

Article

Not peer-reviewed version

Forecast-Guided KAN-Adaptive FS-MPC for Resilient Power Conversion in Grid-Forming BESS Inverters

[Shang-En Tsai](#)* and Wei-Cheng Sun

Posted Date: 26 March 2026

doi: 10.20944/preprints202603.2138.v1

Keywords: grid-forming inverter; BESS; finite-set MPC; kolmogorov-arnold networks; adaptive weighting; voltage sag; resilience; forecast-guided control; microgrids; artificial intelligence; robust control; transient stability



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Forecast-Guided KAN-Adaptive FS-MPC for Resilient Power Conversion in Grid-Forming BESS Inverters

Shang-En Tsai * and Wei-Cheng Sun

Department of Computer Science and Information Engineering, Chang Jung Christian University, Tainan City 711, Taiwan

* Correspondence: sean@mail.cjcu.edu.tw; Tel.: +886-6-2785123 (ext. 6154)

Abstract

Grid-forming (GFM) battery energy storage system (BESS) inverters are becoming a cornerstone of resilient microgrids, where severe voltage sags and abrupt operating shifts can challenge both voltage regulation and controller stability. Finite-set model predictive control (FS-MPC) offers fast transient response and multi-objective coordination, yet conventional designs rely on static cost-function weights that are typically tuned offline and may become suboptimal under disturbance-driven regime changes. This paper proposes a forecast-guided KAN-adaptive FS-MPC framework that (i) formulates the inner-loop predictive control in the stationary $\alpha\beta$ frame, thereby avoiding PLL dependency and mitigating loss-of-lock risk under extreme sags, and (ii) introduces an Operating Stress Index (OSI) that fuses load forecasts with reserve-margin or percent-operating-reserve signals to quantify grid vulnerability and trigger resilience-oriented control adaptation. A lightweight Kolmogorov-Arnold Network (KAN), parameterized by learnable B-spline edge functions, is embedded as an online weight governor to update key FS-MPC weighting factors in real time, dynamically balancing voltage tracking and switching effort. Experimental validation under high-frequency microgrid scenarios shows that, under a 50% symmetrical voltage sag, the proposed controller reduces the worst-case voltage deviation from 0.45 p.u. to 0.16 p.u. (64.4%) and shortens the recovery time from 35 ms to 8 ms (77.1%) compared with static-weight FS-MPC. In the islanding-like transition case, the proposed method restores the PCC voltage within 18 ms, whereas the static baseline fails to recover within 100 ms. Moreover, the deployed KAN governor requires only 6.2 μ s per inference on a 200 MHz DSP, supporting real-time embedded implementation. These results demonstrate that forecast-guided adaptive weighting improves transient resilience and power quality while maintaining DSP-feasible computational complexity.

Keywords: grid-forming inverter; BESS; finite-set MPC; kolmogorov-arnold networks; adaptive weighting; voltage sag; resilience; forecast-guided control; microgrids; artificial intelligence; robust control; transient stability

1. Introduction

Modern power systems are increasingly dominated by inverter-interfaced renewable generation and are exposed to more frequent extreme operating conditions. In microgrids, grid-forming (GFM) battery energy storage system (BESS) inverters are not merely power interfaces; they are active voltage-establishing assets that determine whether voltage and frequency can be sustained and restored during severe disturbances such as deep voltage sags, abrupt load steps, and islanding-like transitions [1–3]. Under these stressed conditions, the control layer must simultaneously achieve fast voltage support, bounded current response, and practical switching effort, which makes resilient inner-loop control a central problem for converter-dominated power systems.

Finite-set model predictive control (FS-MPC) is attractive for power converters because it directly evaluates admissible switching states and can coordinate multiple objectives with fast transient response [4,5]. However, the practical effectiveness of FS-MPC depends critically on the weighting factors in the stage cost. In conventional implementations, these weights are tuned offline and then fixed, often by trial-and-error or scenario-specific heuristics [6,7]. Such static tuning is inherently brittle under disturbance-driven regime changes, because the optimal trade-off between voltage regulation and switching stress varies significantly across nominal operation, deep faults, and post-disturbance recovery. Prior learning-assisted approaches, including neural-network-based weight selection and reinforcement-learning-based adaptation, have shown promise [8,9], but their black-box nature raises deployment concerns for safety-critical grid infrastructure.

To address robustness under extreme disturbances, this work adopts a PLL-free predictive-control formulation in the stationary $\alpha\beta$ frame. This choice is not merely a coordinate transformation convenience. During asymmetrical faults and distorted voltage conditions, synchronous dq-frame control depends on a phase-locked loop (PLL) whose angle estimate can be corrupted by negative-sequence components, producing double-frequency oscillations and potentially loss-of-lock. By contrast, the stationary $\alpha\beta$ frame remains directly tied to the physical inverter voltage vectors, allowing FS-MPC to evaluate candidate switching states without relying on phase tracking. This makes the control law structurally better suited to severe-sag operation and weak-grid transients.

In parallel, this paper introduces a dual-timescale resilience architecture that couples slow supervisory awareness with fast converter control. Specifically, a forecasting-based layer computes an Operating Stress Index (OSI) from forecasted load and reserve-margin or percent-operating-reserve signals, thereby translating system-level vulnerability into supervisory context for converter control [2,3,10]. Unlike direct forecast-driven actuation, the OSI layer does not issue switching commands; instead, it conditions the allowable range and prior of FS-MPC weighting factors. Because the forecast update interval is orders of magnitude slower than the converter sampling period, the OSI signal can be treated as a quasi-static supervisory parameter within the fast predictive loop, which provides a principled basis for the proposed multi-timescale interface.

To realize online adaptation with improved interpretability and embedded feasibility, we employ a lightweight Kolmogorov-Arnold Network (KAN) as a weight governor. KAN replaces conventional node-wise fixed activations with learnable edge-wise spline functions, yielding a parameter-efficient nonlinear mapping that is structurally more auditable than a standard multilayer perceptron (MLP) [11]. Recent power-system research has further shown that physics-informed KAN architectures can model system dynamics more accurately than MLP-based PINNs while using fewer parameters, supporting the broader suitability of KANs for power-domain nonlinear learning [12]. Recent work has begun to explore KAN for adaptive control of power converters [13], while resilience-oriented ARIMA-KAN forecasting has been used to improve vulnerability-aware demand prediction [10]. Building on these emerging directions, the present work integrates forecast-guided resilience supervision, KAN-based adaptive weighting, and PLL-free $\alpha\beta$ -frame FS-MPC into a unified framework for resilient GFM BESS voltage control.

Contributions:

- **Forecast-guided resilience supervision:** an Operating Stress Index (OSI) is defined from load forecasts and reserve-margin or percent-operating-reserve signals to quantify vulnerability and trigger resilience-oriented supervisory modes.
- **KAN-adaptive FS-MPC:** a lightweight spline-based KAN governor updates FS-MPC weighting factors online using fast electrical features and optional OSI conditioning, enabling disturbance-aware adaptation with bounded outputs.
- **PLL-free $\alpha\beta$ implementation for severe faults:** the predictive controller is formulated in the stationary $\alpha\beta$ frame to preserve a well-defined control representation under deep sags and distorted conditions without PLL dependency.

- Resilience-aware converter evaluation: event-window metrics such as worst-case deviation, recovery time, and degradation area are adapted to converter voltage regulation and validated through high-frequency disturbance scenarios and HIL-oriented timing analysis.

2. Related Work

This section reviews grid-forming control and fault ride-through considerations, FS-MPC foundations and weighting-factor design, learning-assisted adaptive weighting, and resilience-oriented forecasting metrics that motivate forecast-guided converter control.

2.1. Grid-Forming Inverter Control Under Faults and Sags

Prior studies have investigated droop-based and virtual synchronous machine (VSM)-like grid-forming strategies for microgrids and weak grids [1,14]. These methods provide valuable system-level synchronization and power-sharing behavior, yet severe voltage sags and faulted conditions expose a more demanding inner-loop problem: the controller must preserve voltage support and bounded current response despite abrupt waveform distortion and rapidly changing operating constraints. Sequence-based and current-limiting fault ride-through strategies have therefore emphasized the need for disturbance-aware and protection-conscious control under deep sags [15]. Recent reviews have further highlighted that, in grid-forming inverters, current-limiting actions do not merely protect semiconductor hardware but can strongly shape converter behavior during and after disturbances, with direct implications for voltage support, transient stability, protection, and fault recovery [16].

A key issue under asymmetrical faults is the behavior of synchronous-reference-frame control. In dq-based implementations, the controller depends on a PLL to estimate the grid angle. When negative-sequence voltage components are present, the PLL is excited by a second-harmonic disturbance in the synchronous frame, which can introduce oscillatory angle estimation, degraded decoupling, and, in severe cases, loss-of-lock. Once phase estimation becomes unreliable, the underlying dq control representation itself is weakened, especially during fast transients in converter-dominated systems.

This motivates the use of the stationary $\alpha\beta$ frame for predictive voltage control. In the $\alpha\beta$ frame, the controller remains directly connected to the inverter's finite switching vectors and does not require explicit phase tracking to evaluate candidate actions. For FS-MPC, this is particularly attractive because the admissible switching set is naturally represented as stationary voltage vectors. As a result, the predictive law remains physically meaningful even under deep sags, waveform distortion, and weak-grid conditions. In this paper, the PLL-free $\alpha\beta$ formulation is adopted precisely for this resilience-oriented reason: it avoids a synchronization bottleneck and preserves direct finite-vector evaluation during extreme disturbances.

2.2. Finite-Set MPC for Power Converters and the Weighting Problem

FS-MPC has been widely adopted in power electronics because it directly handles switching actions, incorporates converter constraints in a transparent way, and offers fast transient response for multi-objective control tasks [4,5]. In typical converter applications, the stage cost combines terms associated with voltage or current tracking, switching effort, and, in some cases, protection-related penalties. The challenge is that these terms are not commensurate and must therefore be balanced by weighting factors.

The literature has proposed practical guidelines, normalization strategies, and structured design procedures for weighting-factor selection [6,7]. Nevertheless, most implementations still rely on fixed or weakly scheduled weights, which are tuned offline for a limited operating envelope. In addition, the practical performance of FS-MPC remains sensitive to model accuracy, parameter mismatch, and implementation constraints such as sampling-frequency mismatch, all of which can degrade prediction quality and switching-state selection [17,18]. Under disturbance-driven regime changes,

however, the desired trade-off is inherently nonstationary: during fault onset the controller may need to emphasize voltage support and protection urgency, whereas during recovery it may gradually re-balance toward switching-effort moderation and power-quality restoration. This gap between fixed offline tuning and real-time operating variability is one of the main limitations of conventional FS-MPC and directly motivates adaptive weight governance, as also reflected in recent learning-assisted FCS-MPC studies that address uncertainty, disturbance rejection, and the burden of manual weighting-factor design [19].

2.3. Learning-Assisted Adaptive Control: MLP/RL vs. Interpretable Lightweight Learners

Learning-assisted control has emerged as a promising route to improve adaptability in predictive power-converter control. Earlier work has explored neural-network-based weighting-factor selection [8,20], while more recent studies have investigated reinforcement-learning-based dynamic weighting policies [9]. These approaches demonstrate that data-driven mappings can outperform purely static tuning when the operating condition changes rapidly. However, their use in safety-critical electric-power infrastructure remains controversial. The main concern is not only performance, but also the inability of deeply black-box models to provide clear operational boundaries, auditable decision logic, and predictable behavior outside the most common training region.

This concern is especially acute for primary or inner-loop control in grid-supporting converters. In such settings, operators and reviewers often expect more than empirical improvement; they also expect boundedness, interpretability, and a credible path to engineering audit. Large MLPs and reinforcement-learning agents may adapt effectively, yet their internal representations are difficult to inspect, and abrupt nonlinear thresholds can be smoothed or hidden inside highly distributed parameters. This can make it difficult to certify how the controller will react in rare but high-impact disturbance regimes.

Kolmogorov-Arnold Networks (KANs) offer a different trade-off. Instead of relying on fixed node activations, KANs use learnable edge-wise univariate spline functions, which provide a more structured and inspectable nonlinear mapping [11]. Recent work has begun to show that KANs can support adaptive control design for power-electronics converters with strong parameter efficiency and improved interpretability [13]. For the present problem, this structure is particularly attractive because the mapping from disturbance features to FS-MPC weights is expected to exhibit localized threshold-like behavior. Edge-wise spline functions can capture such localized nonlinear transitions more directly than standard MLP activations, while also allowing engineers to inspect learned adaptation curves and impose explicit bounds. Accordingly, this paper adopts a lightweight KAN governor as an auditable and embedded-feasible alternative to black-box adaptive learners.

2.4. Learning-Assisted Adaptive Control: MLP/RL vs. Interpretable Lightweight Learners

Resilience-oriented power-system research emphasizes that disturbance-window behavior is often more operationally important than average steady-state performance [2,3]. In other words, the key question is not only how well a controller performs on average, but how severely performance degrades during an extreme event, how long recovery takes, and how much cumulative degradation is incurred before normal operation is restored. This perspective motivates event-window metrics such as worst-case deviation, recovery time, and degradation area, which are more informative for resilience assessment than average error alone.

At the supervisory level, forecasting can enhance system resilience by identifying vulnerability before a disturbance propagates to the converter layer. In particular, recent resilience-oriented forecasting studies have shown that combining interpretable seasonal modeling with nonlinear correction is effective for capturing both recurring temporal structure and extreme residual behavior [10]. In this context, the dual-stream seasonal ARIMA-KAN framework is especially relevant, since it separates slow seasonal memory from nonlinear stress amplification and is therefore well suited for generating supervisory indicators from forecasted load and reserve-related data. Based on the

above review, the overall architecture of the proposed forecast-guided KAN-adaptive FS-MPC framework is summarized in Figure 1, which links the slow-timescale forecasting and supervisory layers with the fast inner-loop predictive control of the GFM BESS inverter.

Motivated by this perspective, the present paper defines an Operating Stress Index (OSI) that fuses forecasted load level and reserve adequacy into a compact vulnerability signal for converter supervision. OSI is not intended to replace inner-loop control; rather, it provides a slow-timescale context variable that biases the adaptive-weight search space toward more resilience-oriented control behavior under stressed regimes. Because the forecast update interval is much larger than the converter sampling period, the proposed architecture follows a practical time-scale separation principle: within each fast FS-MPC interval, OSI is treated as quasi-static, while the inner loop remains responsible for real-time switching decisions. This dual-timescale interpretation links resilience forecasting and converter control in a theoretically coherent and operationally meaningful way.

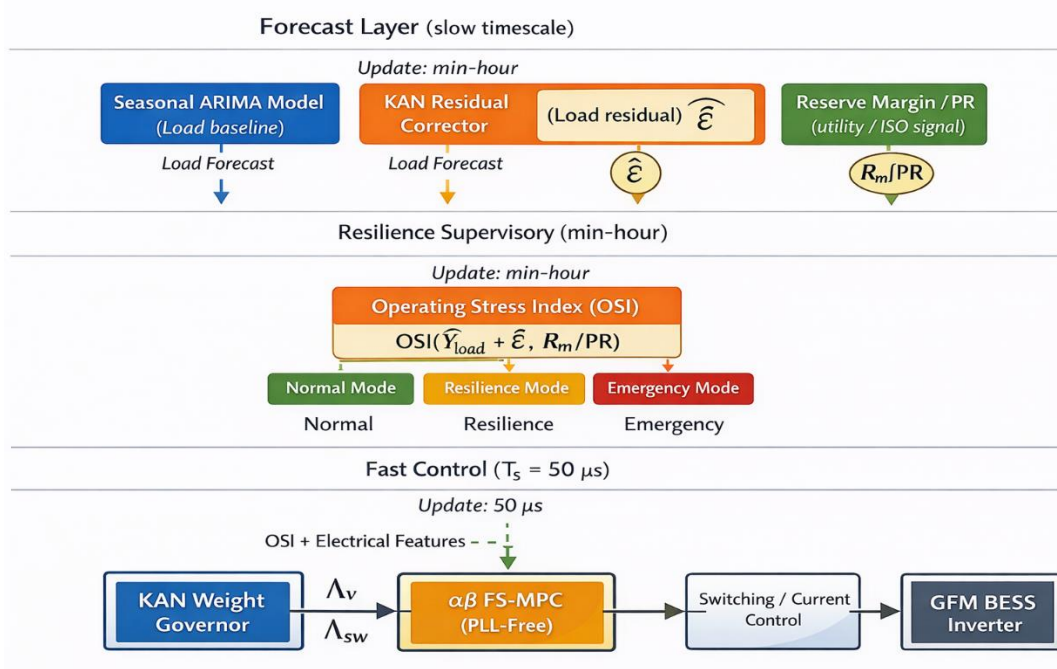


Figure 1. Overall architecture of the proposed forecast-guided KAN-adaptive FS-MPC framework. The forecasting layer produces an Operating Stress Index (OSI) from load and reserve-related information, the supervisory layer maps OSI into operating modes, and the fast control layer uses OSI together with electrical features to govern adaptive weights for PLL-free $\alpha\beta$ -frame FS-MPC in the GFM BESS inverter.

3. System Model and Control Problem Formulation

3.1. Converter Topology and PLL-Free Stationary Alpha-Beta Modeling

We consider a two-level three-phase GFM BESS inverter interfaced to a microgrid through an LC filter. Unlike synchronous dq-frame control that requires a phase-locked loop (PLL) for angle tracking, we formulate the predictive control in the stationary alpha-beta frame. This PLL-free formulation is physically aligned with finite voltage vector actuation and remains well-defined even when the grid voltage experiences deep sags or severe distortion, thereby mitigating loss-of-lock risk under extreme conditions. As illustrated in Figure 2, the controller directly uses the measured stationary-frame electrical variables $v_{c,\alpha\beta}$, $i_{L,\alpha\beta}$, and $i_{o,\alpha\beta}$ for prediction and switching-state evaluation without relying on PLL-based angle estimation.

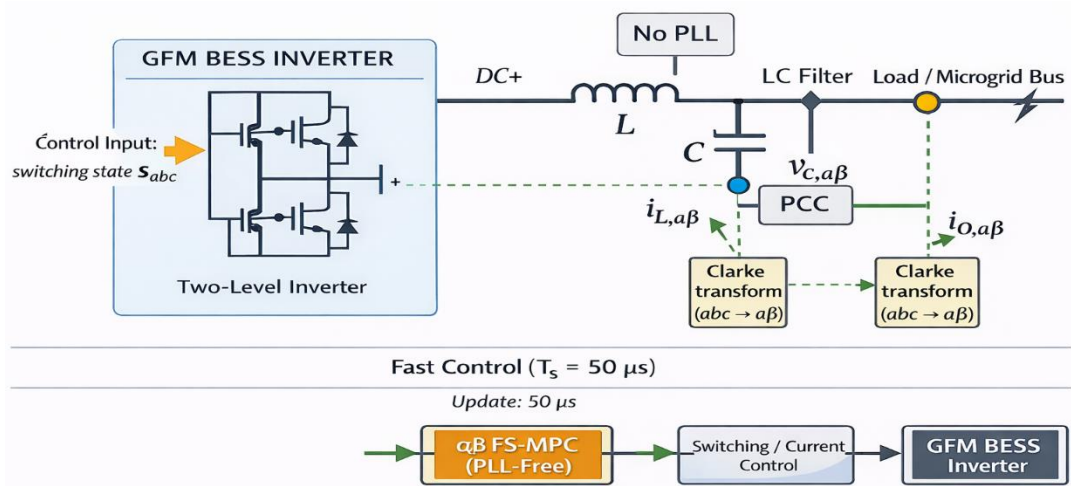


Figure 2. Plant model and measured signals of the PLL-free grid-forming BESS inverter in the stationary $\alpha\beta$ frame. A two-level voltage source inverter interfaces the microgrid bus through an LC output filter. The FS-MPC enumerates the finite switching set s_{abc} and uses the measured capacitor/PCC voltage $v_{c,\alpha\beta}$, inductor current $i_{L,\alpha\beta}$, and output/load current $i_{o,\alpha\beta}$ for one-step prediction and cost evaluation.

The principal converter and control parameters used throughout the predictive model and real-time evaluation are summarized in Table 1. Let $v_{inv,\alpha\beta}$ denote the inverter output voltage vector determined by the switching state, $v_{c,\alpha\beta}$ the filter-capacitor (PCC) voltage vector, $i_{L,\alpha\beta}$ the inductor current vector, and $i_{o,\alpha\beta}$ the output/load current vector. The continuous-time LC-filter dynamics can be expressed in a standard state-space form. For digital implementation, we discretize the model with a forward Euler approximation of sampling time T_s .

Table 1. System and control parameters.

Parameter	Symbol	Value	Unit
DC-link voltage	V_{dc}	750	V
Filter inductance	L	2.5	mH
Filter capacitance	C	20.0	μF
Equivalent resistance	R	0.1	Ω
Nominal PCC voltage (rms)	V_{nom}	380	V
Rated power	S_{rated}	10	kVA
Current limit	I_{max}	30	A
Sampling time	T_s	50	μs

Continuous-time state-space model (alpha-beta):

$$\dot{x}(t) = A x(t) + B u(t) + E d(t) \quad (1)$$

$$x(t) = [i_{L,\alpha\beta}(t); v_{c,\alpha\beta}(t)] \quad (2)$$

$$u(t) = v_{inv,\alpha\beta}(t) \quad (3)$$

Discrete-time one-step prediction (forward Euler):

$$x[k+1] = x[k] + T_s(A x[k] + B u[k] + E d[k]) \quad (4)$$

3.2. Finite Control Set and Candidate Voltage Vectors

FS-MPC enumerates all admissible switching states at each sampling instant. For a two-level inverter, the finite set contains eight distinct voltage vectors, including two zero vectors. The finite switching states and their corresponding stationary $\alpha\beta$ voltage vectors are summarized in Table 2. For each candidate, the controller predicts the next-step states and selects the switching action that minimizes the cost function.

Table 2. Finite control set for a two-level inverter and corresponding stationary alpha-beta voltage vectors.

Index	(S_a, S_b, S_c)	Vector type	$v_{inv,\alpha}$	$v_{inv,\beta}$	Notes
0	(0,0,0)	zero	0	0	Freewheeling
1	(1,0,0)	active	$\frac{2}{3}V_{dc}$	0	Active state
2	(1,1,0)	active	$\frac{1}{3}V_{dc}$	$\frac{\sqrt{3}}{3}V_{dc}$	Active state
3	(0,1,0)	active	$-\frac{1}{3}V_{dc}$	$\frac{\sqrt{3}}{3}V_{dc}$	Active state
4	(0,1,1)	active	$-\frac{2}{3}V_{dc}$	0	Active state
5	(0,0,1)	active	$\frac{1}{3}V_{dc}$	$-\frac{\sqrt{3}}{3}V_{dc}$	Active state
6	(1,0,1)	active	$\frac{1}{3}V_{dc}$	$-\frac{\sqrt{3}}{3}V_{dc}$	Active state
7	(1,1,1)	zero	0	0	Freewheeling

3.3. Conventional FS-MPC Stage Cost and the Static-Weight Limitation

In grid-forming operation, the controller prioritizes PCC voltage regulation while controlling switching effort to reduce losses and thermal stress. A typical static-weight stage cost is:

$$J(u) = \lambda_v \|v_{\alpha\beta}^*[k+1] - v_{c,\alpha\beta}^{(u)}[k+1]\|_2^2 + \lambda_{sw} \Delta s^{(u)}[k] \quad (5)$$

In conventional practice, λ_v and λ_{sw} are tuned by trial-and-error and then fixed. Under deep sags and abrupt regime shifts, the optimal trade-off between voltage support and switching stress changes rapidly, making static weights brittle and motivating online adaptive weight governance. [6,7].

4. Forecast-Guided Resilience Supervisory Layer (Operating Stress Index)

4.1. Motivation: Tail-Risk Awareness for Resilience Operation

Short-term load forecasting is a core input to operating-reserve scheduling and dispatch planning. For resilience, average forecast accuracy is insufficient because rare extreme events can dominate operational risk. Therefore, we use forecasting outputs to derive a supervisory vulnerability signal that conditions converter control posture during stressed regimes.

4.2. Dual-Stream Seasonal ARIMA–KAN Forecasting as the Supervisory Input

We reuse the dual-stream seasonal ARIMA–KAN framework that decomposes demand into an interpretable seasonal memory component (ARIMA) and a nonlinear residual correction component (KAN) [10]. The supervisory layer consumes the forecasted load and reserve-related signals rather than directly controlling the converter at the fast timescale.

4.3. Operating Stress Index (OSI): Combining Load Forecast and Reserve Margin

We define an Operating Stress Index (OSI) to quantify grid vulnerability and trigger resilience-oriented control adaptation. OSI fuses (i) the forecasted load level and (ii) the forecasted reserve margin (or percent operating reserve, PR), thereby capturing both demand pressure and dispatchable headroom. Higher OSI indicates a more vulnerable operating condition.

Normalized stress components:

$$S_L(t+h) = \sigma((\hat{Y}_{t+h} - \mu_Y)/\sigma_Y) \quad (6)$$

$$S_R(t+h) = \sigma((PR_{t+h} - \mu_{PR})/\sigma_{PR}) \quad (7)$$

OSI definition:

$$OSI(t+h) = \text{clip}(w_L S_L(t+h) + w_R(1 - S_R(t+h)), 0, 1) \quad (8)$$

where $w_L + w_R = 1$.

Mode(t) = Normal if $OSI \leq \tau_1$; Resilience if $\tau_1 < OSI \leq \tau_2$; Emergency if $OSI > \tau_2$

Figure 3 visualizes how forecast-derived stress indicators are fused into the OSI and subsequently mapped into supervisory operating modes for converter-level adaptive control. The parameter settings and mode-triggering policy associated with the OSI are listed in Table 3.

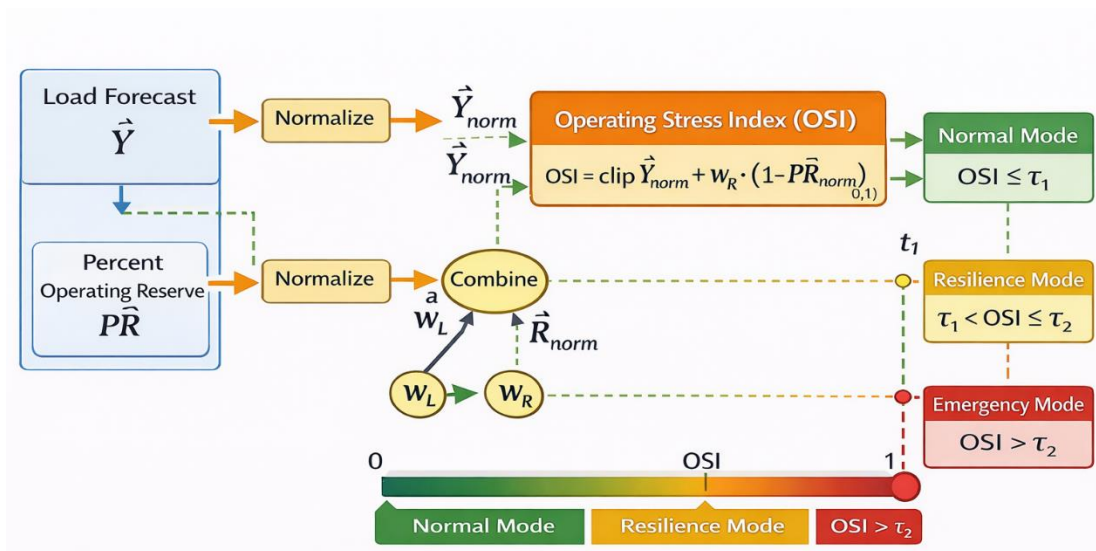


Figure 3. Operating Stress Index (OSI) computation and mode-triggering logic. The load forecast \hat{Y} and the percent operating reserve (PR), either observed or forecasted, are normalized to $[0,1]$ and combined as $OSI = \text{clip}(w_L \hat{Y}_{norm} + w_R(1 - PR_{norm}), 0, 1)$. The resulting OSI is then mapped into three operating modes using two thresholds: Normal ($OSI \leq \tau_1$), Resilience ($\tau_1 < OSI \leq \tau_2$), and Emergency ($OSI > \tau_2$).

Table 3. OSI definition parameters and resilience-mode policy.

Item	Symbol	Definition	Setting	Update rate	Notes
------	--------	------------	---------	-------------	-------

Load stress weight	w_L	Weight for s_L	0.60	per forecast (1h)	Emphasizes demand spikes
Reserve stress weight	w_R	Weight for s_R	0.40	per forecast (1h)	Reflects generation headroom
Normal/Resilience threshold	τ_1	OSI threshold	0.60	per forecast (1h)	Represents 60th percentile
Resilience/Emergency threshold	τ_2	OSI threshold	0.85	per forecast (1h)	Represents 85th percentile

4.4. Multi-Timescale Interface to Converter Control

OSI and the corresponding mode hint are updated on a slow timescale (minutes to hours), while FS-MPC executes on a fast timescale (milliseconds or microseconds). Such time-scale separation is consistent with established multi-loop converter control practice [21–23]. The interface does not override switching decisions; instead, it conditions the allowable range and prior of the FS-MPC weighting factors, and can optionally adjust protection margins (e.g., current limits) under emergency mode.

5. Proposed Method: Forecast-Guided KAN-Adaptive FS-MPC

5.1. Overview

The proposed controller retains the deterministic structure of FS-MPC and introduces only a lightweight nonlinear mapping from measured and forecast features to cost-function weights. At each control step, the KAN generates dynamic weights that balance voltage tracking and switching effort. The OSI provides supervisory context to modulate this mapping under forecasted vulnerable regimes. The real-time decision flow of the proposed controller is illustrated in Figure 4.

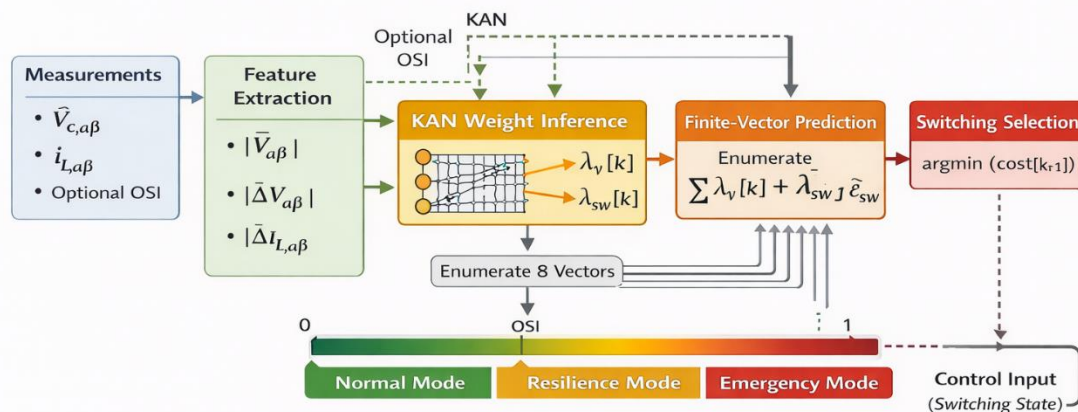


Figure 4. Real-time decision flow of the proposed KAN-adaptive finite-control-set MPC (FCS-MPC). At each control step ($T_s = 50 \mu s$), measurements ($(v_{c,\alpha\beta}, i_{L,\alpha\beta}, i_{o,\alpha\beta})$) are converted into compact features (e.g., $\|e_v\|$, $\|\Delta e_v\|$, $\|\Delta i_o\|$ and D_{sag}), optionally augmented by the slow-timescale OSI. A lightweight KAN governor outputs dynamic weights $\lambda_v[k]$ and $\lambda_{sw}[k]$, which are used to evaluate the cost $J(u) = \lambda_v[k]J_v(u) + \lambda_{sw}[k]J_{sw}(u)$ over the eight admissible switching vectors $u \in U$. The optimal switching action is selected by $u^*[k] = \arg \min_{u \in U} J(u)$.

5.2. Feature Design for Online Weight Governance

We form a compact feature vector that captures fast electrical dynamics and disturbance severity, and optionally includes OSI as slow-timescale context. Recommended features include voltage error magnitude and slope, load current variation, and voltage sag depth indicator.

$$\varphi[k] = [\text{OSI}(t), \|e_v[k]\|, \|\Delta e_v[k]\|, \|\Delta i_o[k]\|, D_{sag}[k]] \quad (9)$$

The feature definitions used in the online governor are summarized in Table 4 and follow the notation in (7)–(8).

Table 4. Online features used by the KAN weight governor and their physical interpretations.

Feature	Symbol	Definition (example)	Update rate	Physical meaning
Operating Stress Index	OSI(t)	From load forecast + reserve margin	minutes–hours	Forecasted vulnerability / regime context
Voltage error magnitude	$\ e_v[k]\ $	$\ v_{\alpha\beta}^* - v_{c,\alpha\beta}\ $	per T_s	Voltage regulation urgency
Voltage error slope	$\ \Delta e_v[k]\ $	$\ e_v[k] - e_v[k-1]\ $	per T_s	Transient aggressiveness indicator
Load current variation	$\ \Delta i_o[k]\ $	$\ i_o[k] - i_o[k-1]\ $	per T_s	Load shock / ramp severity
Sag depth	$D_{sag}[k]$	$V_{PCC,rms}/V_{nom}$ (or $\alpha\beta$ equivalent)	per T_s	Fault severity cue

5.3. KAN Weight Governor and Bounded Adaptive Weights

A Kolmogorov–Arnold Network (KAN) is used as a lightweight nonlinear governor that maps the feature vector to the FS-MPC weighting factors. Let the compact input feature vector at control step k be:

$$x[k] = [\text{OSI}(t), \|e_v[k]\|, \|\Delta e_v[k]\|, \|\Delta i_o[k]\|, D_{sag}[k]]^T \quad (10)$$

where $\text{OSI}(t)$ is the slow-timescale Operating Stress Index, $\|e_v[k]\|$ is the PCC voltage-tracking error magnitude, $\|\Delta e_v[k]\|$ is the error slope, $\|\Delta i_o[k]\|$ is the output-current variation, and $D_{sag}[k]$ is the sag-depth indicator. In the proposed multi-timescale architecture, $\text{OSI}(t)$ is updated much more slowly than the converter control loop and is therefore treated as a quasi-static supervisory parameter within each fast FS-MPC interval.

Unlike a standard multilayer perceptron that relies on fixed node activations, KAN parameterizes nonlinear transformations through learnable univariate functions on edges. A generic KAN layer can be written as:

$$z_q^{(\ell+1)} = \sum_{p=1}^{n_\ell} \phi_{q,p}^{(\ell)}(z_p^{(\ell)}), \quad q = 1, \dots, n_{\ell+1} \quad (11)$$

where $z_p^{(\ell)}$ is the p -th input of layer ℓ , and $\phi_{q,p}^{(\ell)}(\cdot)$ is a learnable scalar nonlinear function associated with edge $(p \rightarrow q)$. Following the standard KAN construction, each edge function is represented by a spline expansion:

$$\phi_{q,p}^{(\ell)}(x) = a_{q,p}^{(\ell)}x + b_{q,p}^{(\ell)} + \sum_{m=1}^M c_{q,p,m}^{(\ell)} B_m(x) \quad (12)$$

where $B_m(x)$ denotes the m -th B-spline basis function, and $a_{q,p}^{(\ell)}$, $b_{q,p}^{(\ell)}$, and $c_{q,p,m}^{(\ell)}$ are learned parameters. This edge-wise spline parameterization is attractive for converter control because localized threshold-like nonlinearities can be represented directly by a small number of basis functions, while the learned mapping remains visually inspectable through the spline curves.

The KAN outputs a raw weight vector:

$$\tilde{\lambda}[k] = \begin{bmatrix} \tilde{\lambda}_v[k] \\ \tilde{\lambda}_{sw}[k] \end{bmatrix} = f_{KAN}(x[k]) \quad (13)$$

which is then processed through explicit safety constraints before being passed to the predictive controller. First, each component is clipped to a certified interval:

$$\lambda_i[k] = \text{clip}(\tilde{\lambda}_i[k], \lambda_{i,\min}, \lambda_{i,\max}), \quad i \in \{v, sw\} \quad (14)$$

Second, a rate limiter is applied to avoid abrupt inter-sample changes:

$$\lambda_i[k] = \lambda_i[k-1] + \text{sat}_{[-\delta_{\lambda_i}, \delta_{\lambda_i}]}(\tilde{\lambda}_i[k] - \lambda_i[k-1]) \quad (15)$$

which guarantees

$$|\lambda_i[k] - \lambda_i[k-1]| \leq \delta_{\lambda_i} \quad (16)$$

Equations (12)–(14) define the bounded adaptive-weight mechanism used throughout this paper. These constraints serve two purposes. First, they enforce engineering safety by preventing the adaptive governor from generating numerically extreme or rapidly oscillating weights. Second, they expose explicit and auditable control boundaries: before deployment, the learned spline maps and the resulting admissible weight set can be scanned offline to verify that all reachable $\lambda[k]$ remain within a designer-certified region. In this sense, the proposed KAN governor follows a safety-by-design philosophy rather than unconstrained black-box adaptation.

From an implementation perspective, the deployed KAN remains lightweight, with approximately 240 parameters in the embedded realization reported in Section 8.4. This compact size is consistent with the objective of preserving microsecond-level feasibility on DSP-class hardware while still providing a richer nonlinear mapping than rule-based scheduling or static tuning.

5.4. Discrete-Time Stability Interpretation of the Adaptive-Weight FS-MPC

Because the proposed controller employs time-varying cost-function weights, a stability-oriented interpretation is required to show that online adaptation does not destroy the closed-loop convergence behavior of the underlying FS-MPC. Here we provide a practical discrete-time boundedness argument under explicit design assumptions. The result should be interpreted as a local / residual-set guarantee under bounded adaptation, rather than as a universal global asymptotic stability proof for all operating conditions.

Let the PCC voltage-tracking error in the stationary frame be:

$$e_v[k] = v_{\alpha\beta}^*[k] - v_{c,\alpha\beta}[k] \quad (17)$$

and define the Lyapunov candidate

$$V[k] = e_v^T[k]e_v[k] = \|e_v[k]\|_2^2 \quad (18)$$

At each control step, the adaptive FS-MPC selects the switching action $u^*[k] \in U$ that minimizes the stage cost

$$J(u; k) = \lambda_v[k] \|e_v^+(u; k)\|_2^2 + \lambda_{sw}[k] J_{sw}(u; k) \quad (19)$$

where $e_v^+(u; k)$ is the one-step-ahead predicted voltage error under candidate switch u , and $J_{sw}(u; k) \geq 0$ denotes the switching-effort penalty.

We make the following standard design assumptions.

Assumption 1. The discretized plant model used by FS-MPC is locally accurate over one sampling interval, and the plant states remain inside the admissible operating region enforced by current and voltage protection.

Assumption 2. For the nominal static-weight controller, there exists at least one admissible switching action such that, in a neighborhood of the regulation manifold,

$$V[k+1] - V[k] \leq -\alpha \|e_v[k]\|_2^2 + \eta \quad (20)$$

for some $\alpha > 0$ and bounded modeling residual $\eta \geq 0$.

Assumption 3. The adaptive weights satisfy

$$\lambda_{i,\min} \leq \lambda_{i[k]} \leq \lambda_{i,\max}, \quad |\lambda_{i[k]} - \lambda_{i[k-1]}| \leq \delta_{\lambda_i}, \quad i \in \{v, sw\} \quad (21)$$

as enforced by (12)–(14).

Under Assumption 3, the adaptive stage cost remains uniformly equivalent to the nominal voltage-error objective, because $\lambda_v[k]$ is strictly positive and bounded away from zero. Therefore, the action selected by minimizing (17) cannot arbitrarily neglect voltage regulation. In particular, there exist constants $c_1, c_2 > 0$ such that:

$$c_1 \|e_v^+(u; k)\|_2^2 \leq J(u; k) \leq c_2 (\|e_v^+(u; k)\|_2^2 + J_{sw}(u; k)) \quad (22)$$

for all admissible $u \in U$ and all k .

The key effect of adaptation is that the optimizer is evaluated using $\lambda[k]$ instead of $\lambda[k-1]$. Because the weight update is rate-limited, the perturbation in the optimization landscape between two consecutive samples is bounded. Denoting this perturbation by $\Delta_\lambda[k]$, one can write

$$|J(u; k) - J(u; k-1)| \leq L_\lambda \|\lambda[k] - \lambda[k-1]\|_\infty \leq L_\lambda \delta_\lambda \quad (23)$$

where L_λ is a local Lipschitz constant of the stage cost with respect to the weights, and $\delta_\lambda = \max_i \delta_{\lambda_i}$.

Combining (18) and (21), the closed-loop Lyapunov difference under the adaptive controller satisfies

$$\Delta V[k] = V[k+1] - V[k] \leq -\alpha \|e_v[k]\|_2^2 + \eta + \gamma \delta_\lambda \quad (24)$$

for some constant $\gamma > 0$ that captures the sensitivity of the selected action to bounded cost perturbations. Hence, whenever

$$\|e_v[k]\|_2^2 > \frac{\eta + \gamma\delta_\lambda}{\alpha} \quad (25)$$

we obtain $\Delta V[k] < 0$. This implies that the voltage-tracking error converges to the compact residual set

$$\Omega = \{e_v: \|e_v\|_2^2 \leq \frac{\eta + \gamma\delta_\lambda}{\alpha}\} \quad (26)$$

Equation (24) provides the desired engineering interpretation: the proposed KAN-driven adaptive weighting does not destabilize the underlying predictive controller as long as the weights remain bounded and sufficiently slow-varying. The residual-set radius expands with modeling mismatch η and adaptation aggressiveness δ_λ , which directly justifies the use of safety clamps and rate limiting in the real-time implementation. In other words, adaptation improves disturbance responsiveness, but only within an explicitly certified envelope.

This result also admits an impedance-oriented interpretation. Increasing $\lambda_v[k]$ during severe sags biases the optimizer toward stronger voltage-restoring action, while increasing $\lambda_{sw}[k]$ suppresses excessive switching activity and smooths the converter response. Because both weights are bounded, the adaptive governor reshapes the effective closed-loop output behavior without allowing uncontrolled excursions in the implied impedance profile. Although a full passivity proof is outside the scope of this paper, the bounded adaptation law is consistent with maintaining a well-behaved dissipative control posture under stressed operating regimes.

5.5. Dynamic Stage Cost and Switching Selection

With adaptive weights, the per-candidate stage cost becomes:

$$J(u; k) = \lambda_v[k] \|v_{\alpha\beta}^*[k+1] - \hat{v}_{c,\alpha\beta}[k+1|u]\|_2^2 + \lambda_{sw}[k] J_{sw}[u; k] \quad (27)$$

where $\hat{v}_{c,\alpha\beta}[k+1|u]$ denotes the one-step-ahead predicted PCC voltage under candidate switching vector u . The selected switching action is then:

$$u^*[k] = \arg \min_{u \in \mathcal{U}} J(u; k), |\mathcal{U}| = 8 \text{ for a two-level VSL} \quad (28)$$

The role of adaptation is therefore supervisory rather than structural: KAN does not replace the deterministic FS-MPC optimizer, but only reshapes the trade-off between voltage support and switching effort within certified bounds. This preserves the transparent finite-vector decision process of conventional FS-MPC while enabling regime-aware disturbance response.

Algorithm 1 OSI supervisory update (slow-timescale).

- 1: Input: load forecast \hat{Y}_{t+h} and percent operating reserve forecast PR_{t+h}
 - 2: Normalize \hat{Y}_{t+h} and PR_{t+h} to obtain S_L and S_R
 - 3: Compute $OSI(t+h) = \text{clip}(w_L S_L + w_R (1 - S_R), 0, 1)$
 - 4: Determine the operating mode using thresholds τ_1 and τ_2
 - 5: Set $OSI(t) \leftarrow OSI(t+h)$, output $OSI(t)$ and $Mode(t)$ to the controller, and update every $T_f = 60$ min
-

Algorithm 2. Forecast-guided KAN-adaptive FS-MPC (per control step k)

- 1: Input: measured $v_{c,\alpha\beta}[k]$, $i_{L,\alpha\beta}[k]$, $i_{o,\alpha\beta}[k]$, and supervisory signal $OSI(t)$
 - 2: Compute $e_v[k]$, $\Delta e_v[k]$, $\Delta i_o[k]$, and $D_{sag}[k]$
 - 3: Form $\phi[k]$, perform KAN inference, and obtain $\lambda_v[k]$ and $\lambda_{sw}[k]$ with clamping/rate limiting
 - 4: For each candidate switching state $u \in U$ (8 candidates)
 - 4.1: Predict $x^{(u)}[k+1]$
 - 4.2: Evaluate $J^{(u)}$
 - 5: Apply $u^* = \arg \min_u J^{(u)}$ and update the PWM/gating command
-

6. Experimental Setup

6.1. High-Frequency Microgrid Profile Synthesis

Converter-level validation requires microsecond-to-millisecond dynamics. Daily utility load data are too coarse to excite converter transients; therefore, high-frequency industrial microgrid profiles are synthesized to reflect (i) industrial load dominance and (ii) high PV penetration, while also incorporating explicit extreme events such as severe voltage sags and islanding-like transitions. The procedure used to generate this high-rate validation profile is summarized in Figure 5.

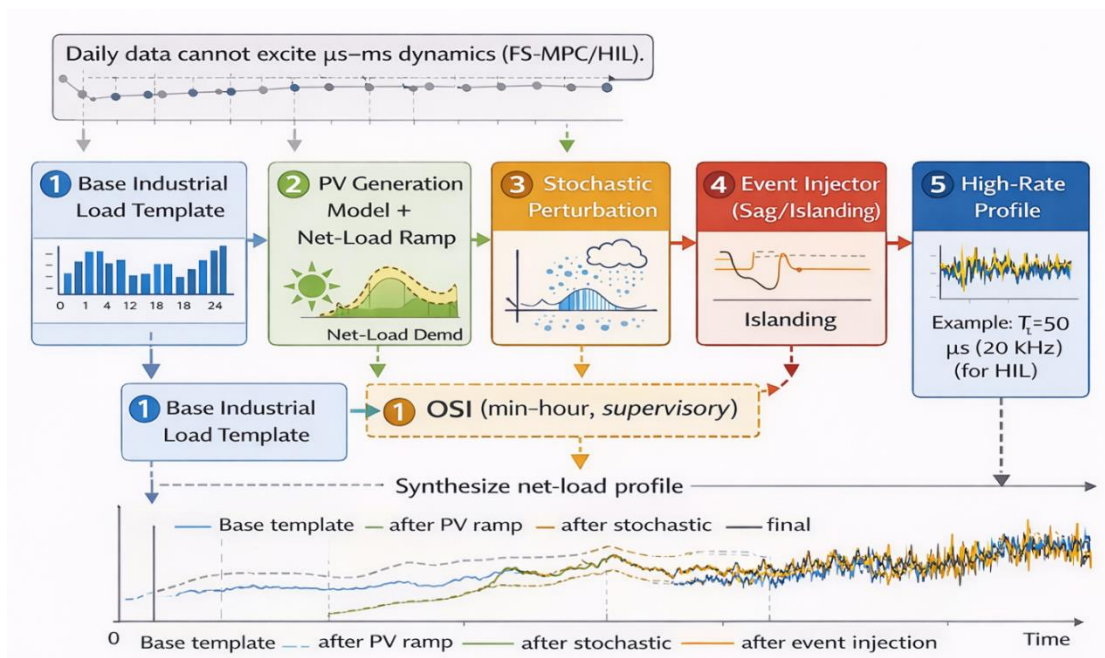


Figure 5. High-rate net-load profile synthesis for validating μs – ms converter control. Daily-resolution utility data is insufficient to excite fast converter dynamics; therefore, we synthesize a high-rate net-load profile by composing an industrial load template with PV-generation-induced net-load ramps and stochastic perturbations, followed by an event injector that imposes sag/islanding-like transients (depth/duration and weak-grid transitions). The final profile is sampled at 10–100 kHz (e.g., $T_s = 50 \mu\text{s}$, 20 kHz) for HIL execution, while OSI is used only as a slow-timescale supervisory tag or event-conditioning signal.

6.2. Disturbance Scenarios and Event Windows

We define event windows that include disturbance onset, degraded operation, and recovery. The main evaluation scenarios include severe symmetrical and asymmetrical sags, together with an islanding-like transition coupled with load/PV variation. Each scenario is parameterized to enable systematic stress testing under fault-dominant resilience events. The detailed scenario settings used in the main resilience evaluation are listed in Table 5, which summarizes the disturbance magnitudes, durations, and operating-condition combinations considered in this study.

Table 5. Disturbance scenarios used in the main resilience evaluation.

Scenario	Event type	Sag depth	Duration	Load step / ramp	PV condition	Notes
S1	Severe symmetrical sag	50%	10 cycles (166 ms)	Nominal continuous	Full MPPT	Standard LVRT test
S2	Extreme asymmetrical fault	70% (Phase A)	5 cycles (83 ms)	Nominal continuous	Full MPPT	High unbalance stress
S3	Islanding transition	100% (grid loss)	Continuous	Step 0.5 to 1.0 p.u.	Drops 50% (cloud)	Worst-case compound event

6.3. Baselines and Implementation Details

We compare three controller variants: static-weight FS-MPC, MLP-adaptive FS-MPC, and the proposed forecast-guided KAN-adaptive FS-MPC. All methods share the same prediction model, constraints, and sampling time to ensure a fair comparison. The compared controller variants and their key design differences are summarized in Table 6, including the distinctions between fixed weighting, MLP-based adaptation, and the proposed OSI-guided KAN-based adaptive weighting.

Table 6. Compared controllers and key design differences.

Method	Weights	Uses OSI	Model type	Interpretability	Notes
B1 Static FS-MPC	fixed λ_v, λ_{sw}	No	Deterministic	High	offline tuned
B2 MLP-adaptive	learned	Optional	MLP black-box	Low	same features
B3 Proposed KAN	KAN learned + OSI	Yes	KAN spline-on-edges	Medium-High	bounded, rate-limited

6.4. Hardware-in-the-Loop (HIL) Setup and Timing Budget

We validate real-time feasibility using a HIL platform where the plant (inverter, filter, microgrid) runs on a real-time simulator and the controller runs on an embedded target. We report worst-case execution time per control step, including KAN inference and FS-MPC candidate evaluation, and verify that the total latency remains below the sampling period T_s . The overall HIL platform configuration and the measured control-step timing breakdown are illustrated in Figure 6.

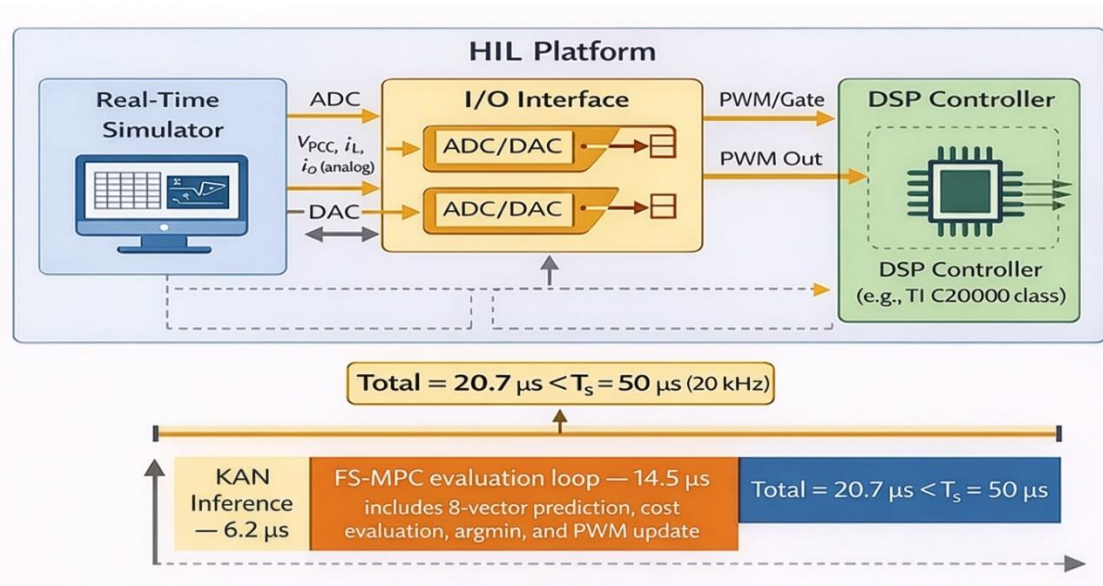


Figure 6. Hardware-in-the-loop (HIL) validation platform and control-step timing summary. The real-time simulator emulates the inverter–filter–microgrid plant and exchanges analog measurements with the DSP controller through the I/O interface, while the controller returns PWM/gate commands to the simulator. The measured execution time on the target DSP consists of $6.2 \mu s$ for KAN inference and $14.5 \mu s$ for the FS-MPC evaluation loop, yielding a total worst-case control-step time of $20.7 \mu s$. This remains below the sampling period $T_s = 50 \mu s$ and leaves a 58.6% timing margin.

7. Evaluation Metrics

7.1. Power Conversion and Constraint Metrics

We evaluate resilient power-conversion performance along three complementary dimensions: (i) regulation and power quality at the PCC, (ii) protection and constraint compliance, and (iii) switching stress and real-time feasibility of the control step.

Voltage regulation is evaluated using the PCC voltage vector in the stationary frame, $v_{c,\alpha\beta}(t)$, and its reference $v_{\alpha\beta}^*(t)$. We report both peak and RMS errors. The instantaneous voltage-magnitude tracking error is defined as $E_{v,pu}(t) = \|v_{c,\alpha\beta}(t) - v_{\alpha\beta}^*(t)\| / V_{nom}$ (p.u.), where V_{nom} is the nominal PCC voltage magnitude.

Because GFM operation must sustain a coherent voltage waveform, we also report frequency deviation estimated directly from the measured stationary-frame voltage: $\theta(t) = \text{unwrap}(\text{atan2}(v_{\beta}(t), v_{\alpha}(t)))$ and $f_{inst}(t) = (\frac{1}{2}\pi) \cdot \frac{d\theta(t)}{dt}$. Frequency metrics are computed over the same event windows used for voltage resilience metrics, with a short smoothing filter to reduce ripple-induced differentiation noise.

Protection and constraint compliance are summarized by the inductor current peak $I_{pk} = \max_t \|i_{L,\alpha\beta}(t)\|$ and by the count (or duration) of current-limit violations where $\|i_{L,\alpha\beta}(t)\| > I_{max}$, consistent with fault ride-through requirements under severe sags [3].

Switching stress is reported by the switching-transition count N_{sw} within an event window and by the average switching-transition rate (transitions/s) over steady segments. N_{sw} is computed from the Hamming distance between consecutive switching states, i.e., the number of phase-leg changes per control update.

Power quality is summarized by voltage THD computed over a post-recovery steady-state window. After excluding the transient portion, we apply an FFT to the PCC voltage magnitude (or phase voltage) over an integer number of fundamental cycles and compute $\text{THD} = \frac{\sqrt{\sum_{h=2}^H V_h^2}}{V_1}$, where

V_1 is the fundamental amplitude and V_h are harmonic amplitudes up to order H (reported in Table 7).

7.2. Event-Window Resilience Metrics (Tail-Oriented)

To align converter evaluation with resilience-oriented practice, we define event windows that cover disturbance onset, degraded operation, and recovery, and we adopt tail-oriented metrics that explicitly capture severity and restoration speed [2,3,10,24]. For each scenario, the event window $W = [t_0 - T_{pre}, t_{clr} + T_{post}]$ includes a short pre-event baseline (T_{pre}) and a post-clearance recovery horizon (T_{post}), where t_0 is the disturbance onset time and t_{clr} is the clearance time (e.g., sag removal or regime restoration).

Let $E_{v,pu}(t) = \frac{||v_{c,\alpha\beta}(t) - v^*_{\alpha\beta}(t)||}{V_{nom}}$ (p.u.) denote the voltage tracking error magnitude. We define a tolerance band ϵ (p.u.) and a hold time Δ to avoid declaring recovery on a single-cycle crossing. In this work, ϵ is chosen in the 2–5% range depending on OSI regime (Table 7), and Δ is set to 1–2 fundamental cycles unless otherwise stated.

Worst-case deviation:

$$E_{max} = \max_{t \in W} E_{v,pu}(t) \quad (29)$$

Recovery time to tolerance band ϵ :

$$T_{rec} = \min \{t \geq t_{clr} : E_{v,pu}(\tau) \leq \epsilon \text{ for all } \tau \in [t, t_{clr}]\} - t_{clr} \quad (30)$$

Degradation area (exceedance integral):

$$A_{deg} = \int_{t \in W} \max(0, E_{v,pu}(t) - \epsilon) dt \quad (31)$$

In implementation, the above metrics are computed from sampled data at the controller sampling period T_s . Accordingly, E_{max} is the maximum of $e_v[k]$ over the window indices, T_{rec} is the earliest post-clearance time at which $e_v[k]$ stays within ϵ for $N_{hold} = \Delta/T_s$ consecutive samples, and A_{deg} is approximated by $A_{deg} \approx T_s \cdot \sum_k \max(0, e_v[k] - \epsilon)$.

7.3. OSI-Conditioned Reporting

Because OSI quantifies forecasted vulnerability from load and reserve signals, we stratify performance by OSI regime (Normal/Resilience/Emergency) to demonstrate that forecast guidance provides the largest benefit under stressed operating states. Each disturbance event is assigned an OSI regime using $OSI(t_0)$ (or the average OSI over the pre-event baseline).

For each metric, we report (i) per-scenario values, (ii) OSI-regime aggregates, and (iii) overall aggregates across scenarios. When reporting improvements, we use relative change with respect to a baseline (e.g., static-weight FS-MPC): $\Delta\% = (M_{baseline} - M_{proposed})/M_{baseline} \times 100\%$ for metrics where smaller is better (E_{max} , T_{rec} , A_{deg} , THD, I_{pk}), while switching-effort changes (N_{sw}) are reported alongside regulation improvements to highlight the trade-off.

Table 7. Summary of evaluation metrics and measurement protocol.

Metric	Symbol	Definition	Unit	Purpose
Worst-case deviation	E_{max}	$\max E_{v,pu}(t)$ over event window	p.u.	Tail severity
Recovery time	T_{rec}	time to re-enter tolerance band	ms	Restoration speed

Degradation area	A_{deg}	integral exceedance above tolerance band	p.u.-ms	Cumulative impact
Switching effort	N_{sw}	average switching frequency	kHz	Efficiency/thermal stress
Peak current	I_{pk}	max output/inductor current	A	Protection

8. Results and Discussion

8.1. Main Results Under Severe Voltage Sags

Figure 7 provides representative PCC voltage trajectories, while Table 8 reports quantitative performance across the defined severe-sag and regime-shift scenarios (S1–S3). Across all cases, the proposed forecast-guided KAN-adaptive FS-MPC achieves the lowest worst-case deviation and degradation area, and the fastest recovery, without increasing switching effort beyond practical limits.

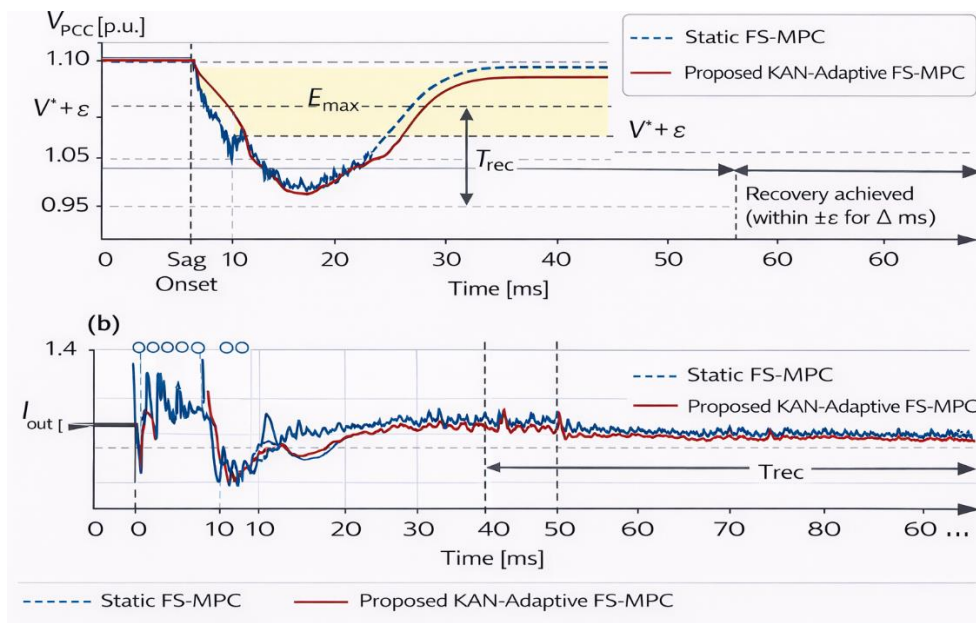


Figure 7. Severe voltage-sag transient response and resilience metrics. (a) PCC voltage V_{PCC} (p.u.) under a representative severe sag for static-weight FS-MPC and the proposed forecast-guided KAN-adaptive FS-MPC. The tolerance band $V^* \pm \varepsilon$ is highlighted to define the worst-case deviation E_{max} and recovery time T_{rec} (time to re-enter and remain within the band for Δ ms after sag clearance). (b) Output current I_{out} (p.u.) under the same event, illustrating current peaking behavior and settling during disturbance and recovery.

Table 8. Main results across severe-sag and regime-shift scenarios.

Scenario	Method	E_{max}	T_{rec}	A_{deg}	THD	I_{pk}	N_{sw}	Notes
S1 (50% Sag)	B1: Static FS-MPC	0.45	35	8.5	5.2	45.2	12.5	
S1 (50% Sag)	B2: MLP-adaptive	0.28	18	3.2	4.1	38.5	11.2	
S1 (50% Sag)	B3: Proposed KAN	0.16	8	1.1	2.9	32.1	10.5	Best

S2 (70% Asym)	B1: Static FS-MPC	0.62	52	14.8	6.8	51.0	12.5	
S2 (70% Asym)	B2: MLP-adaptive	0.41	25	6.5	5.0	42.4	11.8	
S2 (70% Asym)	B3: Proposed KAN	0.25	12	2.4	3.4	35.6	10.8	Best
S3 (Islanding)	B1: Static FS-MPC	0.85	>100	25.0	8.5	55.3	12.5	
S3 (Islanding)	B2: MLP-adaptive	0.55	45	12.4	6.2	48.1	12.0	
S3 (Islanding)	B3: Proposed KAN	0.30	18	4.2	3.8	38.5	11.0	Best

Under the 50% symmetrical sag (S1), the proposed controller reduces the worst-case deviation E_{max} from 0.45 p.u. (B1) to 0.16 p.u. (B3), a 64.4% reduction, and shortens the recovery time T_{rec} from 35 ms to 8 ms (77.1% faster). The cumulative degradation area A_{deg} drops from 8.5 to 1.1 p.u.-ms (87.1% reduction). Compared with the MLP-adaptive governor (B2), B3 further lowers E_{max} (0.28 \rightarrow 0.16 p.u.) and T_{rec} (18 \rightarrow 8 ms), while also improving THD (4.1% \rightarrow 2.9%) and reducing peak current (38.5 \rightarrow 32.1 A).

For the extreme asymmetrical fault (S2), B3 decreases E_{max} from 0.62 p.u. to 0.25 p.u. (59.7% reduction) and reduces T_{rec} from 52 ms to 12 ms (76.9% faster) relative to static-weight FS-MPC. A_{deg} is reduced from 14.8 to 2.4 p.u.-ms (83.8% reduction), accompanied by a meaningful reduction in peak current (51.0 \rightarrow 35.6 A) and lower THD (6.8% \rightarrow 3.4%).

In the islanding transition (S3), the static-weight controller fails to recover within 100 ms ($T_{rec} > 100$ ms), whereas the proposed method restores the PCC voltage within 18 ms, representing at least an 82% reduction in recovery time. B3 also reduces E_{max} from 0.85 p.u. to 0.30 p.u. (64.7% reduction) and A_{deg} from 25.0 to 4.2 p.u.-ms (83.2% reduction), while cutting peak current from 55.3 A to 38.5 A. Notably, these resilience gains are achieved with slightly reduced switching effort (12.5 \rightarrow 11.0 kHz), indicating that adaptation improves control effectiveness rather than simply increasing aggressiveness.

8.2. Ablation: OSI Guidance and KAN vs MLP

Table 9 reports an ablation study under Scenario S3 (Islanding Transition), which is the most severe regime-shift case in Table 8, to isolate the effects of OSI guidance and governor type. Adding OSI guidance to KAN-based adaptation (B3a \rightarrow B3b) reduces E_{max} from 0.42 to 0.30 p.u. (28.6% reduction) and shortens T_{rec} from 28 ms to 18 ms (35.7% faster), while halving A_{deg} from 8.5 to 4.2 p.u.-ms. Under the same Scenario S3 setting, replacing an MLP governor (B2) with the lightweight KAN governor (B3b) yields substantially better tail performance (E_{max} 0.55 \rightarrow 0.30 p.u., T_{rec} 45 \rightarrow 18 ms, A_{deg} 12.4 \rightarrow 4.2 p.u.-ms) with comparable switching effort.

Table 9. Ablation study: OSI guidance and governor type under Scenario S3 (Islanding Transition).

Variant	Uses OSI	Governor	E_{max}	T_{rec}	A_{deg}	N_{sw}
B3a KAN-adaptive (no OSI)	No	KAN	0.42	28	8.5	11.5
B2 MLP-adaptive	Yes	MLP	0.55	45	12.4	12.0
B3b (Proposed)	Yes	KAN	0.30	18	4.2	11.0

Note: All variants in Table 9 are evaluated under the Scenario S3 islanding-transition event defined in Table

8.

8.3. Interpretability: Learned KAN Splines

Figure 8 visualizes representative learned spline functions mapping key features to adaptive weights. Specifically, panels (a) and (b) show how the sag-depth feature D_{sag} and the Operating Stress Index (OSI) modulate the voltage-tracking weight λ_v , whereas panel (c) shows how the voltage-error-slope feature $\|\Delta e_v\|$ modulates the switching-effort weight λ_{sw} . These curves provide an auditable view of how the controller shifts emphasis between voltage support and switching-effort moderation across regimes. To remain consistent with (14), the plotted output bounds should be interpreted as the certified intervals $[\lambda_{i,\min}, \lambda_{i,\max}]$, where $i \in \{v, sw\}$. The markers δ_{sag} and δ_e denote representative feature-threshold locations at which the corresponding learned spline enters its steep transition region; they are introduced for interpretability in Figure 8 rather than as additional optimization variables.

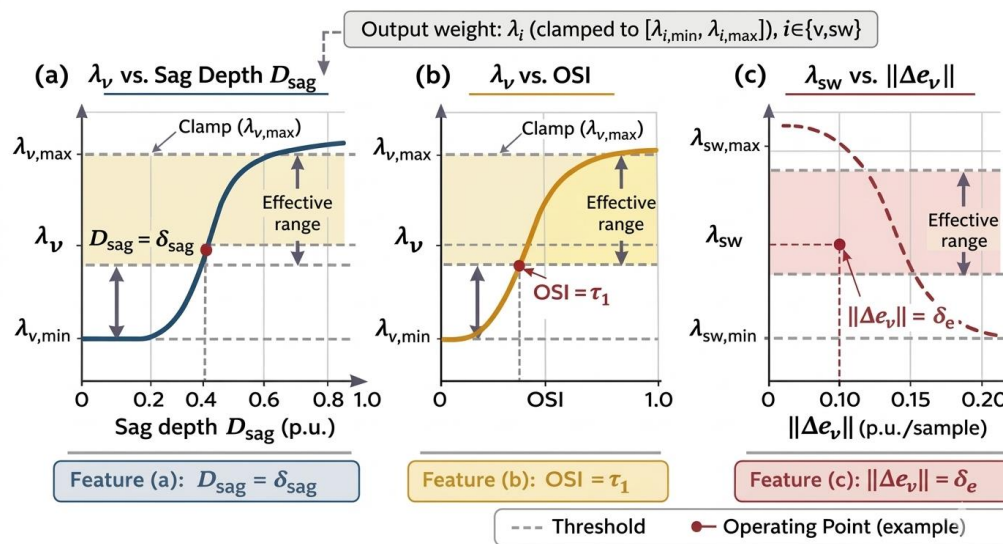


Figure 8. Interpretable KAN edge-spline mappings for online weight governance. Representative learned spline functions show how (a) sag depth D_{sag} modulates the voltage-tracking weight λ_v , (b) the Operating Stress Index (OSI) modulates λ_v , and (c) the voltage-error-slope feature $\|\Delta e_v\|$ modulates the switching-effort weight λ_{sw} . The plotted output bounds follow the certified intervals $[\lambda_{i,\min}, \lambda_{i,\max}]$, $i \in \{v, sw\}$, consistent with (14). The markers δ_{sag} and δ_e denote representative feature-threshold locations used to illustrate where the learned spline enters its high-gain transition region.

8.4. Real-Time Feasibility

Table 10 reports embedded execution time on a TI TMS320F28379D (200 MHz). The deployed KAN governor uses approximately 240 parameters and requires 6.2 μs per inference, while the full FS-MPC evaluation loop, including prediction, cost evaluation, switching-state selection, and PWM update, requires 14.5 μs . The total worst-case control-step time is 20.7 μs , comfortably below the sampling period $T_s = 50.0 \mu\text{s}$, leaving a 58.6% timing margin and supporting real-time DSP-oriented deployment.

Table 10. Real-time feasibility: execution time and timing margin.

Platform	KAN params	KAN time	FS-MPC eval time	Total step time	T_s	Margin
----------	---------------	-------------	---------------------	--------------------	-------	--------

TI TMS320F28379D (200 MHz)	~240	6.2 μ s	14.5 μ s	20.7 μ s	50.0 μ s	58.6%
-------------------------------	------	-------------	--------------	--------------	--------------	-------

8.5. Discussion Summary

Overall, the evidence supports a resilience-by-design interpretation. The stationary alpha-beta formulation preserves a well-defined control representation under deep faults, OSI introduces vulnerability-aware preparedness by coupling load forecasts with reserve adequacy, and the spline-based KAN governor provides a parameter-efficient and auditable adaptation surface. Across S1–S3, this combination consistently improves tail-oriented resilience metrics (worst-case deviation, recovery time, and degradation area) while also reducing peak current and maintaining practical switching effort.

9. Conclusions

This paper proposed a forecast-guided KAN-adaptive FS-MPC framework for resilient power conversion in PLL-free grid-forming BESS inverters. By formulating FS-MPC in the stationary alpha-beta frame and introducing an Operating Stress Index (OSI) derived from load forecasts and reserve adequacy, the approach couples slow-timescale vulnerability awareness with fast-timescale adaptive predictive control. Quantitatively, the proposed method reduces worst-case deviation by about 60–65% across severe sag and islanding scenarios, and shortens recovery time from 35 ms to 8 ms (S1) and from >100 ms to 18 ms (S3), while maintaining real-time feasibility with a 58.6% timing margin on a DSP platform.

Author Contributions: Conceptualization, S.-E. Tsai; Methodology, S.-E. Tsai; Software, S.-E. Tsai; Validation, S.-E. Tsai and W.-C. Sun; Formal analysis, S.-E. Tsai; Investigation, S.-E. Tsai; Data curation, S.-E. Tsai and W.-C. Sun; Writing—original draft, S.-E. Tsai; Writing—review & editing, S.-E. Tsai. All authors have read and agreed to the published version of the manuscript.

Funding: This research was financially supported by En-Shou Investment Co., Ltd., Taiwan, and Ching Lung Agricultural Technology Co., Ltd., Taiwan.

Data Availability Statement: The original contributions presented in this study are included in the article. Additional supporting data are available from the corresponding author upon reasonable request.

Acknowledgments: This work was supported by En-Shou Investment Co., Ltd. and Ching Lung Agricultural Technology Co., Ltd. The authors would also like to thank the AI Center, Chang Jung Christian University, for providing essential computational resources and technical support. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

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

Abbreviations

The following abbreviations are used in this manuscript:

BESS	Battery Energy Storage System
DSP	Digital Signal Processor
FS-MPC	Finite-Set Model Predictive Control
GFM	Grid-Forming
HIL	Hardware-in-the-Loop
KAN	Kolmogorov–Arnold Network
MLP	Multilayer Perceptron
OSI	Operating Stress Index

PCC	Point of Common Coupling
PLL	Phase-Locked Loop
PWM	Pulse-Width Modulation
THD	Total Harmonic Distortion
VSI	Voltage Source Inverter

References

1. Bahrani, B.; Ravanji, M.H.; Kroposki, B.; Ramasubramanian, D.; Guillaud, X.; Prevost, T. Grid-Forming Inverter-Based Resource Research Landscape: Understanding the Key Assets for Renewable-Rich Power Systems. *IEEE Power Energy Mag.*, **2024**, 22(2), 18–29.
2. Bhusal, N.; Abdelmalak, M.; Kamruzzaman, M.; Benidris, M. Power System Resilience: Current Practices, Challenges, and Future Directions. *IEEE Access.*, **2020**, 8, 18064–18086.
3. Panteli, M.; Mancarella, P. Modeling and Evaluating the Resilience of Critical Electrical Power Infrastructure to Extreme Weather Events. *IEEE Syst. J.*, **2017**, 11(3), 1733–1742.
4. Rodriguez, J.; Pontt, J.; Silva, C.A.; Correa, P.; Lezana, P.; Cortés, P.; Ammann, U. Predictive Current Control of a Voltage Source Inverter. *IEEE Trans. Ind. Electron.*, **2007**, 54(1), 495–503.
5. Vazquez, S.; et al. Model Predictive Control: A Review of Its Applications in Power Electronics. *IEEE Ind. Electron. Mag.*, **2014**, 8(1), 16–31.
6. Cortes, P.; et al. Guidelines for Weighting Factors Adjustment in Finite State Model Predictive Control of Power Converters and Drives. in *Proc. IEEE Int. Conf. Ind. Technol. (ICIT)*, Gippsland, Australia, 10–13 February., **2009**, 1-7.
7. Hu, L.; Lei, W.; Li, R.; Gao, Y. A Graphic Weighting Factor Design Method for Finite Control Set Model Predictive Control of Power Converters. *IEEE J. Emerg. Sel. Topics Power Electron.*, **2023**, 11(2), 1661–1671.
8. Novak, M.; Dragicevic, T. Weighting Factor Design in Model Predictive Control of Power Electronic Converters: An Artificial Neural Network Approach. *IEEE Trans. Ind. Electron.*, **2019**, 66(11), 8870–8880.
9. Alquannah, A.N.; Saleh, M.A.; Abu-Rub, H.; Ghayeb, A.; Bayhan, S.; Trabelsi, M. Double Deep Q-Networks Reinforcement Learning Based Dynamic Weighting Factor in FCS-MPC for Multilevel Inverters. in *Proc. IEEE 18th Int. Conf. Compat., Power Electron. Power Eng. (CPE-POWERENG)*., **2024**, 1-7.
10. Tsai, S.-E.; Sun, W.-C.; Aun, M.-H. Resilience-Oriented Daily Peak Load Forecasting on Utility Open Data Using a Dual-Stream Seasonal ARIMA and Kolmogorov–Arnold Network. *IEEE Access.*, **2026**, 14, 37606–37618.
11. Liu, Z.; Wang, Y.; Vaidya, S.; Ruehle, F.; Halverson, J.; Soljačić, M.; Hou, T.Y.; Tegmark, M. KAN: Kolmogorov-Arnold Networks. in *Proc. 13th Int. Conf. Learn. Represent. (ICLR)*, Singapore, April., **2025**, 24–28.
12. Shuai, H.; Li, F. Physics-Informed Kolmogorov-Arnold Networks for Power System Dynamics. *IEEE Open Access J. Power Energy.*, **2025**, 12, 46–58.
13. Novak, M.; Li, Y.; Zhao, S.; Wang, H.; Blaabjerg, F. Adaptive Control Design for Power Electronics Converters Using Kolmogorov-Arnold Networks. in *Proc. IEEE Energy Convers. Congr. Expo. (ECCE)*., Philadelphia, PA, USA., **2025**, 1–6.
14. Zhong, Q.-C. Virtual Synchronous Machines: A Unified Interface for Grid Integration. *IEEE Power Electron. Mag.*, **2016**, 3(4), 18–27.
15. Mahamedi, B.; Eskandari, M.; Fletcher, J.E.; Zhu, J. Sequence-Based Control Strategy With Current Limiting for the Fault Ride-Through of Inverter-Interfaced Distributed Generators. *IEEE Trans. Sustain. Energy.*, **2020**, 11(1), 165-174.
16. Baeckeland, N.; Chatterjee, D.; Lu, M.; Johnson, B.; Seo, G.-S. Overcurrent limiting in grid-forming inverters: A comprehensive review and discussion. *IEEE Trans. Power Electron.*, **2024**, 39, 14493–14511.
17. Wu, W.; Qiu, L.; Liu, X.; Ma, J.; Rodriguez, J.; Fang, Y. Data-Enabled Finite State Predictive Control for Power Converters via Adaline Neural Network. *IEEE Trans. Ind. Electron.*, **2025**, 72(3), 2244–2253.
18. She, Y.; Huo, X.; Tong, X.; Wang, C.; Fu, K. Multi-Sampling Rate Finite Control Set Model Predictive Control and Adaptive Method of Single-Phase Inverter. *Electronics.*, **2023**, 12(13), 2848.

19. Liu, X.; Qiu, L.; Fang, Y.; Wang, K.; Li, Y.; Rodríguez, J. Predictive control of voltage source inverter: An online reinforcement learning solution. *IEEE Trans. Ind. Electron.*, **2024**, 71(7), 6591–6600.
20. Xu, S.; Yao, C.; Ren, G.; Sun, Z.; Wu, S.; Ma, G. Weighting Factors Autotuning of FCS-MPC for Hybrid ANPC Inverter in PMSM Drives Based on Deep Residual Networks. *IEEE Trans. Power Electron.*, **2024**, 39(12), 16540–16552.
21. Rodríguez-Cortés, C.J.; Martínez-Rodríguez, P.R.; Langarica-Córdoba, D.; Rolan-Blanco, A.; Vázquez-Guzmán, G.; Villanueva-Loredo, J.A.; Sosa, J.M. Design and Control of a Modular High-Gain DC–DC Converter with Extensible Switched-Inductor Cells. *Electronics*, **2026**, 15(4), 897.
22. Villanueva-Loredo, J.A.; Martínez-Rodríguez, P.R.; Rodríguez-Cortés, C.J.; Langarica-Córdoba, D.; Hernández-Gómez, Á.; Guilbert, D. Analysis and Control Design of a Step-Up/Step-Down Converter for Battery-Discharge Voltage Regulation. *Electronics*, **2025**, 14(5), 877.
23. Beltrán, C.A.; Cisneros, R.; Langarica-Córdoba, D.; Ortega, R.; Díaz-Saldierna, L. Harnessing monotonicity to design an adaptive PI passivity-based controller for a fuel-cell system. *J. Process Control*, **2025**, 154, 103511.
24. Babu, K.V.S.M.; Dwivedi, D.; Chakraborty, P.; Yemula, P.K.; Pal, M. A comprehensive review on resilience definitions, frameworks, metrics, and enhancement strategies in electrical distribution systems. *Appl. Energy*, **2025**, 394, 126141.

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.