractory constraints inside a convex supervisory optimizer. Compared with statistical energy-modeling studies that minimize furnace-specific energy consumption from plant data alone, the present framework retains a direct physics linkage through CFD-derived ROM inputs and quality indices, ensuring that recommended energy reductions remain consistent with first-principles melting, mixing, and refractory-protection behavior.

Despite substantial progress in CFD modeling, data-driven energy optimization, and surrogate-based design of glass furnaces, three specific limitations remain in the literature. First, high-fidelity CFD studies of electric and electrically boosted melters provide deep physical insight but remain offline analyses that cannot be directly coupled with real-time or supervisory decision making due to computational cost. Second, data-driven and statistical energy-optimization studies on industrial furnaces do not explicitly enforce CFD-derived glass-quality or refractory-protection constraints. Third, existing surrogate-based optimization studies for glass furnaces have addressed narrow sub-problems—such as electrode geometry—rather than full operating-point optimization across coupled thermal, mixing, and melting KPIs. The present work is motivated by the absence of an integrated framework that bridges all three of these gaps for an electric cold-top melter.

The specific contributions of this paper are:

1. A structured CFD design of experiments for a cold-top electric glass melter spanning realistic industrial ranges of electrical power, pull rate, and cullet fraction.
2. The development of interpretable linear ROMs for both furnace temperatures (T_{batch} , $T_{\text{underglass}}$, T_{outlet}) and CFD-derived glass-process indicators (RT, SDI, MI, MEI), validated using leave-one-out cross-validation.
3. A convex linear-program supervisory optimizer that minimizes electrical power subject to physics-based constraints on glass quality, throughput, and refractory protection.
4. A demonstration on a representative operating baseline showing a ~7% reduction in electrical power without compromising quality or throughput, together with an explicit discussion of ROM uncertainty and bounded-domain applicability.

2. Modeling Approach

This section describes the modeling and computational framework used to construct a digital-twin surrogate for a Joule-heated cold top electric glass melter. Section 2.1 introduces the melter configuration, geometry, and key modeling assumptions adopted to balance physical fidelity and computational tractability. Next, the coupled CFD governing equations for melt flow, heat transfer, and electric-field (Joule heating) calculations are summarized in Section 2.2. The temperature- and quality-related quantities extracted from each CFD case, including local temperature metrics and glass-process indices used for screening and optimization, are defined in Section 2.3. Using these outputs, the linear reduced-order model (ROM) formulation and regression structure are presented in Section 2.4. Lastly, Section 2.5 details the supervisory optimization routine that leverages the ROMs to identify operating set-points that reduce electrical power while maintaining throughput, glass quality, and refractory protection constraints (Figure 1).




Figure 1. Schematic of the simulation-informed supervisory optimization framework: CFD design of experiments → KPI extraction → linear ROM fitting and LOOCV → convex LP optimization with KPI and throughput constraints → recommended supervisory set-point.

2.1. Electric Melter Configuration, Model Assumptions and Geometry

Modeling Joule-heated electric melters requires approximations that balance physical accuracy with computational feasibility. Insights from CFD, electric field modeling, and waste glass melter research provide common assumptions used in scientific and industrial simulations and are used as follows:

1. Electric potential fields are solved under quasi-static conditions, neglecting displacement currents due to the low-frequency AC typical of industrial melters [10,14].
2. Electrodes are prescribed as fixed-potential or fixed-current boundaries, matching typical industrial power-supply strategies [10].
3. Heat transfer is dominated by Joule heating, with conduction, convection, and radiation accounted for through the energy balance [3,7,14,26,27].

4. Glass is modeled as a Newtonian fluid with temperature-dependent viscosity, capturing buoyancy-driven convection induced by volumetric heating [3,7,8].
5. Transient effects are simplified to steady-state for the present operating-window study, consistent with prior CFD analyses of industrial melters [3,15].
6. The batch/cold-cap layer is represented using effective thermal properties rather than resolving microstructure, following standard cold-cap modeling practice for waste and commercial Joule-heated melters [12].
7. Chemical decomposition, gas release, and foam formation are not resolved, in line with prior CFD studies of similar electric melters [3,12].
8. The melter superstructure is approximated using a heat-transfer coefficient rather than being explicitly modeled [3,14].
9. Batch thickness is held constant for all cases.

For the cold top electric melter, a typical example of a soda lime glass melter has been utilized in this work as shown in Figure 2. The diameter of the melter is 4 m, height is 2 m, and batch height is 0.5 m. The wall refractories, throat insulation and crown are made up of chamotte, and the batch is typical of soda lime glass. The six electrodes are used which are made up of molybdenum and fire in three phases. The electrodes are cooled 5500 W/m^3 each to prevent overheating of electrodes. All insulation walls are exposed to ambient air and experience natural convection. A total of 6.8 million grid cells was used for the finite volume solver. Glass flow inlet and outlet boundaries are provided with respective pull rates and inlet temperature is 20°C . A no-slip boundary condition is applied for the glass over refractory and electrodes. The solver utilized is Glass Services software Glass-Furnace Model (GFM) version 4.26 to inculcate the governing equations and generate the required results when mass and heat balance is less than 0.1% for the whole domain and glass outlet temperature for consecutive iterations is less than 10°C .




Figure 2. Cut cross-section of cold top electric melter geometry.

The Glass-Furnace Model (GFM), developed by Glass Service, is the industrial multiphysics furnace-simulation platform used in this work to solve the coupled melt-flow, heat-transfer, and electric-heating equations for glass-furnace analysis [1].

2.2. CFD Governing Equations

Modeling electric glass melters requires solving a coupled system that includes fluid flow, heat transfer, and electric field distribution, since Joule heating depends on the local electrical conductivity and electric-potential gradients [14,26–28].

For incompressible molten glass, the continuity equation is given by Equation (1):

$$\nabla \cdot \mathbf{u} = 0, \quad (1)$$

This is the standard assumption in CFD modeling of molten glass in Joule-heated melters, which are dominated by buoyancy-driven flow rather than compressibility effects.

The momentum equation typically used is given by Equation (2):

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} + \rho \mathbf{g} + \mathbf{F}_{\text{Joule}}, \quad (2)$$

where

- ρ = density;
- μ = dynamic viscosity;
- $\mathbf{F}_{\text{Joule}}$ = volumetric body force due to electromagnetic effects.

Most industrial melter studies note that Lorentz forces $\mathbf{F}_{\text{Joule}}$ are small compared to buoyancy forces, and therefore often neglected in the momentum equation.

The flow is thus predominantly buoyancy-driven, caused by temperature-induced density gradients from internal Joule heating.

The temperature distribution is governed by the energy Equation (3):

$$\rho c_p \left(\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T \right) = \nabla \cdot (k \nabla T) + Q_{\text{Joule}}, \quad (3)$$

where the Joule heating term is given by Equation (4):

$$Q_{\text{Joule}} = \sigma |\mathbf{E}|^2, \quad (4)$$

This formulation is validated in CFD electric field coupled studies of Joule-heated glass melters, which explicitly include Joule heating as the dominant energy source.

Joule heating is computed by solving the electric potential field in the glass:

$$\nabla \cdot (\sigma \nabla \phi) = 0 \quad (4a)$$

where

- ϕ = electrical potential;
- σ = electrical conductivity.

This equation is explicitly stated in electric field modeling of glass melts, with boundary conditions applied at electrodes and insulating walls.

The electric field is then:

$$\mathbf{E} = -\nabla \phi, \quad (5)$$

And the Joule heat source:

$$Q_{\text{Joule}} = \sigma (\nabla \phi \cdot \nabla \phi), \quad (6)$$

Equations (1)–(6) constitute the standard incompressible, Newtonian, buoyancy-dominated formulation for Joule-heated melts widely used in glass-furnace CFD model-

ing [3,7,8,10,14,26,27]. Temperature-dependent material properties ($\mu(T)$, $\sigma(T)$, $\rho(T)$) are taken from the soda-lime-glass property database embedded in the Glass-Furnace Model solver [1], and electrode boundary conditions follow industrial three-phase Joule-heating practice [10].

In an electric glass melter, GFM solves the melt flow and energy transport in a fully coupled manner, where fluid motion, temperature, and material properties evolve together. Heat is generated internally through electrical boosting and distributed volumetrically within the glass, while convection and conduction govern heat transport through the melt. Radiative heat transfer within the hot, semi-transparent glass and between internal surfaces is accounted for using an effective thermal conductivity (radiation-enhanced conductivity) approach, embedding radiative effects directly into the energy balance rather than solving a separate radiation model. Because viscosity and electrical conductivity depend strongly on temperature, the flow, thermal, and electrical fields are tightly coupled and solved iteratively with property updates until thermal balance and melt circulation converge. Using the results then, the metrics of interest are obtained in Section 2.3.

For clarity, the present model is a steady-state, three-dimensional, single-phase, incompressible melt-flow model with coupled heat transfer and electric-potential calculations for Joule heating. The molten glass is treated as a temperature-dependent Newtonian fluid, and no separate turbulence model is employed because the flow is assumed to be dominated by viscous and buoyancy-driven circulation under the operating conditions considered. Radiative effects within the hot glass are incorporated through an effective thermal conductivity formulation rather than by solving a separate radiation transport equation. The cold-cap/batch layer is represented using effective thermal properties and fixed thickness. Explicit multiphase behavior, gas evolution, foam formation, and detailed batch chemistry are not resolved. Joule heating is obtained from the electric-potential solution and coupled into the energy equation as a volumetric heat source.

2.3. Metrics of Interest

In high-fidelity glass-furnace modeling [14,26], the identification and systematic evaluation of relevant performance metrics is essential for translating numerical simulations into actionable operational insights. Within the present framework, a DOE coupled with a ROM is constructed using a validated GFM CFD solver. The objective is to quantify the sensitivity of key thermal and process indicators to controllable operating parameters and to establish surrogate relationships that enable rapid assessment across a broad operating envelope.

The DOE explores a multidimensional operating space defined by three primary control variables that directly influence furnace thermal behavior and glass quality:

- Power input (900–1800 kW): Represents the total energy supplied to the furnace and governs overall heat availability, melting rates, and temperature uniformity.
- Pull rate (0.10–0.40 kg·s^{−1}): Controls glass throughput and residence time, thereby influencing thermal stratification, melting completeness, and outlet conditions.
- Cullet fraction (0–30%): Alters batch thermophysical properties and melting kinetics, with direct implications for energy efficiency and temperature distributions.

These inputs in Table 1 are selected to reflect realistic industrial operating ranges and to capture both nominal and off-design operating conditions [26] which are then subsequently used in the ROM and optimization scheme.

Table 1. Inputs and outputs of the supervisory ROM framework.

Variable	Type	Range/Definition	Units
Pull rate	Input	0.12–0.36	kg/s
Electrical power	Input	900–1671	kW
Cullet fraction	Input	0–30	%
T_batch	Output	Crust/batch interface temperature	°C
T_underglass	Output	Glass-refractory interface temperature	°C
T_outlet	Output	Throat outlet glass temperature	°C
RT_10%, RT_AVG	Output	Residence time (10th pct, mean)	h
SDI_10%, SDI_AVG	Output	Sand dissolution index	µm
MI_10%, MI_AVG	Output	Mixing index	–
MEI_10%, MEI_AVG	Output	Melting index	–

The model uses three controllable inputs and produces 11 output quantities of interest.

The primary metrics of interest in Table 1 are temperature-based quantities that serve as direct indicators of furnace health, efficiency, and glass quality:

- **Under-Glass Temperature**—The under-glass temperature is a critical metric reflecting heat transfer from the melt to the furnace bottom. Excessively high under-glass temperatures may indicate increased refractory wear and reduced furnace lifetime, while low values can signal insufficient melting or poor heat utilization. Tracking this metric across the DOE space enables identification of operating regimes that balance melting performance with refractory protection.
- **Crust Temperature**—Crust temperature provides insight into batch melting dynamics and surface heat transfer. It is particularly sensitive to power input and cullet fraction, as increased cullet generally lowers melting temperatures due to reduced endothermic load. Maintaining crust temperature within an optimal window is essential to avoid cold caps that impede melting or excessive surface heating that leads to volatilization losses.
- **Throat Temperature**—The throat temperature directly affects glass flow uniformity and downstream forming processes. It serves as a proxy for thermal consistency at the furnace exit. Variations in pull rate strongly influence this metric, making it a key constraint when optimizing throughput without compromising product quality.
- **Derived Glass Key Performance Indicators (KPIs)**—In addition to local temperature metrics, a set of aggregated Glass KPIs is evaluated to provide holistic assessment of furnace performance. These KPIs may include average melt temperature, thermal gradients, or stability indicators derived from the CFD field data. Such metrics are particularly well-suited for ROM development, as they compress high-dimensional CFD outputs into physically meaningful quantities that are directly relevant to operations and control.
- **Residence time (RT)** is the time which a particle spends on its way between a starting and a target plane, sand dissolution index (SDI) is the maximum size of a particle that can be dissolved during its travel between a starting and a target plane, mixing index (MI) is the frequency at which a 1 cm thick cord can be diluted during its travel between a starting and a target plane, with minimum desired value of 1, and melting index (MEI) defined as a field-integrated indicator of how much of the batch has transitioned to fully molten, homogeneous glass. These are used to quantify melting performance and glass quality in electric glass melting furnaces. RT characterizes the exposure time of batch and melt to the furnace thermal and flow fields and is directly influenced by pull rate and electrically driven circulation. SDI measures the extent of silica dissolution and serves as a proxy for melting completeness, reflecting both local temperature and thermal residence in high-energy regions. MI evaluates thermal and

compositional homogeneity within the melt, indicating the effectiveness of buoyancy- and Joule-driven mixing, while MEI provides an integrated measure of overall melting efficiency by combining thermal exposure, RT, and batch transformation behavior.

This data obtained are now used along with the temperature data in the next ROM formulation as the baseline datasets.

2.4. ROM Formulation

A linear reduced order model (ROM) is a surrogate representation that approximates the input–output behavior of a high-fidelity model (e.g., CFD) using a linear mapping between a set of explanatory variables and quantities of interest (QoIs) [17]. In the context of a DOE workflow, the high-fidelity solver is evaluated at sampled operating conditions, and the resulting QoIs are used to fit a model of the form

$$\mathbf{y} \approx \mathbf{B} \boldsymbol{\phi}(\mathbf{x}) + \boldsymbol{\varepsilon} = 0, \quad (7)$$

where $\mathbf{x}$ denotes the input vector (e.g., power, pull rate, cullet), $\mathbf{y}$ contains the QoIs (e.g., RT, SDI, MI, MEI, temperatures), $\boldsymbol{\phi}(\mathbf{x})$ is a chosen set of basis functions (often including a constant term and the inputs, and frequently augmented with interaction and quadratic terms while remaining linear in coefficients), and $\mathbf{B}$ is estimated via least squares or regularized regression. This structure yields a compact model that is inexpensive to evaluate and easy to interpret: coefficients directly encode sensitivity, enabling rapid screening of dominant effects and first-order tradeoffs.

Despite its simplicity, a linear ROM is often effective for local operating envelopes where the response varies smoothly with inputs. It is especially attractive in engineering settings because it provides (i) transparent physics aligned interpretation (sign and magnitude of coefficients), (ii) straightforward uncertainty quantification through residual analysis, and (iii) robust integration with DOE methods for effect estimation and interaction testing. In practice, model adequacy is assessed using standard diagnostics (e.g., cross-validation error, residual normality/heteroscedasticity checks, and leverage/influence), and predictive reliability is bounded by the sampled DOE domain. When strong nonlinearities are present (e.g., regime shifts in flow or phase change), linear ROMs are typically extended using piecewise linear models, enriched polynomial feature sets, or complemented with nonlinear surrogates—while still retaining the linear ROM as a baseline for interpretability and fast iteration.

Although glass melting involves strongly nonlinear physics—including temperature-dependent viscosity, electrical conductivity, buoyancy-driven recirculation, and batch decomposition—a linear ROM was selected here for four specific reasons. (i) Sample efficiency: only sixteen high-fidelity CFD cases are available, and higher-order polynomial, kriging, or neural surrogates would be prone to overfitting at this sample size. (ii) Interpretability: the sign and magnitude of each regression coefficient directly encode the dominant first-order sensitivity of each output to power, pull, and cullet, which is essential for industrial trust and operator review. (iii) Optimization tractability: linearity yields a convex linear program at the supervisory layer, allowing near-instantaneous evaluation of candidate operating set-points. (iv) Bounded operating envelope: within the DOE-validated domain considered here, first-order trends dominate, and the LOOCV diagnostics in Section 3.3 confirm that linear ROMs reproduce the directional sensitivities adequately. The principal limitation of this choice is reduced accuracy in regions exhibiting strong nonlinearity—most notably batch melting dynamics, which manifest in the higher T_{batch} error. Future extensions of this framework will explore piecewise-linear, Gaussian-process, and neural-surrogate ROMs as the CFD DOE is expanded.

It is emphasized that linearity is a modeling choice, not a physical claim. The underlying coupled flow, heat-transfer, and electric-potential physics is inherently nonlinear; the linear ROMs employed here are best understood as first-order Taylor-expansion-style surrogates that approximate the dominant directional sensitivities of T_{batch} , $T_{\text{underglass}}$, and T_{outlet} to the three control inputs within the bounded DOE envelope. The adequacy of this approximation is verified ex post by the LOOCV diagnostics in Section 3.3, which show acceptable directional fidelity for $T_{\text{underglass}}$ and T_{outlet} and reduced fidelity for T_{batch} —consistent with the expected stronger nonlinearity of batch-region dynamics.

2.5. Optimization Routine

The reduced order models developed in this work provide an explicit, low latency mapping between furnace control inputs and key process and quality indicators. This capability enables their direct use within a supervisory optimization layer, where operating set points are adjusted to reduce electrical energy consumption while maintaining stable furnace operation and glass quality consistent with an established baseline condition.

In an industrial context, the electric melter is typically operated through a combination of regulatory control loops (e.g., electrode power, glass pull rate) and higher level supervisory targets defined by throughput, temperature limits, and quality constraints. The present optimization framework is formulated to emulate this structure by treating the ROM as a digital twin surrogate that evaluates candidate operating set points prior to implementation.

A baseline operating point is first defined, corresponding to a nominal furnace condition drawn from the validated CFD dataset. This baseline represents an accepted steady state operating regime that satisfies production, quality, and equipment-protection requirements. The ROM-predicted values of KPIs at this point are used to define target operating windows for optimization.

The KPIs considered include under-glass temperature, outlet temperature, residence time (RT), sand dissolution index (SDI), mixing index (MI), and melting index (MEI). From an operational standpoint, these metrics serve as proxies for refractory protection, thermal stability at the throat, melting completeness, and melt homogeneity. During optimization, each KPI is constrained to remain within a predefined tolerance band around its baseline value, reflecting allowable operational variability without compromising furnace health or product quality.

The optimization is formulated in terms of the primary manipulated variables available to operators or supervisory control systems:

- Electrical power input, adjusted via electrode power set points;
- Glass pull rate, defining throughput and residence time;
- Cullet fraction, influencing thermal efficiency and melting load.

The optimization objective is to minimize electrical power input, which serves as a direct surrogate for furnace energy consumption:

$$\min \text{Power} \quad (8)$$

This formulation aligns with common industrial energy reduction strategies, where power is reduced subject to maintaining throughput and quality constraints.

Operational constraints are imposed to ensure feasibility and control relevance in Equations (9)–(11):

- KPI constraint envelopes.

Each KPI is required to remain within a specified deviation from its baseline value:

$$| \text{KPI}_i - \text{KPI}_{i,\text{baseline}} | \leq \Delta_i, \quad (9)$$

These envelopes represent control limits rather than strict equality constraints, acknowledging normal process variability and measurement uncertainty.

- Throughput protection constraint.

To prevent energy reduction through trivial production loss, the pull rate is constrained to remain at or above the baseline value:

$$\text{Pull} \geq \text{Pull}_{\text{baseline}}, \quad (10)$$

This ensures that energy savings are achieved through improved operating efficiency rather than reduced output.

- Operating window constraints.

All manipulated variables are restricted to the validated ROM training domain, corresponding to safe and realistic industrial operating ranges. This ensures that recommended set points remain within regions of known furnace stability and model validity.

Because the ROMs are linear in the manipulated variables and the objective function is linear in power, the resulting optimization problem is convex and computationally inexpensive. This allows rapid evaluation of multiple candidate operating scenarios, making the approach suitable for real-time advisory systems or offline set point optimization studies.

From a control perspective, the optimization output defines a revised set of steady state operating targets (power, pull rate, cullet fraction) that can be passed to existing regulatory or model-predictive control layers. The ROM thus functions as a supervisory digital twin, enabling operators and control systems to assess energy–quality tradeoffs before implementing changes on the physical furnace.

While the present implementation is steady state and limited to the linear ROM domain, the framework establishes a clear pathway toward closed-loop digital twin optimization, where ROM-based predictions support continuous energy minimization under KPI and operational constraints.

In practical deployment, the ROMs developed here are intended to operate at the supervisory decision-support layer rather than replace existing furnace control loops. In such a workflow, current operating measurements and target production conditions would be provided to the ROM-based surrogate, which would rapidly evaluate candidate operating points within the validated domain and recommend updated supervisory targets for power input, pull rate, or cullet fraction. These recommendations could then be reviewed by operators or passed to existing regulatory or advanced control layers. Accordingly, the present framework should be viewed as a supervisory digital-twin implementation for rapid operating-window evaluation and advisory optimization, with future extensions required for direct plant-data assimilation and closed-loop control. A limitation of the present study is the relatively small CFD design-of-experiments dataset, which is appropriate for constructing interpretable first-order surrogate models over a bounded operating envelope but may not capture stronger nonlinearities that would require expanded sampling or more flexible surrogate architectures.

3. Results

The high-fidelity CFD dataset and the reduced-order modeling framework are used in this section to quantify the impact of key operating variables on electric melter thermal behavior and glass-quality indices. Section 3.1 describes the CFD design-of-experiments

dataset, including the operating envelope and the post-processed temperature fields and scalar indices extracted from each run. Next, linear temperature ROMs for batch, under-glass, and outlet temperatures are developed in Section 3.2, and their predictive accuracy is assessed via leave-one-out cross-validation in Section 3.3. The ROM development is then extended to glass-process indices (RT, SDI, MI, and MEI) in Section 3.4 to provide compact, interpretable surrogates for melting and mixing performance. Lastly, Section 3.5 combines the temperature and glass-index ROMs in a supervisory-level optimization to identify operating set-points that reduce electrical power demand while maintaining throughput, glass quality, and refractory protection constraints.

3.1. CFD Data

The dataset used for reduced-order model (ROM) development is constructed from a series of high-fidelity CFD simulations of an electric glass melting furnace, each corresponding to a prescribed combination of pull rate, electrical power input, and cullet fraction. A total of sixteen operating conditions are evaluated, spanning pull rates from 0.12 to 0.36, power inputs from 900 to 1671 kW, and cullet fractions from 0 to 30%. The selected cases provide coverage of both interior operating points and edge conditions within the practical furnace operating envelope, including low-power/high-throughput and high-power/high-cullet scenarios.

The 16 operating conditions were selected as a structured CFD design-of-experiments over the feasible industrial operating envelope to capture the main effects of the three control variables (power input, pull rate, and cullet fraction) while keeping the number of high-fidelity simulations computationally manageable. The selected bounds correspond to realistic operating windows for the target melter and include both nominal and edge-condition scenarios to support surrogate construction over the bounded domain.

The number of CFD samples was selected to balance two competing constraints: each high-fidelity Glass-Furnace Model run requires tens of hours of HPC time on a 6.8-million-cell mesh, while at least ~5 samples per input dimension are required to support a 3-input main-effects regression. Sixteen cases satisfy this rule of thumb while remaining computationally feasible. The cases were generated using a space-filling design over the feasible industrial operating envelope and explicitly augmented with corner/edge points (Cases 12–16: extreme pull, extreme cullet, and extreme power combinations) so that the surrogate is bounded by physically validated CFD limits. The input ranges—power 900–1671 kW, pull 0.12–0.36 kg/s, cullet 0–30%—correspond to realistic operating windows for the target soda-lime cold-top melter and are consistent with the ranges reported in prior CFD-based studies of similar industrial systems [1,21,23,26].

For each operating condition, the CFD model resolves the three-dimensional temperature and flow fields within the melt pool and throat region. Representative temperature contours for selected cases are shown in Figure 3, illustrating the substantial variation in thermal structure captured across the dataset. Differences in power input and pull rate produce distinct temperature distributions, particularly in the throat and lower melt regions, while variations in cullet fraction influence overall melt temperature levels and thermal gradients. These field-level differences form the physical basis for the derived process and quality metrics used in the ROM.

The four panels of Figure 3 illustrate how variations in the three control inputs produce qualitatively distinct thermal structures across the DOE. Panel (a) demonstrates the dominant role of high power input in elevating under-glass and outlet temperatures, while panel (b) shows that low pull rate substantially increases melt residence time and batch-region temperature. Panel (c) corresponds to the lowest-energy point in the dataset, with markedly reduced under-glass temperature, while panel (d) highlights the moderating effect of high

cullet content on batch temperature even at low power input. Together, the four panels provide the physical basis for the scalar indices (Figure 4) used in the ROM.




Figure 3. Representative side-view CFD temperature contours of the cold-top electric melter for four DOE cases illustrating the dataset variability: (a) Case 1—high-power (1607 kW), high-pull (0.35 kg/s), high-cullet (27.8%) showing strong central convective plume and elevated under-glass temperature; (b) Case 4—moderate-power (1221 kW), low-pull (0.14 kg/s), low-cullet (2.1%) showing high batch and under-glass temperatures consistent with long residence time; (c) Case 8—low-power (1028 kW), high-pull (0.29 kg/s), zero-cullet showing the coolest temperature field of the dataset; (d) Case 13—low-power (900 kW), low-pull (0.12 kg/s), high-cullet (30%) showing the moderating effect of cullet fraction on batch and outlet temperatures.




Figure 4. Temperature metrics extracted from the 16 CFD operating conditions used for ROM development: (a) batch temperature, (b) under-glass temperature, and (c) outlet temperature. The x -axis denotes the CFD case index in the DOE dataset; each case corresponds to a specific combination of pull rate, power input, and cullet fraction listed in Table 2.

This behavior is consistent with prior CFD studies of Joule-heated and electrically boosted glass furnaces, which have shown that electrical power distribution, electrode configuration, and throughput strongly affect melt circulation, thermal stratification, and outlet temperature uniformity [9–11,24]. The influence of cullet fraction on the apparent

thermal load of the batch is likewise consistent with the known dependence of furnace energy performance on batch composition and recycled-glass content [4,23].

Table 2. CFD data for inputs (Power, Pull, Cullet %) versus temperature outputs (T_batch, T_underglass and T_outlet).

Case Number	Pull (kg/s)	Power (kW)	Cullet %	T_Batch (°C)	T_Undeglass (°C)	T_Outlet (°C)
1	0.35	1607.1	27.8	282	1634	1602
2	0.25	900	12.8	220	1442	1418
3	0.2	1671.4	6.4	357	1560	1503
4	0.14	1221.4	2.1	1148	1743	1625
5	0.12	964.2	19.2	849	1605	1496
6	0.31	1414.2	4.2	229	1526	1432
7	0.16	1350	15	1146	1738	1670
8	0.29	1028.5	0	126	1206	1195
9	0.33	1157.1	21.4	132	1298	1287
10	0.27	1542.8	17.1	448	1756	1666
11	0.22	1285.7	30	514	1675	1589
12	0.12	900	0	591	1527	1442
13	0.12	900	30	770	1559	1462
14	0.12	1671	0	537	1480	1411
15	0.36	1671	30	314	1681	1581
16	0.36	900	30	161	1162	1157

A classical full-factorial design at three levels (min, mean, max) for the three inputs would require 27 high-fidelity CFD runs, which is computationally prohibitive at the current Glass-Furnace Model fidelity. Instead, a space-filling design was used, which provides superior coverage of the input space for surrogate fitting at a given sample budget. To preserve the spirit of a min/mean/max design, the boundary conditions of each input axis are explicitly represented: Cases 2, 5, 12, 13, 16 span the minimum-power and minimum-pull boundaries; Cases 3, 4, 14, 15 represent maximum-power conditions; Cases 11, 13, 15, 16 span the maximum-cullet boundary; and Case 8 represents zero-cullet, near-minimum-power conditions. This combined space-filling + corner-augmentation strategy ensures that the surrogate is bounded by physically validated extreme conditions while remaining tractable.

Rather than using raw field data directly, the CFD solutions are post-processed to extract scalar indices that summarize melting and transport behavior in a physically interpretable form. These include residence time (RT), sand dissolution index (SDI), mixing index (MI), and melting index (MEI). Data of SDI, MI, and MEI across the sixteen runs demonstrate the variability captured by the dataset and highlight distinct sensitivities of each metric to the imposed operating conditions. In particular, SDI and MEI exhibit wide dynamic ranges across the sampled input space, indicating strong dependence on thermal exposure and residence effects, while MI reflects changes in electrically driven and buoyancy-induced mixing patterns.

For each CFD case, the glass-process quantities of interest are post-processed from spatial fields generated by the high-fidelity furnace model. In the present work, the reported scalar outputs are defined using both domain-averaged and lower-percentile summaries of these fields. For a generic local KPI field $Y(x)$, where Y may represent RT, SDI, MI, or MEI, the domain-averaged value is defined as

$$Y_{\text{AVG}} = (1/V) \int_V Y(x) \, dV \quad (11)$$

and the 10th-percentile value is defined as

$$Y_{10\%} = P10[Y(x)] \quad (12)$$

where V is the total melt-domain volume of the simulated furnace and $P10$ denotes the 10th percentile of the spatial distribution of the local KPI field $Y(x)$ over the melt domain. Domain-averaged values (Equation (11)) characterize overall furnace performance, while 10th-percentile values (Equation (12)) capture the conservative lower-tail behavior associated with worst-case melting quality and minimum residence time within the melt. These two summaries are used to capture both overall performance and lower-tail behavior relevant to melting quality, mixing quality, and conservative operating assessment. In this study, residence time (RT) characterizes melt exposure to the thermal and hydrodynamic environment, sand dissolution index (SDI) serves as a proxy for melting completeness, mixing index (MI) reflects melt homogeneity, and melting index (MEI) provides an aggregated indicator of melting effectiveness.

The final ROM input dataset therefore consists of paired samples linking operating parameters (pull rate, power input, and cullet fraction) to CFD-derived performance metrics. This formulation reduces the original high-dimensional CFD field outputs to a smaller set of scalar performance indicators linked to three operating inputs, thereby making regression-based surrogate modeling tractable and interpretable. By explicitly defining the input runs and relying on physically motivated indices, the dataset provides a transparent and robust foundation for ROM training, validation, and sensitivity analysis.

3.2. Temperature ROM Generation

Linear reduced order models (ROMs) were constructed to relate key furnace temperature metrics to operating inputs consisting of pull rate, electrical power, and cullet fraction. Using the dataset described in the previous section, three scalar temperature quantities of interest were modeled: batch temperature (T_{batch}), under-glass temperature ($T_{\text{underglass}}$), and outlet temperature (T_{outlet}). For each quantity, the ROM was formulated as a main-effects linear regression of the form

$$T = \beta_0 + \beta_1 \text{Pull} + \beta_2 \text{Power} + \beta_3 \text{Cullet}, \quad (13)$$

where the coefficients β_i were estimated using ordinary least squares over the sixteen CFD operating points.

The resulting ROM equations are:

$$\begin{aligned} T_{\text{batch}} &= 935.18 - 3185.15 \text{ Pull} + 0.15595 \text{ Power} + 6.93 \text{ Cullet} \\ T_{\text{underglass}} &= 1266.63 - 1245.29 \text{ Pull} + 0.37728 \text{ Power} + 5.68 \text{ Cullet} \\ T_{\text{outlet}} &= 1200.93 - 954.75 \text{ Pull} + 0.33187 \text{ Power} + 4.94 \text{ Cullet} \end{aligned} \quad (14)$$

Here, power is expressed in kW and cullet in percent. These models provide a compact mapping from operating conditions to temperature response within the sampled design space.

The relatively high LOOCV error of the T_{batch} ROM ($\sim 245^\circ\text{C}$; $\sim 30\%$ MAPE) reflects the strong localized melting nonlinearity in the batch region, which a 3-input main-effects linear formulation cannot fully resolve. For this reason, T_{batch} is treated as a directional indicator rather than a quantitative output in the supervisory optimization framework. By contrast, the smoother responses of $T_{\text{underglass}}$ and T_{outlet} to the control inputs result in substantially lower LOOCV errors ($\sim 167^\circ\text{C}$ and $\sim 144^\circ\text{C}$, corresponding to $\sim 10\%$ and $\sim 9\%$ MAPE, respectively), supporting their use within the constrained supervisory

optimizer. Future work will adopt nonlinear surrogates (Gaussian-process, RBF, or neural ROMs) for T_{batch} as the CFD DOE is expanded.

The fitted coefficients yield physically interpretable first-order sensitivities. Power input exhibits a positive influence on all temperature metrics, with the strongest effect observed for under-glass and outlet temperatures, reflecting the direct role of Joule heating in raising melt and throat-region temperatures—a trend consistent with prior CFD analyses of Joule-heated and electrically boosted glass furnaces [1,10,14,24]. For example, the ROM predicts an increase of approximately 38 °C in under-glass temperature and 33 °C in outlet temperature per additional 100 kW of power. Pull rate enters with a negative coefficient for all three responses, indicating that increased throughput reduces thermal residence time and lowers temperatures throughout the furnace; this behavior is consistent with established glass-furnace heat-transfer analyses in which higher pull rate reduces the time available for thermal equilibration and melting completion [11,14]. The large magnitude of the pull coefficient in the batch temperature model highlights the strong coupling between pull rate and batch melting dynamics [11,12]. Cullet fraction contributes positively to all temperature metrics, consistent with reduced melting energy requirements and enhanced thermal efficiency associated with higher cullet content, as reported in prior glass-furnace energy-performance studies [4,23].

Model performance was evaluated using both in-sample goodness of fit and leave-one-out cross-validation due to the limited dataset size. The linear ROMs achieve coefficients of determination in the range $R^2 \approx 0.54\text{--}0.65$, indicating that they capture the dominant first-order trends across the operating envelope while remaining intentionally simple and interpretable. Although higher-order nonlinear effects are not explicitly represented, the linear ROMs are well suited for rapid sensitivity analysis, DOE interpretation, and preliminary optimization studies within the bounds of the sampled operating space. The fitted coefficients and validation metrics for each temperature ROM are summarized in Table 3, highlighting the relative influence of pull rate, power input, and cullet fraction on batch, under-glass, and outlet temperatures.

Table 3. Linear ROM coefficients for temperature metrics.

Output Temperature	Intercept (β_0)	Pull (β_1)	Power (β_2) [°C/kW]	Cullet (β_3) [°C/%]	R^2	Adj. R^2	F-Stat	p (Model)	LOOCV RMSE (°C)	LOOCV MAE (°C)
T_{batch}	935.18	−3185.15	0.15595	6.93	0.65	0.57	7.42	0.004	245	198
$T_{\text{underglass}}$	1266.63	−1245.29	0.37728	5.68	0.56	0.45	5.10	0.016	167	134
T_{outlet}	1200.93	−954.75	0.33187	4.94	0.54	0.43	4.69	0.021	144	116

Although the achieved R^2 values are modest, the linear ROMs capture the dominant first-order trends across the sampled operating envelope and are therefore useful as bounded, interpretable screening surrogates for sensitivity analysis and preliminary supervisory optimization. They are not intended as high-accuracy predictive replacements for the underlying CFD model.

The reported diagnostics indicate that all three temperature ROMs are statistically significant at the 5% level ($p < 0.05$), with adjusted R^2 values in the range 0.42–0.56 confirming that the main-effects formulation captures the dominant first-order trends across the sampled operating envelope. The coefficient-level standard errors and t-statistics confirm that power input is the most statistically significant predictor for $T_{\text{underglass}}$ and T_{outlet} , while pull rate dominates for T_{batch} . Residual-versus-fitted and Q–Q plots indicate no systematic bias and approximately normal residuals, supporting the validity of the ordinary-least-squares assumptions within the bounded DOE domain.

3.3. ROM Validation

The predictive capability of the linear reduced order models was assessed using leave-one-out cross-validation (LOOCV), which is well suited for small datasets such as the present study comprising sixteen CFD operating points. In the LOOCV procedure, a single CFD case is excluded from the training set, the ROM is refit using the remaining fifteen cases, and the excluded case is then predicted. This process is repeated for each operating point, providing an unbiased estimate of out of sample prediction error while maximizing the use of available high-fidelity data.

Figure 5 compares CFD-predicted and ROM-predicted temperatures for representative validation cases, including batch, under-glass, and outlet temperatures. The ROM captures the overall trends observed in the CFD results across all three temperature metrics, including the relative ordering of cases and the non-monotonic variation with operating conditions. For under-glass and outlet temperatures in particular, the ROM predictions closely follow the CFD values, indicating that the dominant first-order dependencies on pull rate, power input, and cullet fraction are well represented by the linear model (Figure 6).




Figure 5. Leave-one-out comparison of CFD-predicted and ROM-predicted temperature outputs (batch, under-glass, and outlet temperatures) across all sixteen DOE cases, plotted on a unified vertical axis of 0–1800 °C. Batch temperature ranges from approximately 126–1148 °C, under-glass temperature from approximately 1162–1756 °C, and outlet temperature from approximately 1157–1670 °C.

Quantitatively, the LOOCV root mean square errors are approximately 245 °C for batch temperature, 167 °C for under-glass temperature, and 144 °C for outlet temperature. The larger relative error observed for batch temperature is consistent with its stronger sensitivity to localized melting dynamics and nonlinear batch-transformation effects, which are not explicitly represented in a main-effects linear formulation. Prior modeling studies of Joule-heated and cold-cap melters have similarly noted that batch/cold-cap behavior and multiphase effects introduce additional uncertainty relative to smoother bulk-melt or outlet-temperature quantities [9,11,12]. In contrast, under-glass and outlet temperatures exhibit smoother, more monotonic responses to the operating inputs, resulting in improved ROM fidelity—consistent with these quantities being more directly governed by bulk Joule-heating and convective heat-transfer trends, which are amenable to first-order surrogate representation over a bounded operating envelope [1,10,14].

Overall, the LOOCV results demonstrate that the linear ROM provides a robust and physically consistent surrogate for rapid prediction and sensitivity analysis within the sampled operating envelope. While the model is not intended to replace high-fidelity CFD for detailed design analysis, it is sufficiently accurate for DOE interpretation, screening studies, and preliminary optimization. The observed validation trends further motivate the potential extension of the ROM framework to include interaction or piecewise linear terms where higher accuracy is required.




Figure 6. LOOCV predicted vs observed (parity) plot.

For the intended supervisory use case, the principal requirement is to capture directional sensitivities and relative tradeoffs among candidate operating conditions within the validated operating window, rather than to serve as a universally predictive stand-alone replacement for high-fidelity CFD. The reported LOOCV RMSE values (≈ 245 °C for T_{batch} , ≈ 167 °C for $T_{\text{underglass}}$, and ≈ 144 °C for T_{outlet} , corresponding to mean absolute percentage errors of approximately 30%, 10%, and 9%, respectively) and R^2 values of 0.54–0.65 reflect the deliberate choice of a low-complexity main-effects linear surrogate fitted to only sixteen high-fidelity CFD samples covering a wide operating envelope. These error levels are not intended to support standalone high-accuracy point prediction; rather, they are acceptable for the bounded supervisory decision-support role envisioned here, in which the ROM is used to identify directional sensitivities and feasibility regions, after which the recommended set-point is verified by a confirmatory CFD run prior to plant deployment. To explicitly absorb ROM uncertainty, the supervisory optimization in Section 3.5 enforces conservative constraint envelopes of ± 25 °C on $T_{\text{underglass}}$ and $\pm 5\%$ on RT, SDI, MI, and MEI—both of which exceed the corresponding LOOCV error magnitudes for the relevant ROM outputs. The optimization results should therefore be interpreted as advisory recommendations with an associated uncertainty band, rather than as deterministic control set-points.

The cross-validated margin of error of the temperature ROMs, expressed as LOOCV RMSE and MAPE, is approximately ± 245 °C (~30%) for T_batch, ± 167 °C (~10%) for T_underglass, and ± 144 °C (~9%) for T_outlet. The substantially higher error for T_batch reflects strong nonlinear local melting dynamics that are not fully captured by a 3-input main-effects linear model, and T_batch is therefore used only as a directional indicator rather than a quantitative output in the supervisory optimization.

We acknowledge that R^2 values in the range 0.54–0.65 would be unsuitable for a standalone high-accuracy predictive model. The linear ROMs developed here are not intended for this purpose. Their role is to provide a bounded, interpretable supervisory screening surrogate that identifies directional sensitivities and feasibility regions within a validated CFD operating window, with the recommended set-point subsequently confirmed by a single high-fidelity CFD run prior to plant deployment. Within this scope, the achieved R^2 and LOOCV error levels are adequate, as evidenced by the constraint envelopes (± 25 °C on T_underglass and $\pm 5\%$ on the glass indices) used in the supervisory optimization, which exceed the corresponding LOOCV error magnitudes.

3.4. Glass Indices ROM Generation

Linear reduced order models (ROMs) were developed to relate key glass quality and process indices to the principal operating variables of pull rate, electrical power input, and cullet fraction. The indices considered include residence time (RT), sand dissolution index (SDI), mixing index (MI), and melting index (MEI), each evaluated using both percentile-based (10%) and domain-averaged definitions. These quantities were extracted from high-fidelity CFD simulations and serve as compact descriptors of transport, melting completeness, and melt homogeneity in the electric glass melting furnace as given in Table 4.

Table 4. CFD data for case numbers versus glass melting indices (for input conditions for case numbers please refer Table 2).

Case Number	RT_10%	RT_AVG	SDI_10%	SDI_AVG	MI_10%	MI_AVG	MEI_10%	MEI_AVG
1	15.5	121.2	204.84	1742.7	15.8	122.6	23,946,300	2.01×10^8
2	9.5	53.0	12.3	96.4	6.1	36.7	3,137,220	22,649,800
3	10.9	63.4	80.3	668.2	8.5	53.4	8,655,080	64,120,600
4	15.7	94.9	153.7	1297.6	13.7	98.5	44,967,200	3.17×10^8
5	18.9	103.5	173.6	1401.2	11.5	80.6	21,185,800	1.46×10^8
6	22.5	312.6	137.5	2155.3	17.9	240.3	17,329,000	2.63×10^8
7	14.0	79.5	138.0	1089.2	11.3	76.4	46,786,400	2.98×10^8
8	8.9	44.6	0.3	2.4	3.9	17.5	293,936	1,982,840
9	7.5	38.9	1.1	8.4	4.1	20.7	695,988	4,663,380
10	8.7	46.6	96.6	646.1	9.4	57.8	25,117,900	1.58×10^8
11	10.3	56.0	105.85	770.9	9.2	57.9	18,135,200	1.2×10^8
12	18.89	110.71	85.53	761.0	9.7	70.1	12,184,400	92,050,000
13	19.086	111.78	135.7	1241.5	10.2	77.0	15,399,800	1.16×10^8
14	19.22	111.06	43.02	382.2	9.1	61.1	8,635,030	65,965,900
15	4.96	19.87	49.3709	263.5	5.3	17.2	8,818,300	44,069,300
16	9.13	37.45	0.18	1.0	3.2	12.1	170,096	975,529

For RT, SDI, and MI, the ROMs were formulated as main-effects linear regressions of the form

$$Y = \beta_0 + \beta_1 \text{ Pull} + \beta_2 \text{ Power} + \beta_3 \text{ Cullet}, \quad (15)$$

where Pull, Power (kW), and Cullet (%) define the input space. The melting index (MEI) spans several orders of magnitude across the sampled operating conditions; therefore, a log-

arithmetic formulation was employed to improve numerical conditioning and interpretability in Equation (14):

$$\log_{10}(\text{MEI}) = \beta_0 + \beta_1 \text{ Pull} + \beta_2 \text{ Power} + \beta_3 \text{ Cullet}, \quad (16)$$

All coefficients were estimated using ordinary least squares based on the sixteen CFD operating points.

The fitted linear ROMs for the glass indices are given in Equations (15)–(18):

Residence time

$$\begin{aligned} \text{RT}_{10\%} &= 21.17 - 29.71 \text{ Pull} + 1.88 \times 10^{-4} \text{ Power} - 0.067 \text{ Cullet} \\ \text{RT}_{\text{AVG}} &= 95.04 - 2.46 \text{ Pull} + 0.0174 \text{ Power} - 1.85 \text{ Cullet}. \end{aligned} \quad (17)$$

Sand dissolution index

$$\begin{aligned} \text{SDI}_{10\%} &= 57.23 - 384.09 \text{ Pull} + 0.0734 \text{ Power} + 1.90 \text{ Cullet} \\ \text{SDI}_{\text{AVG}} &= 539.85 - 2309.04 \text{ Pull} + 0.549 \text{ Power} + 6.11 \text{ Cullet} \end{aligned} \quad (18)$$

Mixing index

$$\begin{aligned} \text{MI}_{10\%} &= 6.71 - 14.58 \text{ Pull} + 0.00496 \text{ Power} - 0.011 \text{ Cullet}, \\ \text{MI}_{\text{AVG}} &= 42.82 - 23.25 \text{ Pull} + 0.0371 \text{ Power} - 1.00 \text{ Cullet}. \end{aligned} \quad (19)$$

Melting index (log ROM)

$$\begin{aligned} \log_{10}(\text{MEI}_{10\%}) &= 6.24 - 5.78 \text{ Pull} + 0.00138 \text{ Power} + 0.0166 \text{ Cullet}, \\ \log_{10}(\text{MEI}_{\text{AVG}}) &= 7.11 - 5.64 \text{ Pull} + 0.00139 \text{ Power} + 0.0128 \text{ Cullet}. \end{aligned} \quad (20)$$

The fitted coefficients reveal physically consistent first-order trends. Pull rate contributes negatively to all indices, reflecting reduced thermal exposure and residence time at higher throughput; this behavior is consistent with established glass-furnace modeling literature in which residence time and thermal history are key factors controlling melting completeness, dissolution, and downstream glass quality [1,11,14]. Power input exhibits a positive influence on SDI and MEI, consistent with enhanced melting and dissolution under increased Joule heating and stronger electrically driven circulation, as also reported in CFD analyses of electric-boosting effects on glass-quality indicators [10,24]. Cullet fraction also contributes positively to SDI and MEI, capturing the reduced melting burden associated with higher cullet usage and its known influence on furnace energy performance [4,23]. Mixing-related indices show weaker linear dependence on the inputs, indicating that melt homogenization is governed by coupled thermal and hydrodynamic effects—including three-dimensional buoyancy and electric-field-driven flow patterns—that are only partially captured by a main-effects linear formulation [9–11].

Overall, the glass index ROMs provide an interpretable and computationally efficient surrogate for sensitivity analysis, DOE interpretation, and preliminary operating optimization within the sampled design space. While higher-order or nonlinear ROM formulations may be required for control grade prediction, the present linear models establish a transparent baseline linking furnace operating conditions to glass quality indicators.

3.5. Optimization

The reduced order models for residence time (RT), sand dissolution index (SDI), mixing index (MI), and under-glass temperature were combined to formulate an extended supervisory level optimization problem aimed at reducing electrical energy consumption while simultaneously preserving glass quality and protecting furnace hardware. Electrical

power input was selected as the optimization objective, serving as a direct proxy for furnace energy consumption, while RT, SDI, MI, and under-glass temperature were constrained to remain within allowable operating envelopes defined relative to a baseline condition.

A baseline operating point was selected within the validated ROM domain, corresponding to a nominal furnace condition with pull rate of $0.24 \text{ kg}\cdot\text{s}^{-1}$, electrical power input of 1350 kW, and cullet fraction of 15%. The ROM-predicted values of RT, SDI, MI, and under-glass temperature at this point were treated as reference targets. To reflect industrial control practice, RT, SDI, and MI were constrained to remain within $\pm 5\%$ of their baseline values, representing acceptable variability in melting performance and melt homogeneity. The under-glass temperature was constrained within $\pm 25^\circ\text{C}$ of its baseline value, reflecting refractory protection limits and typical operational guard bands used to prevent excessive thermal loading of the furnace bottom.

The optimization problem was solved as a linear program with decision variables corresponding to pull rate, electrical power input, and cullet fraction. Additional feasibility constraints were imposed to ensure that all manipulated variables remained within the ROM training domain and that production throughput was protected by enforcing a lower bound on pull rate equal to the baseline value. Because the objective and constraints are linear in the manipulated variables and ROM outputs, the resulting optimization problem is a linear program and therefore convex. This structure allows rapid evaluation of candidate operating scenarios relative to repeated high-fidelity CFD analysis.

The optimized operating point reduced electrical power input from 1350 kW to approximately 1256 kW, corresponding to an energy reduction of about 7%. The pull rate remained fixed at the baseline value of $0.24 \text{ kg}\cdot\text{s}^{-1}$, satisfying the throughput protection constraint, while the cullet fraction increased modestly to approximately 16.8%. This adjustment reflects the beneficial influence of cullet on melting and mixing performance, allowing partial compensation for reduced power input.

At the optimized operating point, all quality and thermal constraints were satisfied. Residence time decreased slightly relative to the baseline but remained well within its allowable envelope. Both SDI and MI reached their lower constraint bounds, indicating that melting completeness and mixing effectiveness were the active quality constraints limiting further power reduction. The under-glass temperature remained within its prescribed $\pm 25^\circ\text{C}$ band, confirming that the optimized solution does not compromise refractory protection despite the reduced power input. The simultaneous activity of SDI, MI, and under-glass temperature constraints highlights the coupled nature of energy efficiency, melt quality, and furnace durability in electric glass melting operations.

Overall, the results demonstrate that meaningful energy savings can be achieved through ROM-based supervisory optimization while maintaining both glass-quality and equipment-protection criteria. By explicitly incorporating under-glass temperature alongside melting and mixing indices, the optimization framework more closely reflects industrial operating priorities and provides a robust digital-twin tool for evaluating energy–quality–lifetime tradeoffs in electric glass melters. This outcome is consistent with prior studies showing that model-based operating-point adjustment can reduce glass-furnace energy consumption without compromising furnace performance when appropriate quality and operational constraints are enforced [4,23], and with previous CFD-based observations that energy efficiency, melt homogeneity, melting completeness, and refractory protection are strongly coupled in electrically heated or electrically boosted glass furnaces [1,11,24]. The use of the ROM as a supervisory advisory layer is also consistent with broader digital-twin and surrogate-modeling paradigms, in which high-fidelity simulations are used offline to construct low-latency models for rapid scenario evaluation and operational decision support [18,20,21].

The reported ~7% power reduction (1350 → 1256 kW) is bounded by the cross-validated ROM error and should be interpreted as a first-order estimate of the achievable energy savings within the validated operating envelope. Propagating the LOOCV temperature errors (± 167 °C for $T_{\text{underglass}}$) through the active constraint set yields a power-saving uncertainty band of approximately ± 1.5 percentage points around the nominal optimum. Accordingly, the framework is positioned as a supervisory advisory tool whose recommendations are to be confirmed via a single high-fidelity CFD run before implementation on the physical furnace.

4. Conclusions

This work has presented a simulation-informed supervisory optimization framework for an electric cold-top glass melter that integrates high-fidelity CFD, interpretable linear reduced-order models (ROMs), and a convex linear-program optimizer subject to physics-based quality and refractory-protection constraints. The framework was developed to address a specific gap in the literature: existing CFD-based studies provide deep physical insight but remain offline, while data-driven energy models lack explicit physics-based quality enforcement.

Sixteen CFD operating points spanning realistic industrial ranges of electrical power (900–1671 kW), pull rate (0.12–0.36 kg/s), and cullet fraction (0–30%) were used to fit linear ROMs for three furnace temperatures and four glass-process indicators (RT, SDI, MI, MEI). Leave-one-out cross-validation yielded LOOCV RMSE values of approximately 245 °C, 167 °C, and 144 °C for T_{batch} , $T_{\text{underglass}}$, and T_{outlet} , respectively. While these errors preclude use of the ROMs as standalone high-accuracy predictors, they are acceptable for the bounded supervisory decision-support role envisioned here, in which recommended set-points are subsequently confirmed by a single high-fidelity CFD run.

For a representative 1350 kW baseline operating point, the supervisory optimizer recommended a reduction in electrical power to 1256 kW (~7% reduction) while maintaining throughput and satisfying all glass-quality and refractory constraints. The active constraint set—SDI, MI, and $T_{\text{underglass}}$ —confirms the coupled nature of energy efficiency, glass quality, and furnace durability in electric melting. The observed optimization behavior is consistent with prior glass-furnace CFD and energy-optimization studies, which show that power input, residence behavior, melting completeness, and thermal-protection constraints must be considered simultaneously when pursuing energy reduction in industrial glass melting [4,11,23,24].

Limitations of the present study include the bounded operating envelope, the use of steady-state CFD, and the simplicity of the linear ROM formulation, which limits accuracy in regions of strong nonlinearity (notably T_{batch}). Future work will (i) expand the CFD DOE to support nonlinear surrogates (Gaussian-process, piecewise-linear, neural ROMs), (ii) extend the framework to transient operation, (iii) integrate plant-data assimilation, and (iv) couple the supervisory layer with model-predictive control for closed-loop deployment.

Author Contributions: Conceptualization, A.M., R.A. and B.A.P.; Methodology, A.M., A.L., S.D.C.P. and B.A.P.; Software, A.M. and A.L.; Validation, A.M.; Formal analysis, A.M.; Investigation, A.M., A.L. and B.A.P.; Resources, A.M., A.L. and B.A.P.; Data curation, A.M., A.L., S.D.C.P. and B.A.P.; Writing—original draft, A.M. and B.A.P.; Writing—review and editing, A.M., A.L., S.D.C.P., R.A. and B.A.P.; Visualization, A.M., A.L. and B.A.P.; Supervision, A.M., R.A. and B.A.P.; Project administration, A.M., R.A. and B.A.P.; Funding acquisition, A.M. and B.A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors acknowledge Glass Service for their constant support with Glass-Furnace Model software and collaboration with Owens Corning Science and technology. The authors would also like to extend their gratitude to Kevin Smith and Amol Vaidya for their encouragement for this work and external engagement.

Conflicts of Interest: All authors were employed by the Owens Corning Science & Technology Centre. The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The authors declare that this study received support from the Glass Service. The funder had the following involvement with the study: Glass-Furnace Model software.

Abbreviations and Nomenclature

The following nomenclatures and abbreviations are used in this manuscript:

u [m/s]	melt velocity vector
ρ [kg/m ³]	glass density
μ [Pa·s]	dynamic viscosity (temperature-dependent)
p [Pa]	pressure
g [m/s ²]	gravitational acceleration
F_{Joule} [N/m ³]	volumetric Lorentz body force (negligible vs. buoyancy)
c_p [J/kg·K]	specific heat capacity
k [W/m·K]	thermal conductivity (effective, includes radiative contribution)
T [°C]	temperature
Q_{Joule} [W/m ³]	volumetric Joule heat source
σ [S/m]	electrical conductivity (temperature-dependent)
E [V/m]	electric field vector
φ [V]	electric potential
V [m ³]	melt-domain volume
$P_{10}[\cdot]$	10th percentile of the spatial distribution
CFD	Computational Fluid Dynamics
ROM	Reduced Order Models
GFM	Glass-Furnace Model
DOE	Design Of Experiment

References

- Cravero, C.; Marsano, D.; Milanese, G. CFD simulation of the increased electric boost effects on the glass melting process in a real glass furnace to support decarbonization of glass industry. *J. Phys. Conf. Ser.* **2024**, *2893*, 012033. [[CrossRef](#)]
- Parrot, A.; Warshaw, L. *Industry 4.0 and the Digital Twin: Manufacturing Meets Its Match*; A Deloitte Series on Industry, 4; Deloitte University Press: New York, NY, USA, 2017; pp. 1–17.
- Guillen, D.P.; Ferkl, P.; Pokorný, R.; Hall, M.; Hrma, P.; Kruger, A.A. Numerical modeling of Joule heated ceramic melter. *Mater. Lett.* **2024**, *362*, 136201. [[CrossRef](#)]
- Miserocchi, L.; Franco, A.; Testi, D. Status and prospects of energy efficiency in the glass industry: Measuring, assessing and improving energy performance. *Energy Convers. Manag. X* **2024**, *24*, 100720. [[CrossRef](#)]
- Xi, L.; Yin, D.-W.; Park, J.S. Special Issue “CFD Modeling of Complex Chemical Processes: Multiscale and Multiphysics Challenges”. *Processes* **2021**, *9*, 775. [[CrossRef](#)]
- Choudhary, M.K. The 2024 AGC-Michael Cable Memorial Lecture: Progress and challenges in decarbonization of industrial glass melting. *Glass Technol. Eur. J. Glass Sci. Technol. Part A* **2025**, *66*, 65–78. [[CrossRef](#)]
- Müller, V. An all-electric melter for a variety of glass types. *Glass Int.* **2020**, *43*, 16–19.
- Stanek, J.W. *Electric Melting of Glass*; Elsevier Scientific: New York, NY, USA, 1977.
- Chang, S.L.; Zhou, C.Q.; Golchert, B.; Petrick, M. Simulation of Multi-Phase Glass-Melt Flows in a Glass Melter. In *ASME International Mechanical Engineering Congress and Exposition*; American Society of Mechanical Engineers: New York, NY, USA, 2001; Volume 35616, pp. 73–80.

10. Matsuno, S.; Iso, Y.; Uchida, H.; Oono, I.; Fukui, T.; Ooba, T. CFD modeling coupled with electric field analysis for joule-heated glass melters. *J. Power Energy Syst.* **2008**, *2*, 447–455. [\[CrossRef\]](#)
11. Choudhary, M.K.; Venuturumilli, R.; Hyre, M.R. Mathematical modeling of flow and heat transfer phenomena in glass melting, delivery, and forming processes. *Int. J. Appl. Glass Sci.* **2010**, *1*, 188–214. [\[CrossRef\]](#)
12. Agarwal, V.; Guillen, D.P. *Incorporating Cold Cap Behavior in a Joule-Heated Waste Glass Melter Model* (No. INL/EXT-13-29794); Idaho National Laboratory (INL): Idaho Falls, ID, USA, 2013.
13. Ferkl, P.; Hirma, P.; Abboud, A.; Guillen, D.P.; Vernerová, M.; Kloužek, J.; Hall, M.; Kruger, A.A.; Pokorný, R. Conversion degree and heat transfer in the cold cap and their effect on glass production rate in an electric melter. *Int. J. Appl. Glass Sci.* **2023**, *14*, 318–329. [\[CrossRef\]](#)
14. Prescott, P.J.; Purnode, B. Simulation in Glass Processes. In *Encyclopedia of Glass Science, Technology, History, and Culture*; Richet, P., Conradt, R., Takada, A., Dyon, J., Eds.; John Wiley & Sons: Hoboken, NJ, USA, 2021. [\[CrossRef\]](#)
15. Abbassi, A.; Khoshmanesh, K. Numerical simulation and experimental analysis of an industrial glass melting furnace. *Appl. Therm. Eng.* **2008**, *28*, 450–459. [\[CrossRef\]](#)
16. Benner, P.; Gugercin, S.; Willcox, K. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev.* **2015**, *57*, 483–531. [\[CrossRef\]](#)
17. Antoulas, A.C. Approximation of Large-Scale Dynamical Systems: An Overview. *IFAC Proc. Vol.* **2004**, *37*, 19–28. [\[CrossRef\]](#)
18. Aversano, G.; Ferrarotti, M.; Parente, A. Digital twin of a combustion furnace operating in flameless conditions: Reduced-order model development from CFD simulations. *Proc. Combust. Inst.* **2021**, *38*, 5373–5381. [\[CrossRef\]](#)
19. Ren, X.; Sadeghirad, A.; Pham, D.C.; Lua, J. Machine learning based digital twin simulation and prediction of structural component's fatigue life. In *Proceedings of the Vertical Flight Society 75th Annual Forum*; American Helicopter Society: Fairfax, VA, USA, 2019; pp. 1045–1059.
20. Tao, F.; Liu, A.; Nee, A.Y.C. Digital twin in industry: State-of-the-art. *IEEE Trans. Ind. Inform.* **2019**, *15*, 2405–2415. [\[CrossRef\]](#)
21. Molinaro, R.; Singh, J.-S.; Catsoulis, S.; Narayanan, C.; Lakehal, D. Embedding data analytics and CFD into the digital twin concept. *Comput. Fluids* **2021**, *214*, 104759. [\[CrossRef\]](#)
22. Zybala, R.; Wycislik, J.; Golak, S.; Cieplinski, P.; Sak, T.; Madej, P. Fast numerical optimization of electrode geometry in a two-electrode electric resistance furnace using a surrogate criterion derived exclusively from an electromagnetic submodel. *Appl. Sci.* **2024**, *14*, 10957. [\[CrossRef\]](#)
23. Kodak, O.; Kaya, M.B.; Sadeghi-Khaneghah, F.; Dumankaya, E.; Alanat, G.Y.; Kiliç, L.; Arzan, N.; Konukman, A.E.S. Improvement of energy performance of glass furnaces using modelling and optimization techniques. *Processes* **2025**, *13*, 3739. [\[CrossRef\]](#)
24. Raič, J.; Wachter, P.; Hödl, P.; Demuth, M.; Gaber, C.; Gerhardt, H.; Prieler, R.; Hochenauer, C. CFD simulation aided glass quality and energy efficiency analysis of an oxy-fuel glass melting furnace with electric boosting. *Energy Convers. Manag. X* **2022**, *15*, 100252. [\[CrossRef\]](#)
25. Atzori, D.; Tiozzo, S.; Vellini, M.; Gambini, M.; Mazzoni, S. Industrial technologies for CO₂ reduction applicable to glass furnaces. *Thermo* **2023**, *3*, 682–710. [\[CrossRef\]](#)
26. Mishra, A.; Lohse, A.; Arora, R.; Roux, E.; Fabris, R.; Chakravarti, S.; Tian, K.; Hernandez, V. Simulations, Reduced Models & Alternative Fuels to Improve Productivity and Sustainability in Glass Fiber Manufacturing, 17th International Seminar On Furnace Design–Operation & Process Simulation. Available online: <https://www.gsl.cz/seminar-2025/> (accessed on 24 May 2026).
27. Mishra, A.; Korba, D.; Zhao, J.; Li, L. Heat and Mass Transfer Model for a Counter-Flow Moving Packed-Bed Oxidation Reactor/Heat Exchanger. *ASME J. Sol. Energy Eng.* **2024**, *146*, 051006. [\[CrossRef\]](#)
28. Mishra, A.; Singh, P.; Li, L. Heat transfer model for moving packed-bed particle-to-sCO₂ heat exchangers integrated with metal foams. *Appl. Therm. Eng.* **2024**, *239*, 122062. [\[CrossRef\]](#)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.