Proc. of the International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME 2025)
16-19 October 2025, Zanzibar, Tanzania

2025 5th International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME) | 979-8-3315-3556-8/25/$31.00 ©2025 IEEE | DOI: 10.1109/ICECCME64568.2025.11277713

Frequency-Constrained UC via Second-Order ODEs with ML Surrogates

Azeez O. Olasoji
Department of Electrical Engineering
University of Cape Town
Cape Town, South Africa
olsaze001@myuct.ac.za

David T.O Oyedokun
Department of Electrical Engineering
University of Cape Town
Cape Town, South Africa
david.oyedokun@uct.ac.za

Mohammad Rajabdorri
IIT, ETSI-ICAI
Universidad Pontificia Comillas
Madrid, Spain
mrajabdorri@comillas.edu

Juan Esteban Sierra-Aguillar
Department of Electrical Engineering
Universidad de Antioquia
Medellin, Colombia
juane.sierra@udea.edu.co

Chukwuemeka E. Okafor
Department of Electrical Engineering
University of Cape Town
Cape Town, South Africa
okfchu002@myuct.ac.za

Mkhutazi Mditshwa
Department of Electrical Engineering
University of Cape Town
Cape Town, South Africa
mdtmkh001@myuct.ac.za

Komla A. Folly
Department of Electrical Engineering
University of Cape Town
Cape Town, South Africa
komla.folly@uct.ac.za

***Abstract*—This paper proposes a fast and interpretable framework for frequency-constrained unit commitment (FCUC) using machine-learning (ML) surrogates. Unlike prior studies that rely on proprietary system-frequency-response tools, we use an open-source second-order differential equation (SODE) model to simulate generator outages, generating a 117 000-scenario dataset. Three linear classifiers —Logistic Regression (LR), Linear Discriminant Analysis (LDA), and a log-loss Stochastic-Gradient-Descent model (SGD-Log)—are trained on this dataset and inserted as linear constraints in a mixed-integer UC model to enforce frequency adequacy. The surrogates detect $\geq$ 98.8 % of unsafe operating points ($\leq$ 32 false negatives) while training in only 0.1 s to 2.3 s. Applied to a spring-week case study on the La Palma island system, they uphold a strict −3 Hz nadir limit, achieve average nadirs of −1.23 Hz to −1.31 Hz, compared with −1.41 Hz in the unconstrained base case, and raise weekly cost by no more than 1.3 % above that baseline. Also, in comparison with a first-order differential equation (FODE) model, the SODE–ML formulations solve 40–390× faster (12.2 s to 101.4 s vs. 4767 s) without linearisation assumptions. These results demonstrate that reproducible, SODE-labelled surrogates enable secure and computationally efficient FCUC for low-inertia grids.**

Keywords*—*Frequency-constrained unit commitment (FCUC), machine learning (ML), frequency nadir (FN), second-order differential equation (SODE), mixed-integer linear programming (MILP), renewable energy sources (RES).

I. Introduction

The penetration of renewable energy sources (RES) in power systems is steadily increasing, driven by the need to mitigate carbon emissions from fossil fuel-based generators. However, the displacement of synchronous generators by integrating RES into power grids presents significant challenges due to their inherent volatility and the uncertainties associated with forecasting their output. Additionally, RES are typically decoupled from the power system, contributing no inertia, thereby amplifying frequency volatility [1]. This issue is particularly pronounced in small-scale systems, such as islands, where low inertia exacerbates frequency volatility and heightens the risk of instability during contingencies [1], [2]. This necessitates advanced frequency stability constraints in unit commitment (UC), as traditional N-1 security criteria become insufficient [2]. This proliferation has led to accelerated rates of change of frequency (RoCoF) and steeper frequency nadirs (FN's) after contingency events. Accurate nadir assessments are essential within unit-commitment (UC) models to ensure system security in low-inertia conditions.

Most UC studies approximate the frequency-nadir constraint derived from the first-order differential equation (FODE;classical swing equation with uniform governor time-constant) by employing various linearization techniques—such as separable programming and piecewise-linear techniques—to preserve MILP tractability [1]. These linearisation approaches simplify turbine-governor and damping dynamics; thus, the possibility of overestimating or underestimating FN requirements in the frequency-constrained unit commitment (FCUC;UC enhanced with RoCoF, steady-state and nadir limits) is very high [2].

In contrast, the second-order differential equation (SODE) formulation proposed in this study inherently captures both inertia and damping dynamics. However, SODE's nonlinear structure has historically prevented its direct embedding in UC models. To address this gap, we propose an SODE-ML surrogate that trains interpretable linear classifiers on a proprietary dataset of 117,000 generator contingency scenarios—catering to variations in inertia, droop settings, and reserve levels—to emulate the SODE solution within a fully linear framework suitable for direct GAMS integration. We validate our approach on the La Palma island power system, a real-time low-inertia grid in the Canary Islands, Spain, to demonstrate practical applicability and performance in realistic UC case studies.

This paper makes three key contributions:

- Development of linear SODE-ML surrogates using Logistic Regression (LR), Linear Discriminant Analysis (LDA), and Stochastic-Gradient-Descent classifier with log-loss (SGD-Log) models trained on a comprehensive generator outage dataset.
- Evaluation of surrogate performance via classification metrics (accuracy, precision, recall, F1) and training cost, with demonstration of direct coefficient export to GAMS.
- Benchmark UC case study on a low-inertia test system comparing objective cost, reserve allocation, and solver runtime between the SODE-ML surrogates and the FODE-based linearization baseline.

The remainder of this paper is organized as follows. Section II reviews existing literature and approaches employed for linearising frequency-nadir constraints. Section III details the underlying approaches used in this study. Section IV presents classification outcomes and UC case-study benchmarks. In section V, we discuss the results and their implications. Finally, Section VI concludes with insights and future research directions.

II. Literature Review

Frequency-constrained Unit Commitment (FCUC) formulations traditionally leverage simplified frequency dynamic models derived from the first-order differential equation (FODE), commonly termed the swing equation [3], [4]. These approaches typically linearize frequency-nadir (FN) constraints using techniques such as separable programming (SP), piecewise-linear (PWL) approximations, or fixed-profile reserve ramps [1], [5]. Although computationally attractive, such linearization strategies inherently simplify turbine-governor dynamics and damping effects, leading to substantial inaccuracies—often manifested as FN overestimations—which subsequently compromise system reliability and operational efficiency.

To overcome the limitations of these analytical methods, recent studies have turned toward data-driven and machine learning (ML) approaches. In [2] demonstrated significant computational efficiency gains by training logistic regression (LR) and support vector machine (SVM) classifiers using simplified system frequency response (SFR) models and proprietary simulation tools, thereby impacting the reproducibility of their approach.

Deep neural network (DNN)-based frameworks have also been proposed [6], providing improved accuracy but introducing considerable computational overhead due to thousands of auxiliary variables from nonlinear activations. Furthermore, methods using multivariate optimal classification trees (OCTs) have shown promising classification accuracy, albeit at the cost of scalability issues arising from exponential computational complexity with larger datasets [7].

Unlike prior ML studies reliant on simplified SFR models or proprietary tools, our framework trains classifiers on data generated directly from second-order differential equation (SODE) models, fully capturing turbine-governor dynamics absent in earlier works. This is achieved using Python's open-source SciPy ODE solvers, which eliminate dependencies on external simulations while preserving fidelity. The resulting SODE-derived frequency nadir estimates are integrated into linear ML classifiers such as LR and SVM, yielding linear constraints suitable for MILP formulations, thus maintaining both dynamic accuracy and computational tractability similar to [2]. This approach ensures scalable, reproducible dataset generation critical for real-world deployment in low-inertia systems.

III. Methodology

A. Classical UC Formulation

The classical Unit Commitment (UC) problem seeks the optimal scheduling of generating units to minimize operational costs while satisfying system constraints. The mathematical formulation is typically represented as a mixed-integer linear programming (MILP) problem:

$$\min \sum_{t \in T} \sum_{i \in I} \left(\mathrm{suc}(x_{t,i}) + \mathrm{gc}(g_{t,i}) \right) \tag{1}$$

Subject to: generator min. and max capacity, demand, ramp up/down, minimum up/down time, transmission, and binary logic constraints (not shown for brevity); refer to [2], [8] for a detailed breakdown .

Here, $\mathrm{suc}(x_{t,i})$ and $\mathrm{gc}(g_{t,i})$ denote the startup and generation costs, $x_{t,i}$ is the binary operational status of unit i at time t, $g_{t,i}$ is its power output, $r_{t,i}$ is the reserve, and D_t is system demand at time t.

B. Frequency Constraint Modelling

1) First-Order Differential Equation (FODE) Approach: The frequency dynamics following a generator outage are modelled using the swing equation:

$$P_a = P_m - P_e \tag{2}$$

$$\frac{2H}{f_0} \frac{d\Delta f(t)}{dt} + D\Delta f(t) = P_a - \sum_i \min(r_i(t), p_i) \tag{3}$$

Post-contingency system inertia:

$$H_{sys} = \sum_{i \in I_{on}} H_i M_i x_{t,i} \tag{4}$$

a) RoCoF Constraint:

$$H_{sys} \geq \frac{P_l f_0}{2\Delta f_{cr}^{RoCoF}} \tag{5}$$

b) QSS Constraint:

$$\sum_{i \in I} r_i \geq P_l - DP_{D,t} \Delta f_{cr}^{QSS} \tag{6}$$

c) Frequency Nadir Constraint: The reserve response during a contingency is modelled as a piecewise function:

$$R_{eff}(t) = \begin{cases} \frac{R_{tot}}{T_g} t & \text{if } t \leq T_g \\ R_{tot} & \text{if } t > T_g \end{cases} \tag{7}$$

Total reserve after outage:

$$R_{tot} = \sum_{i \in I_{on}} r_i \tag{8}$$

The nonlinear frequency nadir constraint becomes:

$$HR_{tot} - \frac{f_0 T_g (\Delta P_l^{\max})^2}{4\Delta f_{cr}^{nadir}} + \frac{DP_{D,t} T_g (\Delta P_l^{\max}) f_0}{4} \geq 0 \tag{9}$$

d) Symbol Definitions::

- f_0: nominal system frequency (Hz)
- Δf_{cr}^{nadir}: critical nadir frequency deviation (Hz)
- ΔP: power imbalance after contingency (MW)
- T_g: governor deadband delay (s)
- $P_{D,t}$: system demand at time t
- D: system damping coefficient (MW/Hz)
- H: aggregated system inertia (s)
- R_{tot}: total spinning reserve available (MW)

2) Second-Order Differential Equation (SODE) Approach: To better capture short-term frequency dynamics, we propose a second-order differential equation (SODE) model, which accounts for turbine-governor interactions:

$$a_{2,i} \frac{d^2 r_i}{dt^2} + a_{1,i} \frac{dr_i}{dt} + r_i(t) = -k_i x_i \left(\Delta f(t) + b_i \frac{d\Delta f(t)}{dt} \right) \tag{10}$$

This model reflects mechanical response characteristics and damping, with parameters derived from generator-specific control systems.

- $r_i(t)$: frequency response from generator i
- $k_i = 1/R_i$: droop gain (inverse of droop constant)
- x_i: binary commitment status
- $a_{1,i}, a_{2,i}, b_i$: parameters of each generating unit i, can be deduced from more accurate models or field tests (In the case of La-Palma, field tests)

a) Model Interpretation: The SODE model captures oscillatory behaviour and inertia-limited frequency support more accurately than FODE, as it is not predicated on the assumption of linear activation of reserves. However, the equation is nonlinear and not directly MILP-compatible; we simulate its response under thousands of contingency scenarios to generate labelled frequency nadir outcomes. These labels are later used to train interpretable ML classifiers that emulate the nadir boundary. These dataset used for these simulations are generated using a synthetic dataset constructed via grid search over operational variables (inertia, droop, reserves, and contingency sizes), enabling scalable and reproducible training data without reliance on proprietary simulation tools discussed next. Fig. 1 shows a schematic of the SODE-based system frequency response model.

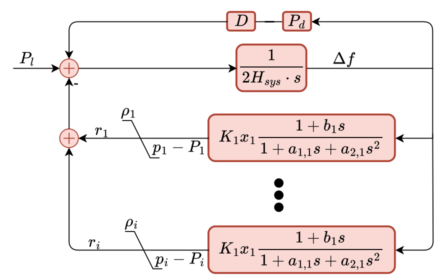

Fig. 1. Schematic of second-order system frequency response model.

C. Dataset Generation via Grid Search

To simulate realistic post-contingency frequency responses with the SODE model, we first construct a synthetic dataset using an exhaustive grid search over critical UC operating points. This dataset enables supervised training of ML classifiers for nadir adequacy.

We adopt a structured grid search (GS) approach, generating candidate operating points by varying generator on/off status, power outputs, and reserve allocations, similar to [2]. Each combination must satisfy key UC constraints: (i) total generation must lie within acceptable demand limits, (ii) sufficient spinning reserve must be maintained, (iii) system inertia before and after any single generator outage must exceed analytically derived thresholds (e.g., from RoCoF or nadir limits), and (iv) total production cost must fall within a target percentile to reflect operational plausibility. These feasibility checks ensure the grid search yields a high-quality subset of dispatches for SODE-based frequency simulation and binary labeling based on nadir violations.For networks larger than La Palma, future work will study other algorithms with different sampling schemes such as Latin Hypercube reducing the computational complexity.

Each data instance generated contains features $x \in X$ that represent post-fault system conditions: post-contingency inertia H_{post}, weighted droop gain k_l, amount of lost generation p_l, and available reserves R_l. These features were selected based on their physical relevance and prior use in frequency stability studies.

The frequency nadir for each scenario is then computed by solving the SODE using Python's `SciPy` library. A binary label $y \in Y$ is assigned based on whether the computed nadir violates a predefined threshold. While GS has exponential computational complexity$\mathcal{O}(n^k)$, it guarantees exhaustive coverage of the operating space—essential for robust learning and downstream UC optimisation. This labelled dataset serves as the training input for the ML models introduced in the next section, which approximate the nadir constraint for integration into UC.

D. ML Surrogates for SODE Frequency-Nadir Constraints

Building on the labelled dataset produced from SODE simulations, we train a set of linear ML models that emulate the nadir security boundary under varying UC operating conditions. These surrogates provide linear decision boundaries that can be directly embedded into the MILP problem without auxiliary variables or non-convex formulations.

In this study, we implement three linear ML models: LR with $c = \infty$, LDA, and SGD-Log algorithm. These models are chosen for their interpretability, compatibility with MILP formulations, and ability to approximate the nadir adequacy boundary using minimal coefficients. Prior work has demonstrated the efficacy of linear ML methods like support vector machines (SVM) and LR in similar FCUC contexts [2].

All models are trained on the dataset generated via grid search and labeled using SODE-based frequency nadir values. The feature vector $x \in X$ includes post-contingency inertia H_{post}, weighted droop gain k_l, lost generation p_l, and available reserves R_l. Labels $y \in Y$ are binary, indicating whether the nadir violates a defined threshold. The three classifiers are trained with their canonical objectives; logistic cross-entropy loss with an effectively negligible L2 penalty for LR ($c = \infty$), the Gaussianian maximum-likelihood criterion with shared covariance for LDA, and logistic loss with an L2 penalty optimised by stochastic gradient descent for the SGD-Log model. Each model outputs a linear constraint of the form:

$$\theta^T x + \theta_0 \geq 0 \tag{11}$$

which is exported into GAMS and enforced as a frequency-security constraint within the UC formulation.

To validate model performance, we apply 70/30 train–test splits and evaluate each ML model using classification accuracy and F1-score. The learned surrogates replace complex dynamic constraints, utilized when deducing FN constraints using FODE, significantly reducing computational overhead while preserving dynamic fidelity.

IV. Case Study and Results

A. System Description

The case study is based on the La Palma island power system, a real isolated grid located in Spain's Canary Islands. The system includes eleven predominantly diesel-based generators, with an installed capacity of approximately 117.7 MW. Roughly 6% of this capacity originates from wind generation, which contributes around 10% of the annual energy demand. The demand dataset used in this study is drawn from publicly available historical data for 2018, where the total yearly demand is estimated at 277.8 GWh (corresponding to an average hourly load of 31.7 MWh).

Unlike prior studies that evaluate seasonal variation, we focus here on a representative week in a single season to establish the proof of concept. Future extensions will expand this evaluation to all four seasons. The under-frequency security criterion for this system is a minimum nadir threshold of 3.0 Hz, which is enforced during the SODE simulations to distinguish secure versus insecure operating points.

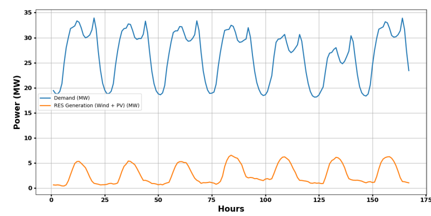

Fig. 2. Aggregate demand and RES generation profiles over a representative spring week in La Palma.

B. Dataset Summary and Feature Insights

Table I presents the time and scale involved in generating and labeling the dataset used to train ML models. After applying feasibility filters, 20,000 operating points were retained from the Grid Search (GS) space. Each point was then evaluated (single generator outages) as contingency events using the SODE model, resulting in a labelled dataset of 115,000 secure/insecure instances for ML training.

TABLE I. Dataset generation and labeling summary

Stage	Size	Runtime (s)
Grid Search Operating Points	20,000	2789
Labeled Dataset via SODE	117,000	2042

Table II shows the Pearson correlation coefficients (ρ) between the input features and SODE-derived frequency nadir; while reserve availability has traditionally been used as a proxy for frequency security, our analysis confirms that inertia and droop gain are significantly more correlated with nadir deviation. The lost generation p_{loss} exhibits the strongest (negative) association. These findings align with recent FCUC literature [2], supporting the inclusion of physically grounded features such as H_{sys} and K_{on} in ML-based surrogate modelling.

C. ML Performance Evaluation

Table III reports seven key metrics for the three linear ML models trained on the SODE-labelled dataset: *training runtime*, *recall*, number of *false negatives* (FN), *precision*, *F1-score*, *area under the ROC curve* (AUROC) and overall *accuracy*. Recall, FN, and AUROC are most critical for FCUC, because they quantify how reliably a surrogate detects unsafe

TABLE II. Pearson correlation between input features and frequency nadir

Feature	Correlation
Lost generation (p_{loss})	-0.9514
Post-contingency inertia (H_{sys})	$+0.6417$
Aggregate droop (K_{on})	$+0.6410$
Available reserves (R_{l})	$+0.5854$

TABLE III. Performance of ML models

Model	Runtime (s)	Recall (%)	FN	Prec. (%)	F1 (%)	AUROC (%)	Acc. (%)
Logistic (C = ∞)	0.262	98.96	28	17.4	29.5	98.4	89.2
LDA	0.115	98.81	32	9.4	17.2	95.8	78.2
SGD-Log	2.335	99.00	27	16.2	27.8	98.0	88.2

operating points (very shallow frequency nadirs). Precision reflects the converse cost risk—how often the model flags a schedule as unsafe when it is actually safe—while F1 balances those two margins. Accuracy offers a coarse check, and wall-clock runtime determines whether a surrogate can be retrained on the fly as system conditions evolve.

Balancing speed and accuracy: All three surrogates train in under 2.3 s, yet achieve *recall* $\geq$ 98.8 %—crucial because missed nadir violations (*false negatives*) would compromise security. The LR ($C = \infty$) model is therefore adopted as default: it cuts training time by a factor of ≈9 compared with SGD-Log while preserving all but one additional contingency (28 vs 27 FN) and attaining the highest AUROC (98.4 %). Precision naturally drops to 9–17 % because the threshold is biased toward eliminating FN; the result is a slight increase in cost rather than a violation of FN or initiating an under frequency event . In other words, the 100× scheduling speed-up is obtained *at no material loss of frequency security or cost.*

D. Computational Setup and Cross-Study Benchmark

All optimisation runs were executed in GAMS 24.9.2 using the CPLEX solver on a 4-core Intel® Core™ i5-6500 (3.2 GHz) desktop with 16 GB RAM.[1]

Per-hour runtime: To neutralise the different horizon lengths used in recent FCUC studies, we report solver time per simulated hour (s h^{-1}):

This study's (SODE-LR, spring week, 24 h horizon); 38 s/24 h ≈ 1.6 s h^{-1}. In [2] (LR surrogate, 168 h horizon, seasonally averaged); 302 s/168 h ≈ 1.8 s h^{-1}. While the analytical FODE MILP (baseline in [2], [8])- 25 506 s/168 h ≈ 152 s h^{-1}.

Thus, even after horizon normalisation, the proposed surrogate delivers ≈ *12 % faster per-hour performance* than the previous ML state-of-the-art and remains more than two orders of magnitude faster than the analytical FODE MILP.[2]

[1] Processor and memory are stated to enable per-interval normalisation; In [2] hardware characteristics where not specified.

[2] In [2] runtimes are weekly averages over winter, spring, summer and autumn. Extending the analysis to the remaining seasons is left for future work.

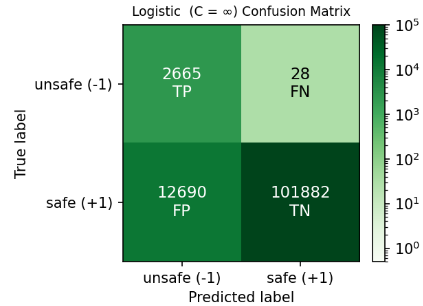

(a) LR ($C = \infty$)

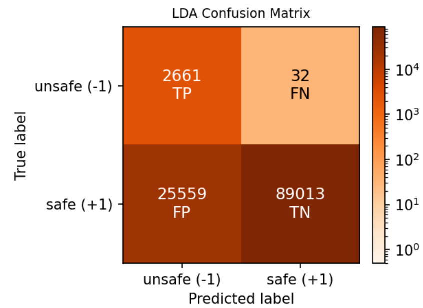

(b) LDA

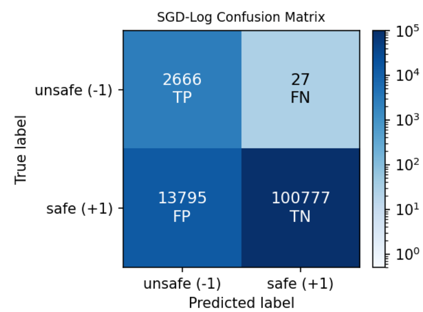

(c) SGD-Log

Fig. 3. Confusion matrix for selected ML models.

Such sub-second evaluation per UC interval enables; real-time rescheduling after contingencies on ordinary desktop hardware, a task formerly mostly reserved for high performance computing clusters.

Furthermore, the analytical FODE MILP needs 70 min for a 24-h La Palma schedule; the proposed SODE-LR surrogate cuts this to ≈40 s. An operator can hypothetically recompute a secure commitment within a real time cycle after a line or generator outage, without having waiting for the next dispatch horizon. Also, Planners can test topology upgrades, inertia-sharing schemes or reserve-policy tweaks in real time on a standard 4-core desktop, avoiding costly high-performance computing (HPC) resources and shortening study lead-times.

V. FCUC Simulation Results and Discussion

Table IV presents the average weekly UC cost for the La Palma system in spring, comparing the SODE-ML surrogates against the FODE-based baseline. The ML-based UC formulations were solved for each surrogate model using GAMS with CPLEX, and runtimes reflect complete MILP solution time.

TABLE IV. Average weekly cost, runtime and frequency nadir of each FCUC formulation

Method	Runtime (s)	Startup Cost (k€)	Op. Cost (k€)	Total Cost (k€)	Avg Fnadir (Hz)
Base Case (BC)	29.65	6.16	109.84	116.00	-1.40699
LR ($c = \infty$)	101.40	6.04	111.07	117.11	-1.31048
LDA	12.20	5.98	110.97	116.95	-1.22890
SGD-Log	75.67	6.09	111.41	117.51	-1.21544
FODE*	4767.27	6.67	110.23	116.90	-1.31547

* FODE baseline re-computed in this study using the analytical model of [2], [8] with a −3.5 Hz nadir limit.

FCUC discussion: Three insights emerge from Table IV. *(i) Cost–security trade-off*- The unconstrained *Base Case* attains the lowest weekly cost 116.00 k€ but also the deepest average nadir (−1.407 Hz). Imposing a tight −3 Hz frequency limit through ML surrogates raises cost by only 0.82 % (LDA), 0.96 % (LR), and 1.30 % (SGD-Log) while improving nadir to −1.229 Hz to −1.310 Hz. *(ii) Computational advantage*- All three surrogates solve in 12.2 s to 101.4 s, slashing runtime by two to three orders of magnitude compared with the analytical FODE baseline (4767 s). *(iii) Baseline comparison*- The FODE case (marked with "*") was evaluated under the looser −3.5 Hz criterion and still incurs 0.7 k€ more in startup cost and a $\geq$ 4700 s runtime penalty relative to the fastest ML surrogate (LDA). The study deliberately adopts the stricter −3 Hz setting to emulate a "no-surprises" policy aimed at averting apocalypse-like events—such as the recent Iberian-peninsula cascade, where voltage instability initiated widespread outages. Although voltage, not frequency, triggered that blackout, a stricter frequency nadir provides extra headroom should exigent frequency excursions arise. The SODE-trained surrogates thus deliver enhanced resilience at negligible additional cost and with a lower computational burden than legacy analytical models.

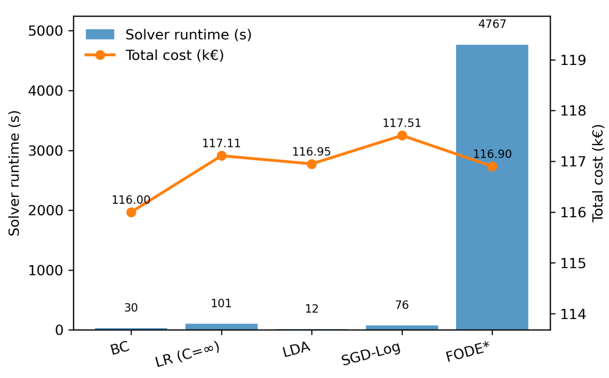

Fig. 4. Weekly UC cost and solve time across methods.

As stated prior we follow similar methodology in [2], where post-contingency nadir values are used to label dispatch adequacy. However, while Rajabdorri's work used a Simulink-based SFR simulator, we implement an equivalent second-order dynamic model in open-source Python using `SciPy`. The resulting nadir values, validated in post-UC simulations (Table IV), demonstrate functional parity. This reinforces the reproducibility of the FCUC framework and demonstrates that open, reproducible implementations can preserve both modelling fidelity and practical utility.

VI. Conclusion

This paper presents FCUC framework that replaces proprietary swing-equation simplifications with an open-source SODE model and three linear ML surrogates (LR, LDA, SGD-Log). Trained on 117k SODE-simulated contingencies, the surrogates detect $\geq$ 98.8 % of unsafe operating points with no more than 32 missed cases and require 0.1 –2.3 s to fit. Embedded as linear constraints, they enforce a strict −3 Hz nadir limit on the La Palma island system, raising weekly costs by at most 1.3 % relative to the unconstrained base case yet solving the FCUC MILP **40–390× faster** than an analytical first-order (FODE) formulation. Validated on real island power system, the ML-surrogate is immediately transferable to other isolated systems.

Because the entire schema—SODE simulation, GS and ML surrogates—relies only on Python/SciPy and GAMS(which can be replaced with pyomo), the approach is fully reproducible and avoids proprietary frequency-response software.

This study evaluates a single representative week under deterministic assumptions. Future work will extend to seasonal variability, stochastic UC, and integration of non-conventional frequency services such as BESS and synthetic inertia. Future studies would also test the validity of the approach on larger test systems. Overall, SODE-labelled, MILP-compatible surrogates provide a practical, computationally efficient path to secure real-time scheduling in low-inertia grids.

References

[1] A. O. Olasoji, D. T. O. Oyedokun, S. O. Omogoye, and C. Thron. Review of frequency response strategies in renewable-dominated power system grids: Market adaptations and unit commitment formulation. *Scientific African*, 26:e02357, 2024.

[2] M. Rajabdorri, B. Kazemtabrizi, M. Troffaes, L. Sigrist, and E. Lobato. Inclusion of frequency nadir constraint in the unit commitment problem of small power systems using machine learning. *Sustainable Energy, Grids and Networks*, 36:101161, 2023.

[3] L. Badesa, F. Teng, and G. Strbac. Simultaneous scheduling of multiple frequency services in stochastic unit commitment. *IEEE Transactions on Power Systems*, 34(5):3858–3868, 2019.

[4] C. J. Ferrandon-Cervantes, B. Kazemtabrizi, and M. C. M. Troffaes. Inclusion of frequency stability constraints in unit commitment using separable programming. *Electric Power Systems Research*, 203:107669, 2022.

[5] H. Ahmadi and H. Ghasemi. Security-constrained unit commitment with linearized system frequency limit constraints. *IEEE Transactions on Power Systems*, 29(4):1536–1545, 2014.

[6] Y. Zhang et al. Encoding frequency constraints in preventive unit commitment using deep learning with region-of-interest active sampling. *IEEE Transactions on Power Systems*, 37(3):1942–1955, 2022.

[7] D. T. Lagos and N. D. Hatziargyriou. Data-driven frequency dynamic unit commitment for island systems with high res penetration. *IEEE Transactions on Power Systems*, 36(5):4699–4711, 2021.

[8] A. O. Olasoji, D. T. O. Oyedokun, M. Rajabdorri, and J. E. Sierra-Aguillar. Investigating inclusion of linear sensitivity factors in analytical frequency constrained unit commitment formulation. In *2025 33rd Southern African Universities Power Engineering Conference (SAUPEC)*, pages 1–6, Pretoria, South Africa, 2025.