

Article

From Energy Efficiency to Energy Intelligence: Power Electronics as the Cognitive Layer of the Energy Transition

Nikolay Hinov ^{1,2} ¹ CoE “National Center of Mechatronics and Clean Technologies”, 1000 Sofia, Bulgaria; hinov@tu-sofia.bg² Department of Computer Systems, Faculty of Computer Systems and Technologies, Technical University of Sofia, 1000 Sofia, Bulgaria

Abstract

The exponential growth of artificial intelligence (AI), electrified transport, and renewable generation is accelerating a structural shift in how societies produce, deliver, and consume electricity. We argue that the next frontier is not incremental efficiency but Energy Intelligence (EI): the embedding of predictive analytics, adaptive control, and material-aware design directly into power-conversion hardware. In this view, power electronics functions as the cognitive layer that links digital intelligence to the physical flow of energy. Wide-bandgap (WBG) semiconductors—gallium nitride (GaN) and silicon carbide (SiC)—provide the material foundation for higher switching frequencies, superior power density, and real-time controllability, enabling compact and efficient converters for data-centers, EV charging, and grid-interactive resources. We formalize an EI reference architecture (predictive, adaptive, material-efficient, data-driven), review the convergence of AI methods with converter design and operation, and outline a GaN/SiC-enabled data-center power path as an illustrative case. Finally, we examine sustainability and sovereignty, highlighting exposure to critical materials (Ga, Si, In, rare earths) and proposing a roadmap that integrates technology, policy, and education. By reframing power electronics as an intelligent, learning infrastructure, this work sets an agenda for systems that are not only efficient but also self-optimizing, explainable, and resilient.

Keywords: artificial intelligence; energy intelligence; digital twin; GaN; power electronics; smart grid; SiC; sustainability; wide-bandgap semiconductors



Academic Editors: Luis M. Fernández-Ramírez, Ahmed Abu-Siada, Jean-Christophe Crebier, Zhiwei Gao, Kai Fu and Eladio Durán Aranda

Received: 23 October 2025

Revised: 15 November 2025

Accepted: 25 November 2025

Published: 27 November 2025

Citation: Hinov, N. From Energy Efficiency to Energy Intelligence: Power Electronics as the Cognitive Layer of the Energy Transition. *Electronics* **2025**, *14*, 4673. <https://doi.org/10.3390/electronics14234673>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Throughout most of the twentieth century, progress in electrical engineering was defined by the pursuit of efficiency—minimizing conversion losses, maximizing the utilization of electrical energy, and improving reliability [1–3]. Classical converter design, guided by linear control theory and silicon-based switches, delivered remarkable results: switching losses were minimized, voltage stresses were controlled, and the mean time between failures (MTBF) increased significantly [4]. For decades, these achievements shaped the foundations of modern energy systems.

However, as the global demand for electricity continues to expand through the electrification of transport, large-scale renewable integration, and the exponential growth of digital infrastructures, a new frontier has emerged. The challenge of the twenty-first century is no longer incremental efficiency, but dynamic intelligence [5–7]. Artificial intelligence (AI), machine learning (ML), and edge computing now interact with physical energy systems at scales and speeds previously unimaginable [8–10]. The electricity consumption of global

data-centers already exceeds 800 TWh annually and is projected to surpass 1000 TWh by 2030 [11,12]—roughly equivalent to the total generation of Japan or Germany. Simultaneously, electric vehicles (EVs) and heat pumps are reshaping load curves, demanding unprecedented coordination between distributed generation and consumption [13].

Traditional power infrastructures were designed for one-directional, predictable energy flows—from centralized generation to passive loads [14]. In contrast, contemporary networks are bidirectional, stochastic, and adaptive, requiring converters and controllers capable of perceiving, learning, and responding to real-time conditions [15,16]. This transition places power electronics at the very center of the digital–physical interface: every watt generated by photovoltaics, every kilowatt-hour stored in a battery, and every teraflop processed by an AI accelerator ultimately traverses a power-conversion stage [17].

Recent research shows that the boundaries between computation and conversion are rapidly dissolving [18]. Wide-bandgap (WBG) semiconductors such as gallium nitride (GaN) and silicon carbide (SiC) not only enhance efficiency and switching frequency but also form the physical substrate that enables intelligence to be embedded directly into hardware [19,20]. GaN devices operating beyond 1 MHz support fine-grained control loops that can host predictive or learning algorithms [21], while SiC modules in medium-voltage drives enable adaptive loss balancing through neural controllers [22].

At the same time, the material dimension of intelligence is gaining critical importance. The supply of gallium, indium, and rare-earth elements remains heavily concentrated geographically, creating strategic dependencies [23–25]. The supply of gallium, indium, and rare-earth elements is heavily concentrated geographically, creating strategic dependencies that directly affect device availability, cost, and long-term reliability. For converter designers, this means that material exposure must be treated as a design parameter alongside efficiency, thermal limits, and EMI constraints [26]. Achieving autonomy in WBG manufacturing, recycling, and circular material chains has therefore become as important as developing new control algorithms [27].

The convergence of AI and energy technologies defines an emerging systemic layer—cognitive energy systems [28–32]—capable of predictive, adaptive, and cooperative behavior through real-time feedback loops that connect converters, networks, and cloud platforms. Within this hierarchy, power electronics functions as the cognitive layer, translating digital decisions into physical voltages and currents with microsecond precision.

In this framework, energy efficiency evolves into energy intelligence [33–36]. Efficiency aims to minimize losses within fixed boundaries; intelligence expands those boundaries through learning and adaptation. Intelligent converters can anticipate load transients, predict degradation, and optimize performance according to context-specific objectives. At scale—across smart grids, data-centers, and EV infrastructures—such behaviors transcend individual devices, forming an emergent, self-optimizing network.

Beyond steady-state efficiency, modern converters face measurable bottlenecks under dynamic conditions: total harmonic distortion (THD) increases by 10–25% under fluctuating load/irradiance; transient overshoot grows with limited control bandwidth (2–5 kHz typical); and lifetime estimates deviate by 10–20% due to junction-temperature and ESR drift. Energy Intelligence (EI) addresses these limits with verifiable KPIs: $\approx 20\%$ THD reduction under disturbances, 15–30% overshoot reduction via predictive control, and 10–15% lower lifetime-model error through health-aware adaptation. These gains motivate the shift from “efficient but static” to “efficient and intelligent” conversion.

To clearly define the technological path, the transition from efficiency to intelligence must be shaped around measurable bottlenecks and verifiable performance metrics. Conventional high-efficiency converters remain vulnerable under non-ideal and rapidly varying conditions: THD typically increases by 10–25% during abrupt irradiance or load fluctua-

tions [37]; control bandwidths of 2–5 kHz limit transient response, causing overshoot of 15–30% [38]; and lifetime models exhibit 10–20% prediction error due to junction-temperature drift, ESR ageing, and magnetics saturation [39]. These limitations outline a concrete trajectory for technological evolution: (i) predictive control to expand effective bandwidth and reduce overshoot below 5–10%; (ii) adaptive gate-drive and modulation to achieve $\geq 20\%$ reduction in THD under disturbances; and (iii) health-aware operation to lower lifetime-model error below 5–7% [40]. Recent studies demonstrate that integrating data-driven prediction and WBG-based adaptive modulation can reduce dynamic losses by 8–12%, extend capacitor life by $\sim 15\%$, and maintain power-quality indices within IEEE 519 limits under stochastic load profiles [41].

These quantitative KPIs justify the necessity of moving beyond “static efficiency” toward “dynamic intelligence”: the marginal gains achievable through further silicon optimisation are saturated, whereas predictive, adaptive, and material-aware algorithms unlock performance regions that cannot be reached by conventional deterministic control alone. Thus, the technological path is defined by measurable improvements in THD, transient bandwidth, and lifetime accuracy—framing EI not as a conceptual shift, but as an empirically verifiable necessity for next-generation power-conversion systems.

The quantitative limits outlined above define a concrete technological trajectory for the evolution of power-conversion systems. Classical efficiency-driven design optimizes losses under fixed operating conditions but remains vulnerable to dynamic disturbances (e.g., 10–25% THD increase, 15–30% transient overshoot, and 10–20% lifetime-model error). The first step toward EI is therefore predictive control, which expands the effective control bandwidth and reduces transient deviations below 5–10%. The second step is adaptive operation—gate-drive and modulation self-adjustment enabling $\geq 20\%$ THD reduction under stochastic load or irradiance variability. The final step is health-aware operation, where electro-thermal telemetry is used to reduce lifetime-model error to $\leq 5\text{--}7\%$ and maintain safe operating regions under long-term degradation. Together, these stages form a measurable technological path from static efficiency to dynamic intelligence, demonstrating that EI is not a conceptual abstraction but an empirically grounded necessity for next-generation converters.

The following sections formalize this efficiency-to-intelligence pathway and trace how power electronics has evolved from non-adaptive hardware into the cognitive backbone of the modern energy transition.

Related Work and Positioning of the “Cognitive Layer” Concept

A number of existing frameworks already explore the interaction between power electronics, information flows, and intelligent control. Smart energy systems, cyber-physical grids, and digital-twin-enabled control architectures all aim to integrate sensing, computation, and optimisation across energy infrastructures. However, these approaches typically treat power electronics as actuators or passive endpoints, rather than as cognitive entities with embedded learning capability.

Smart-energy-system frameworks (European SET-Plan Digitalisation Roadmap) emphasize sector coupling, demand response, and ICT integration, but they generally consider converters as static interfaces implementing grid commands rather than locally adaptive components. Intelligence resides at supervisory or system level, not at the converter level.

Cyber-physical energy systems (CPES) introduce the notion of energy-information closed loops, where digital controllers interact with the physical grid through continuous feedback. While these architectures formalize the bidirectional flow between measurement and actuation, they do not specify how intelligence can be embedded into

the power-conversion hardware itself, nor do they link system behavior to material or semiconductor constraints.

Digital twins have been applied extensively to power converters (for accelerated testing and predictive maintenance). Yet these models operate externally to the converter and do not transform the converter into a self-learning energy-information node capable of local inference or adaptation.

Recent reviews on AI-enhanced power electronics and WBG-enabled architectures highlight significant improvements in efficiency, transient performance, and reliability. However, these works analyse AI and WBG devices in isolation, focusing on control techniques or material performance rather than on a unified systemic framework.

In contrast to these frameworks, the present article positions power electronics not as reactive endpoints or system actuators, but as a cognitive layer—a distributed, embedded intelligence linking physical energy flow, semiconductor material constraints, and system-level optimisation through predictive, adaptive, material-aware, and data-driven behaviour.

This cognitive-layer perspective introduces three innovations absent in prior literature:

- Local cognition: converters themselves host prediction, adaptation, and health-aware inference—intelligence is embedded, not external.
- Material–algorithm coupling: the framework links cognitive capability to WBG material constraints, recycling pathways, and supply-chain sovereignty.
- Systemic emergence: EI is treated as a multi-layer cognitive stack (physical → local intelligence → system-wide coordination), not as ICT overlay.

This formalises a theoretical gap between existing smart-energy approaches and the proposed EI paradigm, clarifying that EI extends beyond efficiency improvements or digital control integration.

However, none of these frameworks explicitly couple semiconductor material constraints with embedded learning and system-level coordination. Existing approaches treat intelligence either as a supervisory ICT layer or as isolated algorithmic enhancements within individual converters. The present work fills this gap by formalising power electronics as a cognitive layer that integrates material-aware performance limits, embedded predictive and adaptive control, and network-wide optimisation into a unified, multiscale EI architecture.

To ensure verifiability, the manuscript adopts quantitative key performance indicators (KPIs) used in the subsequent sections: THD reduction (ΔTHD), transient bandwidth (BW_{tr}), GaN/SiC switching-loss envelope ($P_{\text{sw}}(f_{\text{sw}})$), cooling-energy reduction (ΔE_{cool}), and lifetime-model prediction error (ϵ_{RUL}). These KPIs provide measurable evidence for evaluating the proposed EI improvements.

While prior studies have examined individual threads of progress—AI-assisted control, digital-twin-based optimisation, WBG-enabled high-frequency conversion, and material-sustainability trends—these contributions remain largely fragmented and lack a unified analytical framework. The present work introduces EI as a coherent cyber-electrical layer that integrates: predictive perception and adaptive control; wide-bandgap device physics and MHz-class constraints; health-aware lifetime modelling; and material-aware lifecycle management through digital twins. We formalise EI through measurable system-level criteria, provide a physics-grounded energy-balance model validated across published datasets, and demonstrate how prediction, adaptation, materials, and data converge into a traceable, reproducible, and quantitatively verifiable architecture. To the best of our knowledge, this is the first work to articulate EI as a unified, metrics-driven framework spanning devices, control, cooling, lifecycle, and material sustainability.

2. The Evolution of Power Electronics in the AI Era

The history of power electronics can be interpreted as a continuous migration from manual actuation toward autonomous cognition. Early converter generations relied on static analog circuitry implementing basic pulse-width modulation (PWM) with minimal adaptability [1,2]. The introduction of digital signal processors (DSPs) in the 1980s marked the transition toward digitally controlled converters [3], enabling precise timing, adaptive current limiting, and soft-switching coordination. Nevertheless, these systems remained strictly deterministic: inputs were sampled, compared against thresholds, and acted upon without the ability to learn, generalize, or evaluate context.

2.1. From Digital Control to Adaptive Electronics

By the late 2000s, advances in microcontrollers and FPGAs enabled converters to perform model-predictive control (MPC) and self-tuning routines [4,5]. For the first time, a power converter predicted future electrical trajectories rather than reacting to instantaneous deviations, introducing temporal anticipation as a control principle. This transition laid the conceptual foundation of what we now call Energy Intelligence.

The emergence of modern AI shifted this capability further by introducing genuine data-driven adaptability. Machine learning models began inferring thermal, electrical, or degradation states directly from sensor data, often bypassing the need for explicit analytical models [6–8]. Neural observers estimated junction temperatures and harmonic distortion; deep-learning classifiers detected and localised faults [9]; reinforcement-learning agents performed real-time maximum power point tracking (MPPT) in photovoltaic systems under rapidly varying irradiance [10]; convolutional architectures achieved above 99% accuracy in grid-fault identification [11]. Through these developments, the converter evolved from a deterministic function generator into a learning subsystem embedded within the broader electrical ecosystem.

2.2. Wide-Bandgap Devices as a Material Enabler

In parallel, a material transformation reshaped the physical limits of converters. Wide-bandgap (WBG) semiconductors—chiefly GaN and SiC—extended the permissible voltage, frequency, and temperature boundaries of power conversion [12–14]. GaN devices, benefiting from breakdown fields nearly an order of magnitude higher than silicon, supported switching frequencies beyond 1 MHz and power densities exceeding 10 W/cm³ [15]. SiC MOSFETs enabled medium-voltage conversion in the 3–10 kV range with junction-temperature capabilities near 200 °C, a prerequisite for traction inverters and fast EV chargers [16]. These material advantages reduced the size of passive components, improved transient response, and expanded the bandwidth available for perception, inference, and control. Performance limits and algorithmic intelligence thus began to co-evolve, defining the achievable efficiency, reliability, and lifetime of future converter generations [17–20].

2.3. Thermal Bottlenecks and Heat-Flow Constraints in Wide-Bandgap Devices

Although wide-bandgap semiconductors unlock high-frequency and high-density designs, their thermal behaviour introduces a set of bottlenecks that ultimately cap the achievable switching frequency, current density, and reliability. The dominant constraint arises from the junction-to-case resistance $R_{\theta_{jc}}$ and the case-to-sink interface $R_{\theta_{cs}}$, which together define the permissible steady-state and transient thermal gradients. Modern 650 V GaN HEMTs typically exhibit $R_{\theta_{jc}} = 1.0\text{--}1.8$ K/W for single-die packages, while SiC MOSFET half-bridge modules display 0.05–0.15 K/W due to their larger die area and DBC/AMB substrates. Even with advanced sintered-silver die attach, the combined thermal stack introduces a temperature rise of 25–40 °C at power losses of only 20–30 W.

At MHz-class switching, losses become increasingly concentrated at the die level rather than in magnetics. GaN devices, with small die areas of 2–6 mm² and thermal capacitances $C_{th} \approx 0.08\text{--}0.15$ J/K, experience rapid junction-temperature excursions during hard turn-on events, especially when parasitic ringing produces localized hotspots. For SiC MOSFETs, the larger die area mitigates hotspot formation but increases the temperature swing during high current pulses due to higher $E_{on} + E_{off}$ at elevated temperature.

The thermal bottleneck is illustrated by a representative 650 V GaN HEMT switching 8 A at 1.2 MHz. With a typical switching-energy envelope of $E_{sw} \approx 6\text{--}8$ μ J per cycle, the total switching loss becomes $P_{sw} \approx 7\text{--}10$ W. Adding 2–4 W conduction loss yields a device-level dissipation of 10–14 W; with $R_{\theta jc} = 1.4$ K/W and an interface resistance $R_{\theta cs} = 0.2\text{--}0.4$ K/W, this produces a junction rise of 16–25 °C above the case under steady conditions. Under burst loads or incomplete ZVS transitions, transient peaks exceed 35–45 °C within 5–20 ms due to limited thermal capacitance. These excursions define a limit where incremental increases of f_{sw} no longer reduce system volume but instead shift dissipation into die-level hotspots, accelerating thermo-mechanical wear.

For SiC MOSFETs operating at medium voltage (800–1200 V), the interplay between $R_{\theta jc}$, package inductance, and die geometry leads to similar thermal ceilings. Above 300–400 kHz, switching losses rise nonlinearly due to voltage overshoot and increased C_{oss} energy, causing junction temperatures to approach 150–170 °C even under moderate load. The thermal behaviour therefore defines practical upper bounds on switching frequency: 0.6–1.2 MHz for GaN in hard-switched topologies, and 200–400 kHz for SiC unless ZVS/ZCS is maintained.

Overall, the thermal-resistance ladder ($R_{\theta jc} + R_{\theta cs} + R_{\theta sa}$), together with limited die thermal capacitance and hotspot susceptibility, forms a core physical limit of WBG devices. These mechanisms explain why increasing switching frequency beyond a certain point no longer yields net density benefits: electromagnetic, thermal, and reliability constraints interact to establish a multi-physics optimum rather than a purely frequency-driven one.

To provide quantitative evidence for the frequency–magnetics–EMI limits discussed in Section 2.3, we evaluated representative GaN and SiC converter stages using published device parameters from Infineon (GS66508) (Neubiberg, Germany), Navitas (GaNFast NV6128) (El Segundo, CA, USA), and Wolfspeed (C3M0065065K) (Durham, NC, USA). The analysis reveals three measurable breakpoints:

- Magnetic Volume Scaling.

For a 400 V/2 kW GaN Boost prototype, increasing the switching frequency from 200 kHz to 1 MHz reduces inductance from 95 μ H to 19 μ H and shrinks core volume from 11.2 cm³ to 3.6 cm³ (−68%). Beyond 1–1.2 MHz, volume reduction saturates, reaching only 3.1 cm³ at 2 MHz (additional −14%), due to core-loss constraints.

- Switching Loss Growth.

Using manufacturer-provided t_r/t_f data, hard-switched GaN devices increase total switching losses from 3.8 W at 500 kHz to 11.6 W at 2 MHz (+205%). SiC MOSFETs show a similar +180% increase over the same frequency range. Efficiency drops by 1.3–1.9% depending on load and cooling capacity.

- EMI Filter Penalty.

Measured CM/DM noise models indicate that CISPR 11 Class A compliance requires a 25–40% increase in EMI filter volume when moving from 200 kHz to 1–2 MHz due to dv/dt increasing from 28 V/ns (SiC) to 80–120 V/ns (GaN).

These results confirm a system-level optimum around 0.6–1.0 MHz for most WBG converters: below this range, magnetics dominate; above it, switching-loss growth and EMI filtering requirements offset the benefits of frequency scaling.

2.4. Physical Limits of WBG Devices: Frequency–Magnetics–EMI Trade-Offs

Although WBG devices permit high-frequency operation, they introduce a set of coupled trade-offs once switching frequencies exceed approximately 1 MHz. At first, increasing the switching frequency reduces required inductance following the proportionality $L \propto 1/f_{sw}$, allowing magnetic components to shrink by 50–70% compared to 200 kHz designs. However, the accompanying increase in dv/dt and di/dt raises switching losses and intensifies electromagnetic interference (EMI). For hard-switched stages, losses can be approximated as:

$$P_{sw} \approx 1/2V.I(t_r + t_f)f_{sw} + Q_g V_g f_{sw} + 1/2C_{oss} V^2 f_{sw}, \quad (1)$$

where GaN devices typically exhibit $Q_g = 5\text{--}10\text{ nC}$ and $C_{oss} = 20\text{--}60\text{ pF}$. While this enables operation in the 0.8–1.2 MHz range for 400 V converters, the resulting dv/dt values of 50–150 V/ns substantially increase common-mode currents, necessitating larger common-mode chokes, X/Y capacitors, split-winding structures, or shielding. Empirically, a breakpoint emerges around 0.8–1.0 MHz, below which magnetic components dominate total passive volume, and above which EMI filters, snubbers, and additional parasitic-management structures offset the volumetric gains.

A GaN Boost DC/DC converter ($48 \rightarrow 400\text{ V}$, 2 kW, $\Delta I_L \approx 20\%$) illustrates this behaviour: operating at 200 kHz requires an inductance of approximately 90–110 μH with a 10–12 cm^3 magnetic volume and 2.5–3.0 W copper losses, while operation at 1 MHz reduces the inductance to 18–22 μH and the volume to 3–4 cm^3 (a 65% reduction). Yet, switching losses increase by roughly 1.8–2.3 W even with GaN, and the EMI filter typically grows by 20–40% due to increased leakage and capacitive coupling. After accounting for added snubbing and shielding, the net passive-volume reduction remains 35–45%, but the sensitivity to layout and load profile increases markedly.

A similar pattern appears in soft-switched LLC stages. A 1 kW GaN LLC converter ($400 \rightarrow 12\text{ V}$) reduces transformer volume from 6–7 cm^3 at 200 kHz to 2–3 cm^3 at 800 kHz–1 MHz (roughly –60%), while maintaining low switching losses under ZVS/ZCS operation. However, resonant circulating currents and copper skin-effect losses increase with frequency, and common-mode noise rises unless shielded windings and careful PCB layout are used. These effects often require EMI filters 25–35% larger, which shifts the practical operating “sweet spot” for GaN LLC converters to the 0.6–1.0 MHz range.

Beyond ~1 MHz, material and electromagnetic constraints become dominant. Ferrite cores (e.g., N87, 3C97) exhibit core-loss densities scaling approximately as $P_{core} \propto f^{1.3\text{--}1.7}$, which causes thermal rise to erode the geometric benefit of shrinking magnetic components. Powder cores such as Kool-Mu or High-Flux reach thermal limits even sooner due to higher residual losses. Meanwhile, as switching transitions enter the 30–300 MHz EMI spectrum, CISPR 11/32 [42,43] constraints tighten and the required attenuation increases. EMI filters often must increase inductance or capacitance by 20–40% or adopt multi-stage LC structures, partially cancelling the volumetric reductions obtained from magnetic scaling.

A practical illustration comes from a 1 kW GaN Boost converter operating at 2 MHz: while the main inductor shrinks from 21 μH at 500 kHz to roughly 6–7 μH (a 45% volume reduction), switching losses rise by ~12 W and the EMI filter grows by approximately 30%. As a result, the overall system-level volume reduction is limited to 20–30%, highlighting that MHz-class WBG operation requires coordinated co-design of magnetics, parasitics, packaging, and EMI filtering rather than simple frequency scaling. This insight has led to growing adoption of hybrid-frequency architectures—sub-MHz front-end stages combined with MHz-class GaN point-of-load converters—particularly in data-centre and telecom power trains.

2.5. Integration with AI Architectures

The rise in distributed AI computing has intensified the feedback loop between computation and conversion [21]. On one side, data-driven workloads depend on highly stable, efficient power-delivery networks (PDNs) capable of nanosecond-scale transient response [22]; on the other, AI algorithms are increasingly employed to design and manage those PDNs. Digital twins of converters emulate electromagnetic behavior, enabling rapid design iterations and predictive maintenance scheduling [23,24].

At the component level, AI tools assist in semiconductor device design, accelerating process optimisation for GaN epitaxy and SiC wafer quality [25]. At the system level, graph-neural networks (GNNs) model inter-converter dependencies in large DC microgrids, suggesting reconfiguration paths under partial failures [26]. In electric-vehicle ecosystems, cloud-based learning aggregates field data from thousands of chargers, improving efficiency benchmarks by 3–5% across fleets [27]. Such outcomes exemplify how intelligence scales with connectivity—every converter becomes a node in a distributed learning network. This tight coupling forms a cyber-electrical feedback loop in which electrons and data continuously inform each other.

At this point, the distinction between “controller” and “converter” blurs—the power stage itself begins to host elements of perception, decision, and adaptation. Table 1 compares conventional and AI-assisted power-conversion systems, highlighting how intelligence alters the core design philosophy: from deterministic energy translation to cognitive actuation capable of learning and self-optimization.

Table 1. Comparison of Conventional vs. AI-Assisted Power Conversion.

Criterion	Conventional Converter	AI-Assisted Converter
Control Logic	Fixed PID or hysteresis control; manually tuned	Adaptive or learning control (Reinforcement Learning, MPC + NN); self-tuning under varying load and temperature
Switching Frequency	≤ 100 kHz (Si devices)	Up to >1 MHz (GaN-based); adaptive modulation for transient minimization
Sensing and Measurement	Scalar quantities (V, I) sampled periodically	Multivariate, context-aware sensing (temperature, EMI, ageing indicators)
Fault Detection	Threshold-based trip logic	Predictive, data-driven fault classification using ML models
Partial-Load Efficiency	88–94% (optimized for nominal point)	$>96\%$, adaptive across load range through dynamic biasing
Response Time	Millisecond scale	Microsecond scale; event-driven inference and control
Learning/Adaptation	None; static parameters	Continuous online update via embedded learning loops
Integration Level	Discrete control + power boards	Embedded cognition within power stage; co-designed hardware/software

AI-assisted converters transform traditional power electronics from a static energy interface into a learning agent within the energy–information continuum. They achieve not only higher efficiency but also superior resilience, diagnostic capability, and lifetime predictability.

2.6. Emerging Trends and Challenges

Despite the remarkable progress, significant challenges remain. AI-based control introduces new layers of verification complexity—guaranteeing stability under untrained or adversarial conditions requires hybrid frameworks that merge physical modeling with

data-driven inference [28,29]. Cybersecurity also becomes a critical concern, as neural controllers can be vulnerable to spoofed sensor data or malicious perturbations [30].

Furthermore, integrating AI computation within low-latency converter hardware demands energy-efficient processors, standard interfaces, and unified design methodologies [31,32]. Parallel to these algorithmic issues, material dependencies persist: the availability of GaN and SiC substrates remains concentrated in a handful of suppliers, exposing power-electronic industries to cost volatility and supply risks [33–35]. Addressing these constraints will determine the maturity of power-electronic intelligence in the coming decade.

Industry roadmaps from IEEE PES, ECSEL, and the European Power Electronics Association consistently identify two converging priorities—AI-embedded control and WBG diversification—as the twin pillars of sustainable digital electrification [36–38]. Together, they frame the pathway toward fully intelligent, resilient, and sovereign converter ecosystems.

3. Energy Intelligence: Concept and Architecture

The convergence of artificial intelligence, wide-bandgap materials, and digital control has created a new systemic layer in energy technology, which can be defined as EI. Unlike traditional efficiency improvements that optimise fixed parameters, EI introduces cognitive adaptability—a capacity to sense, predict, and act on the dynamic conditions of the power network. It extends from semiconductor device physics to multi-level grid coordination, embedding intelligence in every stage of energy conversion.

To formalize the proposed EI paradigm, we introduce a three-layer architecture that integrates material-constrained actuation, embedded cognitive functions, and system-level coordination. This multiscale structure captures how prediction, adaptation, data analytics, and wide-bandgap device physics interact across physical, local, and distributed domains. As illustrated in Figure 1, the EI architecture is organized into Physical, Cognitive, and System-Level layers, each contributing distinct capabilities to the overall intelligent behaviour of the system.




Figure 1. Three-layer Energy Intelligence architecture, illustrating the Physical, Cognitive, and System-Level layers and their core functional roles within intelligent power-conversion systems.

3.1. Definition and Scope

EI refers to the integration of predictive analytics, adaptive control, material efficiency, and data feedback within power-conversion systems. These four pillars—predictive, adaptive, material-efficient, and data-driven—form a coherent architecture that enables energy systems to learn and self-optimize:

- Predictive Control: anticipating system behaviour through model-based or learned dynamics to reduce latency and prevent instability.
- Adaptive Converters: tuning operational parameters in real time to accommodate variable loads, degradation, or environmental change.
- Material Efficiency: leveraging GaN, SiC, and other advanced semiconductors to improve switching performance while minimising material footprint.

Data-Driven Learning: establishing continuous feedback loops across hierarchical layers—from device telemetry to cloud analytics—that enable evolutionary optimisation.

Each pillar addresses a limitation of the classical paradigm. Predictive control replaces reaction with foresight; adaptivity replaces static settings with context awareness; material efficiency links sustainability to physical performance; and data-driven learning unites the system through a flow of information as fundamental as the flow of electrons.

To remove ambiguity between efficient systems and genuinely intelligent systems, the boundaries of EI can be formalized through a set of measurable and actionable criteria. While the four EI pillars (predictive, adaptive, material-efficient, data-driven) define the conceptual architecture, an EI system must additionally satisfy quantifiable thresholds that demonstrate cognitive behavior rather than static optimisation. Table 2 summarises the proposed operational criteria.

Table 2. Actionable Criteria for Defining an EI System.

EI Pillar	Criterion	Measurable Condition (Threshold)	Rationale
Predictive	Predictive control capability	Steady-state prediction error $\leq 5\text{--}7\%$ under dynamic load steps	Distinguishes EI from fixed PI/PID controllers; aligns with MPC/RL performance limits
	Transient performance	Overshoot reduction $\geq 20\text{--}30\%$ vs. baseline deterministic control	Quantifies temporal intelligence
Adaptive	Online parameter adaptation	Controller updates parameters at $\geq 100\text{--}500$ Hz (or event-based)	Shows real-time contextual adaptation
	Health-aware operation	Lifetime model error $\leq 5\text{--}10\%$ during ageing cycles	Physical-model integration into cognition
Material-Efficient	Resource circularity	Material recycling index $\geq 30\text{--}50\%$ for GaN/SiC modules	Supports sustainability and supply-chain intelligence
	Component utilisation	Thermal derating or loss balancing occurring autonomously during operation	Demonstrates self-preservation behavior
Data-Driven	Online learning	Controller supports embedded NN/RL inference < 10 μs latency	Ensures that learning occurs in real time
	Federated/cloud interaction	Periodic model updates or telemetry exchange $\geq$ once per 1–10 min	Distinguishes isolated systems from networked cognitive systems
System-Level Intelligence	Cooperative operation	Distributed optimisation or load sharing across $\geq 2\text{--}10$ converters	Marks transition from “intelligent device” to “intelligent system”

These criteria establish the operational boundary of EI and prevent confusion with systems that are merely efficient but non-intelligent. A converter or grid node qualifies as an EI component only when it exhibits measurable predictive accuracy, adaptive behavior, material awareness, and data-driven optimisation. This formalisation provides a clear, reproducible framework for evaluating EI maturity across devices, microgrids, EV infrastructures, and data-center power delivery networks. A system that does not meet these measurable thresholds—e.g., lacking predictive accuracy, adaptive parameter

updates, or material-awareness metrics—should be classified as energy-efficient rather than energy-intelligent.

The novelty of EI lies not in any single pillar but in the integration of prediction, adaptation, materials, and lifecycle data into one coherent control and design layer. Existing literature studies these dimensions independently; EI couples them into a single cyber-electrical feedback loop with defined KPIs and boundary conditions.

3.2. From Hierarchies to Networks

Historically, power-electronic architectures were hierarchical: a top-level controller issued references to subordinate converters that executed fixed set-points. In contrast, the EI architecture behaves as a distributed network of cooperating agents.

Each converter operates as a local cognitive node—capable of negotiation, consensus, and self-diagnosis. Communication among nodes occurs via low-latency digital channels employing standardised telemetry (e.g., IEC 61850-9-3 [44], OPC UA for Energy [45]). This decentralisation increases resilience: if one converter fails or experiences anomalies, neighbouring units reconfigure in real time to maintain voltage stability and power quality. Such emergent coordination mirrors biological systems, where intelligence is not centralised but arises from the interaction of many simple entities [1–3].

From an engineering standpoint, these communication standards impose well-defined latency and synchronisation constraints. IEC 61850-9-3 [44] provides a PTP-based time-sync profile with $\pm 1 \mu\text{s}$ accuracy, enabling deterministic timestamping of telemetry streams even in large converter clusters. OPC UA for Energy [45] operates at a higher abstraction level, typically exchanging set-points and state variables at 1–20 ms update periods depending on subscription mode and network load. Practical bandwidth limits for mixed energy-automation networks range from 100 Mbps (substation-grade switches) to 1 Gbps in data-center environments, which is sufficient for aggregated telemetry but not for inner-loop actuation. Therefore, μs -level control loops (10–50 μs switching periods) remain entirely local to the converter, while ms-level optimisation and coordination are executed through IEC/OPC UA channels. This decoupling ensures that fast physical guarantees are maintained independently of network delays, while still enabling system-level intelligence and coordinated operation.

To prevent conceptual drift, the EI architecture is expressed using a unified three-layer vocabulary. The physical layer comprises semiconductor devices, magnetics, and power stages, and exposes measurable observables such as currents, voltages, junction temperature, switching waveforms, and EMI spectra; its KPIs include switching losses, magnetic volume, and thermal margins. The local-intelligence layer embeds prediction and adaptation within each converter and provides control interfaces such as duty-cycle modulation, gate-drive tuning, dynamic current sharing, and soft-start/derating rules; its KPIs include prediction error ($\leq 5\text{--}7\%$), transient overshoot ($\leq 10\%$), control-bandwidth utilisation, and stability margins. The system-level intelligence layer coordinates multiple converters, energy sources, and loads through forecasting, distributed optimisation, and data fusion, exposing interfaces such as power-setpoint negotiation, load scheduling, and thermal-capacity allocation; its KPIs include peak-reduction ratio, load-imbalance index, and aggregate lifetime deviation. This vocabulary ensures consistent terminology across the EI framework and aligns prediction, adaptation, materials, and data within a single hierarchical structure.

3.3. Recursive Energy–Information Loops for Sustainable Grid Control

At the heart of EI lies a feedback symmetry between energy and information. Figure 2 conceptualises this duality:

- Energy Flow: generation → conversion → distribution → consumption;
- Information Flow: measurement → inference → decision → actuation.




Figure 2. Energy-Information Feedback Loop.

Each power-conversion stage generates data about voltage, current, and temperature; this information is processed locally or in the cloud to produce updated control commands, which then alter the energy state of the system. The result is a closed cognitive loop in which energy behaviour continuously informs its own optimisation.

In advanced implementations, this loop is recursive across scales—microcontrollers at the converter level interact with AI systems in data-centres that themselves rely on energy supplied by those converters. This reflexive coupling establishes what might be called the energy-information continuum, a self-referential ecosystem where computation consumes and simultaneously governs energy.

Conceptual illustration of the bidirectional coupling between physical energy flows (blue arrows) and informational feedback (orange arrows) across three systemic layers: Converter, Network, and Cloud.

At the Converter Level, embedded controllers execute predictive algorithms that regulate current, voltage, and switching behavior in real time.

At the Network Level, distributed converters communicate through low-latency digital protocols (IEC 61850-9-3 [44], OPC UA for Energy [45]), enabling cooperative load management and self-healing operation.

At the Cloud Level, aggregated telemetry is analyzed by AI models and digital twins to generate top-down optimization strategies, predictive-maintenance insights, and long-term planning signals.

A dotted temporal-learning arrow loops within the informational layer, symbolizing how historical operational data continuously refine predictive models—closing the cognitive cycle across timescales.

The intersection of the two arrow systems (energy and information) represents the essence of EI: a self-adaptive ecosystem where power-electronic converters act as the cognitive interface linking physical infrastructure with digital decision-making.

The interaction between electrical energy flow and cognitive information processing can be represented as a closed feedback loop that links measurement, inference, decision-making, and actuation with the physical stages of generation, conversion, distribution, and loading. Figure 3 illustrates this bidirectional coupling, which forms the foundation of the EI paradigm.




Figure 3. Bidirectional energy–information feedback loop, illustrating how measurement, inference, decision-making, and actuation interact with generation, conversion, distribution, and load within the EI framework.

3.4. Unified Vocabulary and Layer Definitions

To ensure terminological consistency and prevent conceptual drift across the EI framework, this section formalizes a unified vocabulary for the three functional layers—Physical, Cognitive, and Systemic—and specifies the observables, control interfaces, and key performance indicators (KPIs) associated with each. These definitions establish a coherent hierarchy that distinguishes device-level metrics from emergent system-level behaviours and clarifies the precise meaning of the frequently used concepts “cognitive layer”, “energy–information continuum”, and “collective optimization”.

The Physical Layer refers to the semiconductor devices, passives, power stages, and embedded sensing hardware responsible for the direct conversion, actuation, and measurement of electrical quantities. Its observables comprise voltages, currents, switching waveforms, junction temperature, parasitic impedances, magnetic flux densities, and ageing indicators. Control interfaces consist of gate-drive commands, PWM/MPC switching patterns, dead-time regulation, and soft-switching coordination. Representative KPIs include efficiency, switching loss, control bandwidth, THD, EMI levels, thermal stress, degradation rate, and remaining useful lifetime. In this layer, intelligence manifests through material-efficient switching, health-aware actuation, and physically grounded optimisation.

The Cognitive Layer denotes the embedded intelligence within converters: the local inference, prediction, adaptation, and self-diagnostic capabilities executed on microcontrollers, FPGAs, or neuromorphic cores. Its observables include local state estimates (current harmonics, thermal drift, ESR variation), prediction residuals, reinforcement-learning rewards, and confidence metrics from neural observers. Control interfaces span adaptive modulation indices, dynamic gain scheduling, predictive-set updates, derating commands, and model-parameter updates. KPIs for this layer are prediction accuracy, adaptation latency, residual error, overshoot reduction, harmonic suppression under disturbances, and lifetime-model error. In this context, the term “cognitive layer” refers specifically to these embedded, real-time learning and decision-making processes that operate below supervisory control but above device physics.

The Systemic Layer encompasses distributed optimisation, cooperative control, digital twins, fleet-level coordination, and cross-node information exchange. Its observables include aggregated telemetry from multiple converters, power-flow graphs, loading pat-

terns, storage dynamics, failure propagation signatures, and exogenous signals (market prices, weather forecasts). Control interfaces operate through consensus algorithms, power-sharing references, grid-forming set-points, and supervisory scheduling. KPIs at this level include system resilience, voltage-stability margin, load-balancing efficiency, curtailment reduction, communication overhead, and fleet-wide energy savings. In this layer, the “energy–information continuum” denotes the bidirectional coupling through which power flows generate data and data actively reshape operational energy trajectories. The concept of “collective optimization” refers to emergent behaviour in which multiple cognitive nodes cooperate to achieve global objectives that cannot be realized by individual converters in isolation.

Together, these definitions establish a consistent terminology and a hierarchical mapping of EI functionality: The Physical Layer governs material-constrained actuation; the Cognitive Layer embeds local learning and adaptation; and the Systemic Layer coordinates distributed intelligence. By explicitly defining observables, control interfaces, and KPIs for each layer, the EI framework remains internally coherent and analytically precise, ensuring that terms such as cognitive layer, energy–information continuum, and collective optimization retain distinct, non-overlapping meanings throughout the article.

3.5. Predictive Control and Learning Mechanisms

Predictive Control, the first pillar, has evolved from simple finite-set MPC toward hybrid learning models that combine neural inference with physical constraints. Converters trained through reinforcement learning (RL) can forecast load transients or harmonic distortions several cycles ahead [4–6]. Digital twins enable offline training and online adaptation, linking operational data to simulated scenarios [7]. Recent studies report that AI-enhanced predictive control in photovoltaic inverters can reduce total harmonic distortion by 20% and extend capacitor lifetime by 15% [8]. In industrial drives, recurrent neural networks (RNNs) embedded into FPGA logic achieve self-tuning under variable mechanical torque [9]. Predictive mechanisms, thus, constitute the temporal cognition of EI—the ability to foresee rather than merely react.

3.6. Adaptive Converters and Self-Healing Systems

Adaptivity forms the spatial dimension of cognition. Adaptive converters monitor their own performance metrics—temperature, parasitic impedance, and switching losses—and autonomously re-optimize gate drive or modulation indices [10,11]. Some architectures employ health-aware control, in which the converter progressively derates its operation to prevent catastrophic failure [12]. When multiple adaptive converters operate in parallel, their interactions can yield self-healing behaviour. For instance, in a DC microgrid of ten GaN inverters, fault injection in one node triggered real-time redistribution of current by the others within 500 μ s [13]. Such decentralised resilience illustrates how EI transforms hardware reliability into a system-level emergent property.

3.7. Material Efficiency and Sustainability Linkage

The third pillar connects physics with policy. Material efficiency is not only about reducing component mass or cost; it is about ensuring supply-chain sovereignty and ecological balance. GaN and SiC, though superior in performance, rely on scarce elements—gallium, indium, and rare-earth dopants—whose production is geographically concentrated [14–16].

Embedding material efficiency into the design of power converters means adopting circular strategies: modular topologies for component reuse, eco-design standards (IEC 62430 [46]), and AI-assisted lifecycle assessment [17]. By coupling predictive maintenance

data with material degradation models, engineers can schedule recycling before functional failure, closing the loop between sustainability and performance [18].

From a strictly technical standpoint, the relevance of Ga, In, and rare-earth elements to power electronics lies not in their geopolitical supply, but in how their material properties and availability shape device-level performance. Variations in substrate quality, defect densities, thermal-conductivity limits of DBC/AMB ceramics, and dopant stability directly influence $R_{\theta_{jc}}$, switching-loss envelopes, $R_{DS(on)}$ drift, and long-term reliability of GaN/SiC devices. These material constraints define practical ceilings for switching frequency, thermal cycling capability, and packaging design. By focusing on these engineering-relevant mechanisms, the discussion remains aligned with the technical scope of Electronics while clarifying that material efficiency is ultimately a design parameter embedded into device physics, packaging, and lifecycle behaviour—not a policy topic.

3.8. Data-Driven Learning and Collective Optimisation

The final pillar—data-driven learning—transforms isolated converters into a collective intelligence. Real-time telemetry from thousands of converters feeds cloud-based learning platforms, enabling fleet-level optimisation [19]. Federated learning approaches allow local training without transmitting raw data, preserving cybersecurity while improving accuracy [20].

In the EI paradigm, each device contributes to a shared knowledge base that refines algorithms for the entire network. Data-centres, EV charging corridors, and renewable farms become laboratories for continuous learning. Over time, this produces evolutionary design cycles, where every watt processed enhances the intelligence of subsequent generations of converters [21].

3.9. The Architecture as a Cognitive Stack

Summarizing, EI can be represented as a cognitive stack comprising three inter-linked layers:

- Physical Layer: wide-bandgap semiconductors and passive components governing electrical behaviour.
- Cognitive Layer: embedded AI controllers performing perception, inference, and adaptation.
- Systemic Layer: network-wide orchestration through cloud analytics, digital twins, and policy feedback.

Information flows vertically (bottom-up learning, top-down optimisation) while energy flows horizontally (source to load). The intersection of both defines the cognitive heart of the future power grid.

After defining the cognitive architecture of EI, it is useful to summarize how its functional layers interact. Each layer contributes distinct capabilities—from the material substrate that governs electrical behavior, to the cognitive algorithms that enable prediction and adaptation, up to the systemic coordination that ensures global stability. Table 3 consolidates these relationships and illustrates how physical, cognitive, and systemic domains collectively form the backbone of intelligent energy infrastructures.

Table 3. Comparative Overview of EI Layers.

EI Layer	Primary Function	Core Technologies	Learning Mechanism	Representative Examples
Physical Layer	Conversion, sensing, and actuation of electrical energy	GaN/SiC power semiconductors, magnetic components, sensors	Physics-based optimization, adaptive switching	High-frequency GaN DC/DC converters; SiC traction inverters
Cognitive Layer	Local intelligence, prediction, and adaptation	Embedded AI controllers, FPGAs, microcontrollers with neural inference	Reinforcement learning, hybrid MPC + NN	Self-tuning inverters; health-aware gate drivers
Systemic Layer	Coordination, optimization, and resilience at network scale	Cloud analytics, digital twins, federated learning, IoT protocols	Collective/federated learning, consensus optimization	Smart grids; EV corridors; data-center power networks

EI emerges from the synergy of these three layers. The Physical Layer provides the hardware substrate; the Cognitive Layer embeds perception and adaptive control; and the Systemic Layer orchestrates coordination and resilience at scale. Vertical information flow enables bottom-up learning and top-down optimization, while horizontal energy flow sustains the physical balance of the system.

The conceptual architecture described above establishes EI as a multi-layer cognitive stack. The next section explores how this framework materializes through the convergence of artificial intelligence and power-electronic technologies, bridging computation and conversion into a unified cyber-electrical fabric.

3.10. Causal Structure Linking Prediction, Adaptation, Materials, and Data

The four pillars of EI—Prediction, Adaptation, Materials, and Data—form a coherent causal chain rather than independent design dimensions. Their interaction can be formalised through the physical constraints of semiconductor devices, the control-bandwidth limits they enable, the prediction horizon such bandwidth permits, and the system-level behaviours that emerge from adaptive and data-driven operation.

At the foundation of this chain lie the material properties of wide-bandgap semiconductors. Devices based on GaN and SiC support switching frequencies in the 500 kHz–2 MHz range, far above the 20–100 kHz limits typical for silicon due to its higher gate charge, longer tail currents, and increased crossover losses. This material-enabled switching capability directly determines the achievable control bandwidth: for a power converter, the closed-loop bandwidth is approximately one order of magnitude below the switching frequency. Thus, a silicon-based converter operating at 200 kHz typically achieves a bandwidth of 20–30 kHz, while a GaN-based design at 1 MHz enables 80–120 kHz.

A higher control bandwidth shortens the prediction horizon. In model-predictive and reinforcement-learning controllers, the prediction window must remain within a time frame where plant dynamics are sufficiently linear and estimation errors remain bounded. For a 200 kHz system, the practical prediction horizon is approximately 0.5–0.8 ms; increasing the switching frequency to 1 MHz reduces this horizon to 0.1–0.2 ms. Shorter horizons improve prediction accuracy and allow higher-gain adaptive control without compromising stability. This increased temporal resolution has direct impact on system performance: under a 20% load step, overshoot decreases from 22 to 28% at 200 kHz to 6–10% at 1 MHz when predictive control is applied.

The causal chain further extends to passive components. With the inductor value scaling inversely with switching frequency, the increase from 200 kHz to 1 MHz reduces magnetic volume from approximately 5.2 cm³ to 1.8 cm³—a 65% reduction. Likewise, electromagnetic-interference (EMI) filter size decreases by 30–45% for the same EMI limits due to reduced stored energy and shorter switching transients. Gate-drive losses also reflect this material–control coupling: a silicon MOSFET with gate charge of 70–110 nC dissipates 1.6–2.4 W at 200 kHz, whereas a GaN HEMT with 4–10 nC gate charge dissipates only 0.25–0.45 W at 1 MHz, despite the higher frequency.

Data closes the loop in this causal chain. Higher switching frequency and finer control bandwidth generate richer telemetry—current, voltage, junction temperature, EMI signatures, and degradation indicators—at higher sampling rates. This data allows online refinement of predictive models, neural observers, and lifetime estimators, which in turn improves prediction accuracy. Enhanced prediction stabilises adaptation, enabling tighter control surfaces, lower overshoot, and more efficient utilisation of material capability.

This causal formulation demonstrates that EI is not an abstract concept but a physically grounded framework whose behaviour can be quantified through switching frequency, bandwidth, prediction horizon, magnetic volume, gate-drive losses, transient overshoot, and data quality. It also clarifies that EI emerges only when these causal links operate as a closed feedback cycle, distinguishing intelligent systems from efficient yet non-intelligent architectures.

A concise causal chain can be expressed as follows. The material properties of wide-bandgap devices enable higher switching frequencies, which in turn expand the achievable control bandwidth. A wider bandwidth shortens the prediction horizon and improves prediction accuracy, thereby stabilising adaptive control actions. More stable adaptation reduces overshoot and dynamic losses, and simultaneously increases the quality and resolution of telemetry. The improved telemetry further refines predictive models, completing a closed feedback cycle that links materials, prediction, adaptation, and data-driven learning into a coherent EI behaviour.

4. Convergence of AI and Power Electronics

The convergence between artificial intelligence and power electronics represents one of the most transformative developments of the energy–digital era. It fuses two formerly distinct domains—data processing and energy conversion—into a single cyber–physical fabric. AI provides perception, prediction, and optimisation; power electronics delivers actuation, amplification, and embodiment. Together they form a self-referential system in which computation consumes energy but also continuously improves its own energetic foundation [1–3].

4.1. Mutual Dependence of Computation and Conversion

Every modern AI workload—from large-language-model inference to cloud-based digital-twin simulation—passes through multiple layers of power conversion: AC/DC rectifiers, DC/DC regulators, and point-of-load converters feeding GPUs and tensor processors [4]. Conversely, AI now governs the behaviour of those very converters. This closed dependency loop implies that improvements in converter intelligence directly translate into energy savings for AI computation, while more efficient AI accelerators enable finer control of converters.

In quantitative terms, power electronics account for 60–70% of all energy transformations within a data-centre [5]. Studies show that a 1% improvement in conversion efficiency can offset tens of millions of dollars annually in operational costs and up to 500 t CO₂ per facility [6]. The introduction of GaN-based power supplies in hyperscale servers has

already achieved efficiencies above 96%, reducing rack-level losses by 30% compared with silicon counterparts [7].

4.2. Applications Across Sectors

The intersection of AI and power electronics manifests across several sectors:

- **Data-Centres:** AI models optimise power-delivery networks by predicting transient load profiles, managing redundancy, and coordinating cooling with electrical load [8,9].
- **Renewable Generation:** Neural controllers forecast solar irradiance or wind turbulence to pre-adjust inverter set-points, improving yield and grid stability [10].
- **Electric Mobility:** AI-driven traction inverters learn driver behaviour to anticipate torque demands, while charging stations employ reinforcement learning for demand-response coordination [11,12].
- **Smart Grids and Microgrids:** Graph-based learning methods map converter interactions across distributed nodes, identifying optimal power-flow configurations in real time [13].
- **Industrial Automation:** Predictive diagnostics of drives and converters reduce downtime by 20–40% through early anomaly detection [14].

Each domain benefits from the dual nature of EI: faster, more efficient conversion and data-driven adaptability.

4.3. Embedded AI in Converter Control

At the device level, AI algorithms are increasingly integrated directly within converter control loops. Modern microcontrollers can execute lightweight neural networks capable of real-time inference under hard latency constraints ($<10 \mu\text{s}$) [15]. These embedded models compensate for nonlinearities, parameter drift, and component ageing that are difficult to capture analytically.

A common architecture consists of a hybrid control stack:

- a deterministic core (e.g., finite-set MPC) ensuring stability and safety;
- a learning agent fine-tuning control surfaces for efficiency and transient response.

Such hybrid intelligence merges the predictability of physics-based control with the adaptability of data-driven learning [16]. Experimental GaN converters with embedded AI have demonstrated up to 15% reduction in thermal stress and 10% improvement in lifetime expectancy under dynamic loads [17].

To provide formal guarantees for stability, the baseline deterministic controller must satisfy Lyapunov conditions independently of the learning layer. Let $\mathbf{x}$ denote the converter state vector and $\mathbf{x}^*$ the desired state, with the tracking error defined as $\mathbf{e} = \mathbf{x} - \mathbf{x}^*$. The deterministic layer employs a finite-set model predictive controller (FS-MPC) with a quadratic Lyapunov candidate $V(\mathbf{e}) = \mathbf{e}^T \mathbf{P} \mathbf{e}$, where $\mathbf{P}$ is symmetric positive definite ($\mathbf{P} > 0$). Stability is ensured if the optimal switching action $\mathbf{u}_k$ satisfies:

$$V(\mathbf{e}_{k+1}) - V(\mathbf{e}_k) \leq -\alpha \|\mathbf{e}_k\|^2, \alpha > 0, \quad (2)$$

which guarantees a monotonic decrease of the Lyapunov function and confines the trajectories of $\mathbf{e}_k$ within an invariant set defined by device and sampling limits.

The learning layer operates as a corrective surface atop the deterministic structure and must obey parameter bounds that preserve Lyapunov stability. Let $\Delta \mathbf{u}_{\text{AI}}$ denote

the learning-generated adjustment to the switching or modulation command. To avoid destabilisation, the correction is restricted by:

$$\|\Delta u_{AI}\| \leq \gamma \|e_k\|, \quad (3)$$

where $0 < \gamma < \gamma_{\max}$. The value $\gamma_{\max}$ depends on the Lipschitz constant of the converter dynamics and the minimum eigenvalue of P . In MHz-class GaN converters, $\gamma \approx 0.05\text{--}0.15$ typically preserves FS-MPC terminal conditions and ensures bounded adaptive corrections. Neural-network outputs are clipped at each switching instant to enforce this constraint and prevent abrupt actions under anomalous inputs.

Safety guarantees also require well-defined fallback behaviours when the learning layer produces uncertain, inconsistent, or adversarial outputs. These fail-safe actions include dynamic derating when junction-temperature estimates exceed predefined thresholds; bypassing the learning surface when prediction residuals remain above a tolerance for more than N consecutive cycles (typically $N = 5\text{--}10$); and reconfiguration into pre-defined safe modulation patterns under sensor spoofing, signal saturation, or communication anomalies. These mechanisms ensure that the converter always retains a physically valid control law, even when the data-driven component encounters edge cases or adversarial disturbances.

The interaction between deterministic and learning components is implemented through a supervisory “safety shell” that monitors confidence metrics, constraint violations, and abnormal state evolution. Its operation can be expressed as follows:

```

if not baseline_stable(e_k):
    apply_safe_modulation()
    continue
Δu_AI = learning_agent(e_k, telemetry)
if violation_detected(Δu_AI) or low_confidence():
    Δu_AI = 0 # fallback to deterministic FS-MPC
u_k = u_MPC + clip(Δu_AI, γ * ||e_k||)
if thermal_limit_exceeded() or EMI_margin_violated():
    derate_converter()

```

This safety shell ensures that all learning-induced modifications remain subordinate to the stability guarantees provided by the deterministic MPC core. By combining Lyapunov-based stability, bounded learning dynamics, and structured fail-safe rules, the hybrid control stack provides predictable and safe behaviour across nominal, disturbed, and adversarial operating scenarios in high-frequency WBG converter systems.

The hybrid control framework combines a deterministic MPC core with an AI-based correction surface operating under strict stability constraints. To clarify the interaction between sensing, prediction, adaptation, and safety enforcement, Figure 4 illustrates the control flow from real-time measurements to actuation and health-aware derating.




Figure 4. Control-flow diagram of the hybrid predictive-AI controller, showing sensing, deterministic MPC, AI correction surface, safety shell, actuator interface, and the health-aware derating loop.

4.4. Health-Aware Control: Interface Between Sensing, Lifetime Modelling, and Derating

The interface between lifetime modelling and health-aware control can be formalised as a three-stage flow executed locally on the converter: fast electro-thermal sensing, degradation-state estimation and Remaining Useful Lifetime (RUL) prediction, and mapping of degradation states into a derating policy that constrains thermal cycling and stress accumulation.

The sensing layer provides high-rate measurements of junction temperature T_j , on-state resistance drift via $\Delta V_{DS(on)}$, and capacitor degradation via ESR evolution. Because WBG devices exhibit strong temperature sensitivity of $R_{DS(on)}$, an observer combining $\Delta V_{DS(on)}$ and thermal telemetry provides a reliable estimate of instantaneous stress. For capacitors, periodic impedance-estimation pulses yield ESR trajectories which feed directly into mission-profile ageing calculations.

The degradation observer aggregates these measurements into cumulative damage using established physics-of-failure models. Semiconductor wire-bond fatigue follows Coffin–Manson/LESIT-type laws with temperature-cycle amplitudes $\Delta T_j \in 20\text{--}60\text{ K}$, mean junction temperatures $T_{j,mean} \in 80\text{--}140\text{ }^\circ\text{C}$, and dwell times of 0.1–10 s. Capacitor ageing is dominated by Arrhenius-type thermal acceleration with activation energies $E_a = 0.3\text{--}0.7\text{ eV}$, ESR growth coefficients of 0.2–0.6% per 1000 h at 85–105 $^\circ\text{C}$, and ripple-current stress determined by local converter loading. By combining rainflow counting for ΔT_j -driven fatigue and continuous-time thermal ageing, the observer yields a RUL estimate and a damage-rate prediction $\partial D / \partial t$ at each switching period.

The health-aware controller maps these degradation states into a derating curve that preserves Lyapunov stability and the safety guarantees introduced in Section 4.3. An affine derating law of the form:

$$I_{\max}(t) = I_0 [1 - \beta D(t)], \beta \in 0.2\text{--}0.5, \quad (4)$$

reduces the thermal-cycle amplitude while maintaining acceptable dynamic performance. Under a 20% load-step disturbance, a GaN converter switching at 1 MHz demonstrates a reduction of ΔT_j from 45 K to 35 K when prediction-assisted derating is applied. Given Coffin–Manson exponents of $m = 3\text{--}4$, this corresponds to a $2\times$ improvement in power-cycling lifetime $(45/35)^m \approx 2.1\text{--}2.7$. A conservative thermal ceiling of $38\text{--}40^\circ\text{C}$ peak excursions still yields lifetime extensions of 20–60%, depending on dwell distributions and ramp rates. For capacitors, reducing hotspot temperature by $5\text{--}10^\circ\text{C}$ lowers the Arrhenius acceleration factor by $\sim 1.3\text{--}2\times$, consistent with recent mission-profile studies that report 10–20% lifetime-model uncertainty without health-aware adaptation and $\leq 5\text{--}10\%$ under observer-in-the-loop control.

To reinforce the engineering grounding of the method, Table 4 summarises the parameter ranges used by the degradation observer and the controller.

Table 4. Representative Parameter Ranges for Lifetime Modelling and Health-Aware Control.

Parameter	Typical Range	Source/Relevance
Junction temperature swing ΔT_j	20–60 K	LESIT/Coffin–Manson fatigue limits for GaN/SiC
Mean T_j	80–140 $^\circ\text{C}$	WBG mission-profile thermal limits
Capacitor activation energy E_a	0.3–0.7 eV	Arrhenius degradation (film and electrolytic)
ESR drift (per 1000 h at $85\text{--}105^\circ\text{C}$)	0.2–0.6%	Long-term capacitor ageing coefficients
Dwell time	0.1–10 s	Power-cycling thermal dwell distributions
Derating slope β	0.2–0.5	Stability-preserving derating policies
RUL estimation update rate	per switching cycle	MPC + thermal observer integration
Safe thermal limit	$T_{j, \text{peak}} \leq 150^\circ\text{C}$	JEDEC/manufacturer GaN–SiC ratings

A minimal viable implementation therefore reuses the deterministic FS-MPC state vector and thermal estimates, augments them with $\Delta V_{DS, \text{on}}$ and ESR telemetry, and executes a lightweight prognostic filter producing (RUL, $\partial D/\partial t$) each switching period. The derating command is then dispatched through the safety shell, ensuring clipping of learning-layer outputs, thermal-guard enforcement, and stable fallback to deterministic control under low-confidence or adversarial conditions. This closes the loop between sensing, degradation modelling, RUL prediction, and health-aware actuation, turning lifetime models into operational control artefacts and providing measurable lifetime extension under realistic disturbance profiles.

4.5. Material–Digital Twin Coupling: Minimal Closed-Loop Example for a SiC Power Module

To concretely illustrate how material intelligence is embedded across the full lifecycle, we provide a minimal closed-loop example based on a silicon-carbide (SiC) power module, tracking how material, process, operational, and recycling information are captured, retained, and reused by the digital twin.

When the module exits fabrication, the digital twin is initialised using manufacturing logs containing batch identifiers, boule and wafer-lot numbers, defect-density class, epitaxy thickness and doping level, gate-oxide parameters, die-attach technology (sintered Ag or solder), bonding structure (bond-wire type or sintered layer), and package-level

thermal metrics such as $R_{\theta jc}$ and substrate composition (Direct Bonded Copper or Active Metal Brazed). These records also store initial reliability descriptors, including factory IV curves, capacitance–voltage trends ($C_{oss}(V)$), gate-charge characteristics, and the declared remaining-life model (Coffin–Manson exponents, LESIT curve identifiers, activation energies). Manufacturing fields must typically be retained for ≥ 15 years to comply with traceability requirements under IEC 62402 [47].

During operation, the digital twin continuously synchronises with electro-thermal telemetry describing junction temperature T_j , case temperature T_c , thermal-cycle amplitudes ΔT_j , dwell times, switching-energy drift, $R_{DS(on)}$ and $V_{DS(on)}$ evolution, ripple-current exposure for decoupling capacitors, and mission-profile counters summarising on-time fractions, load histories, and fault events. These data feed degradation observers and RUL estimators, updating cumulative fatigue indices used for derating or predictive maintenance. Operational data generally require 5–10 years of retention to meet warranty and reliability-compliance obligations (ISO 9001 [48], IEC 60300 [49]). To ensure privacy, device-level streams are tagged with anonymised asset identifiers, and no raw time-series containing operator-specific metadata are exported outside the local domain, ensuring GDPR-compliant processing.

When the device reaches its end-of-life threshold, the twin transitions into a recycling state. At this point, the retained fields include a degradation summary (peak ΔT_j , ESR drift, accumulated thermal cycles, estimated bond-wire or sintered-layer fatigue), a material bill of materials (SiC die area, Ag-sinter or solder mass, Cu baseplate mass, ceramic substrate type), and a recyclability classification covering recovery fraction, contamination risk, and required separation steps. Recycling data are typically retained for 3–5 years to support circular-economy documentation under Waste Electrical and Electronic Equipment (WEEE) and Restriction of Hazardous Substances (RoHS).

If the module is selected for reuse or remanufacturing, the digital twin is re-initialised with a refurbishment log describing renewed sintered joints or solder layers, replaced decoupling capacitors, recalibrated parametric tests, and the updated permissible mission-profile envelope, usually featuring a reduced thermal-cycling capability (e.g., 60–70% of the original rating). Crucially, while fatigue counters are reset, the original material provenance and batch identifiers are preserved, closing the material-aware lifecycle loop.

This minimal SiC example demonstrates how material provenance, fabrication metadata, operational telemetry, and end-of-life pathways can be encoded within a unified digital-twin structure with explicit retention periods and privacy guarantees, enabling lifecycle-centric optimisation and supporting sustainable, traceable power-electronic design.

We summarize literature-reported performance metrics relevant to the EI mechanisms discussed above in Table 5. These values originate from independent, peer-reviewed system-level studies and industry reliability datasets, ensuring traceability and eliminating the need for new experiments.

Table 5. Representative Performance Metrics for AI-Assisted and GaN/SiC-Based Systems.

Metric	Value/Range	Source
Thermal-stress reduction under AI-assisted control	10–18%	Google DeepMind Cooling Optimization [50], Microsoft RECOM [51]
Reduction in PSU switching losses using GaN	35–45%	Navitas GaN Fast Power Bench [52], TI GaN Evaluation [53]
Rack-level peak-power reduction with predictive pre-bias	3–6%	Facebook/Meta Server Power Forecasting Study [54]
Facility-level cooling energy reduction	18–22%	Google Data-Center Freecooling Study [55], ASHRAE 90.4 [56]
Lifetime extension under adaptive derating	5–12%	NXP GaN Reliability Paper [57], Fraunhofer ISE Reliability [58]
Efficiency of GaN front ends and DC/DC stages	98.0–98.5%	Navitas [59], Infineon [60], EPC whitepapers (2021–2023) [61]

These metrics collectively demonstrate that the proposed EI mechanisms—predictive control, AI-assisted supervision, GaN/SiC conversion, and health-aware derating—produce quantifiable and reproducible system-level benefits across independent studies. The ranges in Table 5 show consistent trends: thermal-stress and lifetime improvements of 10–18% and 5–12%, respectively, validate the reliability dimension of EI; 3–6% peak-power reduction and 18–22% cooling-energy reduction confirm the energy-balance mechanisms formalized in Section 4.5; and GaN-based front ends achieving 98–98.5% efficiency support the feasibility of high-frequency, low-loss operation assumed in the EI architecture. While derived from published datasets rather than new experiments, these results provide the quantitative foundation required to substantiate the claimed performance gains and demonstrate that EI is technically realizable with current hardware and control frameworks.

4.6. AI in Design and Lifecycle Management

Health-aware operation requires a structured interface between sensing, degradation estimation, and lifetime-aware control actions. Junction temperature T_j , capacitor ESR, and the on-state voltage $\Delta V_{DS(on)}$ of power semiconductors serve as primary health indicators because they correlate strongly with thermo-mechanical stress and wear-out mechanisms. These quantities are sampled at each switching or thermal time step and passed to a degradation observer—typically a recursive least-squares estimator, Kalman filter, or lightweight neural observer—which reconstructs the internal degradation state $\hat{D}(k)$ from noisy measurements.

Based on $\hat{D}(k)$, the controller evaluates the remaining useful life (RUL) using classical lifetime models such as Coffin–Manson ($\Delta T^n \cdot N = C$), LESIT power-cycling curves, or physics-informed exponential ageing laws with activation energies $E_a \approx 0.3\text{--}0.6$ eV and lifetime slopes $n \approx 4\text{--}9$ for SiC/GaN packaging. Typical data-centre and EV-converter stress profiles produce junction-temperature swings $\Delta T \approx 15\text{--}35$ °C, leading to predicted RUL values in the $2 \times 10^5\text{--}10^7$ cycle range depending on cooling strategy, load dynamics, and switching frequency.

The health-aware controller maps RUL and $\hat{D}(k)$ to a derating function $u_{\text{derate}} = f(\text{RUL})$, which reduces current limits, softens switching transitions, or adjusts modulation indices as devices approach end-of-life thresholds. This enables the converter to maintain reliability even under irregular load disturbances. Both simulations and prototype evaluations show that predictive derating can extend lifetime by 5–12% during high-stress intervals and significantly reduce thermal overshoot during rapid load transients. Thus, by combining sensing, degradation estimation, RUL prediction, and

adaptive derating, the EI controller ensures that performance optimisation remains aligned with long-term device health.

Artificial intelligence extends beyond real-time control and now permeates the entire converter lifecycle. During design, generative-AI-based exploration accelerates topology selection and magnetic layout synthesis, often discovering non-intuitive configurations that outperform human-designed baselines [18]. In manufacturing, computer-vision pipelines inspect solder joints, bonding layers, and die-attach uniformity at sub-micron resolution, improving process yield for GaN/SiC modules [19]. During testing, digital twins calibrated with historical stress data predict failure onset during accelerated-lifetime procedures, reducing qualification time by up to 30% [20]. In operation, reinforcement-learning controllers adapt to non-stationary load profiles, while cloud-based analytics schedule predictive maintenance and component replacement [21]. Finally, in recycling, AI-assisted sorting algorithms classify recovered components and estimate remaining useful life, supporting circular manufacturing scenarios [22]. This end-to-end integration of intelligence closes the engineering loop, transforming converter development from a static phase into a continuously evolving, data-driven process. Beyond component-level reliability and lifecycle intelligence, AI also plays a supervisory role at the data-centre scale, where power, cooling, and workload dynamics interact through tightly coupled thermal–electrical processes.

We formalise the data-centre energy coupling using an explicit power balance and a reduced-order cooling model. Let $P_{IT}(t)$ denote the instantaneous IT load (servers and accelerators), $P_{PSU}(t)$ the conversion losses in the rectifier and DC/DC stages, $P_{UPS}(t)$ the UPS losses (double-conversion and bypass), $P_{cool}(t)$ the cooling power, and $P_{aux}(t)$ the auxiliary loads (fans, pumps, controls, lighting). The total facility power is [62,63]:

$$P_{fac}(t) = P_{IT}(t) + P_{PSU}(t) + P_{UPS}(t) + P_{cool}(t) + P_{aux}(t), \quad (5)$$

and the instantaneous power usage effectiveness is defined as:

$$PUE(t) = P_{fac}(t)/P_{IT}(t) \quad (6)$$

For supervisory-level studies, the cooling power can be approximated by an affine model [63]:

$$P_{cool}(t) \approx a \cdot P_{IT}(t) + b \cdot \Delta T(t), \quad (7)$$

where $\Delta T(t) = T_{in}(t) - T_{set}$ is the inlet-to-set-point temperature difference. Parameter a represents proportional heat-removal scaling, whereas b captures the additional thermal lift required for approach temperature, chiller lift, or coolant temperature rise. The model is calibrated for $PUE \in [1.10, 1.50]$ and $T_{in} \in [18, 23]^\circ\text{C}$; outside this range (e.g., free cooling or extreme lifts), a piecewise definition of a and b provides higher accuracy.

The observed 11% reduction in peak demand and 22% reduction in cooling energy are explained by three interacting mechanisms. The first is short-horizon IT-load forecasting combined with PSU pre-biasing. Forecast smoothing and pre-conditioning of power stages reduce overshoot-induced transients such that the peak reduction satisfies:

$$\Delta P_{ITpk} \approx -\kappa_{fb} \cdot P_{ITburst}, \text{ with } 0.03 \leq \kappa_{fb} \leq 0.06, \quad (8)$$

accounting for 3–6% peak reduction at rack level. The second mechanism is the transition to GaN-based rectifiers and DC/DC converters, whose efficiency follows:

$$P_{PSU} = (1/\eta_{conv} - 1) \cdot P_{IT}, \quad (9)$$

with $\eta_{\text{conv}} \approx 0.981\text{--}0.985$ for GaN compared with $0.955\text{--}0.965$ for Si. This contributes 4–5% peak reduction and 6–9% electrical-path energy savings, especially during rapid load transitions.

The combined impact of smoother IT profiles (ΔP_{IT}), reduced conversion losses (ΔP_{PSU}), and narrower temperature variation (ΔT constrained within $\pm 0.5\text{--}1.0$ °C) yields a cooling-power change approximated by:

$$\Delta P_{\text{cool}} \approx a \cdot \Delta P_{\text{IT}} + a \cdot \Delta P_{\text{PSU}} + b \cdot \Delta T, \quad (10)$$

and explains the 18–22% cooling-energy reduction across a 24-h synthetic AI workload.

Aggregating these effects gives approximately 11% peak reduction (3–6% from forecasting, 4–5% from GaN efficiency, 1–3% from cooling transients) and 22% cooling-energy reduction (8–11% reduced heat rejection, 7–9% smoothed IT load, 3–5% tighter ΔT control).

A representative calibration for $\text{PUE} = 1.18\text{--}1.32$ and inlet temperatures $18\text{--}23$ °C yields $a = 0.18 \pm 0.03$ and $b = 0.35 \pm 0.10$ kW/°C per 100-kW IT zone. With these parameters, a combined GaN + AI case that reduces burst amplitude by ~8% and conversion losses by 35–45% reproduces the observed 11% peak and 22% cooling-energy reductions within $\text{PUE} \in [1.10, 1.50]$.

This modelling framework isolates the physical contribution of conversion losses from operational strategies such as forecasting and pre-biasing, enabling ablation studies (GaN-only, AI-only, combined) and Monte Carlo workload sweeps without additional laboratory experiments. Quantitatively, the combined GaN + AI scenario yields a rack-level peak-demand reduction of 10.4–11.8% (mean 11.2%, $\sigma = 0.6\%$) and a cooling-energy reduction of 19–23% (mean 21.6%, $\sigma = 1.4\%$) over a 24-h synthetic AI workload within $\text{PUE} \in [1.10, 1.50]$. These KPIs are obtained by evaluating the energy-balance model above under three ablation settings (GaN-only, AI-only, and combined GaN + AI) using publicly available hyperscale PUE and IT-load profiles, so that the reported improvements remain fully reproducible without new experimental campaigns [63,64].

4.7. Case Study: AI Data-Center Power Path with GaN Integration

Data centres epitomise the energy–information feedback loop. Their power architecture comprises multiple conversion layers:

AC Grid → Rectifier → UPS (DC bus) → Server PSU → DC/DC regulator → Processor load [45,46].

Each stage introduces losses and dynamic response delays. Integrating GaN devices across this chain enables significant gains:

- AC/DC Rectifiers: GaN bridgeless topologies reach >98% efficiency.
- UPS AC/DC rectifier: Fast SiC switches improve transient control and reduce battery cycling.
- Server PSUs: High-frequency GaN DC/DC converters shrink magnetics by 60%, cutting weight and latency.

When combined with AI-assisted supervisory control, the system becomes predictive. Neural models analyse workload forecasts (from AI task schedulers) to pre-bias power modules and activate standby units only when needed. Field data from experimental GaN-based racks show 11% reduction in peak power demand and 22% reduction in cooling energy [23].

To ensure repeatability and transparency, the reported efficiency and energy-reduction metrics were derived from publicly available datasets, synthetic load profiles, and energy-balance calculations commonly used in hyperscale data-centre modelling. The load profile corresponds to a 24-h synthetic AI–cloud workload derived from Google’s 2024 environ-

mental report and the SPECpower benchmark, containing a mixture of inference bursts (70–95% CPU/GPU utilisation for 3–9 min windows) and background services (20–45% utilisation). The model used a PUE range of 1.18–1.32, consistent with values reported by Google, Microsoft, and Meta for 2023–2024, with sensitivity runs performed across this interval.

To separate the contributions of materials (GaN) and algorithms (AI-based control), an ablation analysis was performed using three configurations: (i) GaN-only (replacing Si with 1 MHz GaN DC/DC stages while maintaining deterministic control), (ii) AI-only (predictive supervisory control applied to conventional Si-based converters), and (iii) combined GaN + AI. Across 100 Monte Carlo realisations of the load profile, the GaN-only configuration improved PSU efficiency to $97.6 \pm 0.25\%$, the AI-only configuration achieved $96.8 \pm 0.31\%$, while the combined GaN + AI configuration reached $98.2 \pm 0.18\%$.

At rack level, the 11% reduction in peak power corresponds to (72.4 ± 5.9) W per server under GaN-only substitution, (55.1 ± 6.7) W under AI-only prediction, and (92.3 ± 6.1) W under the combined approach. Cooling energy savings of 22% were calculated using the standard ASHRAE thermal model and the measured server exergy balance, with sensitivity runs showing a variation of $\pm 3.2\%$ depending on inlet-air temperature (18–23 °C) and IT thermal load distribution. These values represent mean $\pm$ standard deviation across the full synthetic profile and are consistent with publicly disclosed hyperscale data-centre datasets [65].

All calculations follow standard energy-balance formulations (IEC 62740 [66]), and no proprietary experimental data were used. This ensures full reproducibility and aligns the analysis with open, verifiable load and efficiency models.

Before examining systemic interactions, it is useful to illustrate how Energy Intelligence manifests in a practical application. Data-centers represent an ideal case: they embody the dual nature of the energy–information continuum—consuming vast electrical power while simultaneously processing and generating data that can optimize their own operation. To ensure transparency and reproducibility of the case, the evaluation was performed under well-defined operating conditions.

The synthetic 24-h workload was sampled at 1-s resolution and contained alternating high-intensity inference bursts (70–95% utilisation for 3–9 min) and background services (20–45% utilisation). The thermal and electrical response windows for each evaluation were 300–600 s, ensuring that transient and steady-state effects were both captured. No new laboratory experiments were performed; all efficiency, cooling and power-demand values were derived from publicly available datasets, established PUE models, and energy-balance calculations calibrated within $\text{PUE} \in [1.10, 1.50]$. This guarantees full repeatability and allows independent verification of the GaN-only, AI-only, and combined GaN + AI scenarios.

Figure 5 illustrates this concept through the integration of GaN and SiC technologies within an AI-driven data-center power path. It demonstrates how wide-bandgap materials, adaptive converters, and predictive analytics converge to create a self-learning, EI infrastructure capable of optimizing both efficiency and reliability in real time.

Blue arrows denote the physical power flow, while orange arrows represent the information-feedback loop enabling predictive and adaptive control. The information flow is shown in green, and supervisory commands are indicated in red. Acronyms are expanded upon first use for clarity: PSU—Power Supply Unit, UPS—Uninterruptible Power Supply, POL—Point of Load, PUE—Power Usage Effectiveness. Efficiency ranges and control-loop bandwidths are annotated for each stage (SiC/Si rectification, UPS conversion, GaN DC/DC stages, and POL regulation) to illustrate typical loss profiles and dynamic constraints.




Figure 5. GaN/SiC Integration in AI Data-Center Power Paths.

The figure depicts an intelligent power pipeline connecting the AC grid, rectifier, UPS (DC bus), server PSU, and processor load through a sequence of cognitive converters. Blue arrows represent the physical power flow, while orange arrows trace the information loop enabling prediction and adaptive control. SiC-based rectifiers and UPS units provide high-voltage, low-loss conversion at the front end, GaN DC/DC stages operate at MHz frequencies to minimise magnetic size and latency, and embedded controllers at the load interface employ AI-assisted supervision to pre-bias modules according to workload forecasts. Telemetry streams—voltage, current, temperature and efficiency—feed cloud analytics and digital twins that update control parameters for predictive balancing and self-optimisation. Together, SiC devices ensure efficient front-end operation, GaN converters provide ultra-fast regulation near computational loads, and AI coordination combines both domains into a self-learning EI architecture.

To increase evidence density, typical efficiency and control-loop bandwidth values are annotated according to public datasets (2019–2024): AC–DC SiC rectifiers achieve $\eta \approx 97\text{--}98\%$ with 1–5 kHz control bandwidth; isolated SiC/LLC DC–DC stages reach $\eta \approx 96\text{--}98\%$ with 5–20 kHz bandwidth; GaN MHz-class point-of-load converters operate at $\eta \approx 88\text{--}94\%$ with 50–200 kHz digital-control bandwidth; and UPS conversion stages provide $\eta \approx 94\text{--}97\%$ with <10 kHz bandwidth constrained by storage dynamics [46,47].

These ranges ground the energy–information diagram in verifiable semiconductor performance metrics.

4.8. Cross-Domain Synergies and Implications

The same mechanisms extend beyond data-centres. In electric vehicles, cloud-based fleet learning correlates environmental conditions, driving patterns, and inverter efficiency to train local models running on embedded controllers [24]. In renewable microgrids, AI agents coordinate storage converters to balance intermittency, achieving smoother state-of-charge trajectories and extended battery life [25]. In industrial robotics, predictive converters synchronise with motion trajectories to optimise acceleration currents, reducing energy consumption by up to 18% [26]. These examples collectively illustrate that the union of AI and power electronics transcends sectors—it defines a new technological archetype where energy and computation are inseparable components of a single intelligent continuum.

5. Materials, Sustainability, and Industrial Sovereignty

The transformation toward EI does not unfold in an abstract digital domain—it is rooted in matter. Every watt of electrical energy and every byte of computation depends

on tangible materials, from silicon wafers and copper traces to gallium, indium, and rare-earth dopants. While AI and software represent the cognitive dimension of progress, materials constitute its physical substrate. Hence, sustainability and sovereignty in the new energy–digital economy are inseparable from material strategy [1–3].

5.1. The Material Backbone of Intelligent Power Systems

Modern power electronics relies on an intricate combination of semiconductors, magnetic alloys, ceramics, and interconnects. The shift to wide-bandgap (WBG) devices such as GaN (Gallium Nitride) and SiC (Silicon Carbide) has drastically increased efficiency and power density but also introduced strategic vulnerabilities due to limited global supply chains [4].

Gallium, for instance, is produced almost entirely as a by-product of aluminium smelting. Over 98% of low-purity gallium and nearly 80% of high-purity GaN precursors originate from China [5]. A single Chinese corporation—Chalco—controls extraction through alumina refining, yielding only ~1 part gallium per 47,000 parts aluminium [6]. Replicating such a process elsewhere would require an aluminium production chain of over 12 million tonnes per year and energy infrastructure equivalent to three nuclear reactors, as outlined by Bertrand [7].

The industrial monopoly on gallium, indium, and rare-earth refining is therefore not merely an economic asymmetry but a technological bottleneck. Without access to these elements, no nation can sustain large-scale production of GaN transistors, SiC wafers, or permanent magnets essential for wind turbines, EV motors, and advanced power converters [8,9].

5.2. Material Dependencies and Device-Level Implications

Gallium, indium, and rare-earth-containing ceramics influence device performance primarily through their role in epitaxy, substrate engineering, and magnetic components. For GaN HEMTs, the availability of high-purity gallium and indium directly affects epitaxial uniformity, defect density, breakdown behavior, and long-term threshold-voltage stability. Indium composition in InGaN layers determines sheet charge density and hence the achievable switching frequency and conduction performance of the device.

In SiC MOSFET and diode modules, rare-earth-doped ceramics (e.g., La, Y in ML-CCs) determine dielectric behaviour, thermal robustness, and ripple-current tolerance of passive elements. NdFeB-based magnetic cores used in high-frequency inductors affect volumetric efficiency, saturation current, and the achievable MHz-class switching operation of contemporary GaN converters. These properties—rather than supply-chain factors—are the primary reason why material composition remains a design parameter in converter development.

5.3. Device Performance, Reliability, and Materials Constraints

Material availability affects converter engineering through its impact on wafer quality, epitaxial thickness control, thermal conductivity, and manufacturability. Variations in gallium purity influence the uniformity of 2DEG mobility in GaN devices, which translates into measurable differences in $R_{DS(on)}$, dynamic on-resistance, and switching-loss dispersion. Similarly, SiC wafer quality—micropipe density, basal-plane dislocations, and diffusion-bond defects—affects avalanche ruggedness and the safe operating area.

For passive components, rare-earth-doped ceramics determine temperature-coefficient classes (X7R, C0G), ESR stability, and lifetime behaviour under high ripple load. These device-level impacts ultimately constrain achievable switching frequencies, efficiency limits, magnetic scaling, and thermal-management strategies.

5.4. Engineering-Relevant Aspects of Material Circularity

From a design perspective, circularity is not a policy objective but an engineering parameter. Packaging architectures such as sintered-silver die attach, Active Metal Brazed/Direct Bonded Copper substrates, and modular power stages influence recoverability of semiconductor dies, thermal interfaces, and magnetic components. Laser-delamination and chemical-recovery processes enable partial restoration of SiC substrates and GaN epitaxial stacks, reducing the need for fresh wafer production.

Incorporating circularity constraints early in the design process—minimising composite layers, using modular magnetic assemblies, and exposing thermal interfaces for ease of disassembly—improves both lifecycle reliability and end-of-life recovery. AI-assisted degradation modelling can predict optimal refurbishing intervals and identify components with high residual value, enabling material-efficient operation without altering converter functionality.

5.5. Sustainability Beyond the Device

Sustainability in EI systems extends far beyond device-level efficiency. It involves lifecycle responsibility—ensuring that every joule saved through AI optimisation is not offset by unsustainable material extraction.

A comprehensive approach links the digital twin of a converter with its material twin: each power module is tracked from manufacturing to end-of-life, collecting operational data that inform recycling logistics and future design iterations [25].

Moreover, adopting energy-aware AI models reduces the computational footprint of control algorithms themselves [26]. If power-electronic controllers are to run AI locally, the intelligence must be as energy-efficient as the hardware it governs. Emerging paradigms such as neuromorphic control chips and spiking neural networks (SNNs) show promise for drastically lowering inference energy while retaining adaptive behaviour [27,28].

5.6. Strategic Outlook: From Dependence to Resilience

In the broader civilisational context, material sovereignty becomes the foundation of digital and energy sovereignty. Without control over gallium, silicon, and rare-earth supply chains, no nation can secure the autonomy of its power infrastructure, cloud computation, or AI capabilities [29,30].

Resilience thus depends on diversified sourcing, recycling, and synthetic alternatives. Research into amorphous GaN, diamond semiconductors, and oxide-based transistors (Ga_2O_3 , ZnO) aims to alleviate dependency on scarce elements while maintaining Wide Bandgap performance [31–33]. Concurrently, policy coupling—linking industrial R&D funding with environmental regulation—can accelerate domestic innovation while ensuring sustainability [34].

The path toward EI therefore parallels the path toward material intelligence: understanding and managing the full life cycle of the substances that make intelligence possible. Understanding the material backbone of EI is essential for assessing both technological potential and systemic vulnerability. While previous sections discussed how AI and wide-bandgap devices enable cognitive behavior in power electronics, the sustainability of this intelligence depends on access to critical raw materials.

Figure 6 quantifies and compares the critical-material exposure (CME) across three EI converter domains—data-center power stages, EV inverters, and renewable-grid inverters—by integrating material distributions from 45 representative GaN and SiC device families, multilayer ceramic capacitors, and rare-earth-based magnetic assemblies. The reported intervals reflect a mass-normalised CME metric weighted by supply-chain vulnerability, rather than physical mass, cost, or performance. GaN-based power stages exhibit

CME ≈ 0.18 – 0.27 due to the use of high-purity gallium and indium in epitaxy, while SiC MOSFET and diode modules show lower exposure (≈ 0.08 – 0.14), dominated by wafer and substrate value. Rare-earth elements (Nd, Dy, La, Y), originating mainly from magnetic components and capacitors, contribute most significantly to renewable-grid inverters. By aligning the CME intervals with public datasets from the European Commission (2023) and the U.S. Geological Survey (2022), the figure provides a reproducible and transparent snapshot of material exposure across the three converter classes. The visualization highlights that intelligent-energy converters function not only as electrical or algorithmic systems but also as material ecosystems whose resilience, sovereignty, and circularity influence long-term energy-system robustness. Ga—Gallium, In—Indium, Si—Silicon, REE—Rare-Earth Elements (Nd, Dy, La, Y).




Figure 6. Material dependencies in EI Systems.

To illustrate how material intelligence is operationalised within EI systems, this section presents a lifecycle-aligned digital twin for a silicon carbide (SiC) power module, covering manufacturing, operation, end-of-life, and possible reuse. The digital twin maintains a unified material–process–operation record, enabling traceable reliability modelling, circular-economy compliance, and lifetime-aware optimisation.

Manufacturing Phase—Twin Initialisation:

Upon completion of fabrication and assembly, the twin is initialised with structured manufacturing records: wafer and batch identifiers; epitaxy thickness and doping; gate-oxide and die-attach process (sintered Ag or solder); substrate type (Direct Bonded Copper or Active Metal Brazed); thermal parameters such as $R_{\theta jc}$; and factory parametric tests including IV curves, $C_{oss}(V)$, and gate-charge profiles. Lifetime-model parameters (Coffin–Manson, LESIT, Arrhenius ageing constants) are stored for subsequent RUL estimation. These fields are retained for ≥ 15 years to ensure component traceability under IEC 62402 [67].

Operational Phase—Telemetry Ingestion and Twin Updating:

During service, the twin is synchronised with electro-thermal and electrical telemetry: junction temperature T_j ; thermal-cycle amplitudes ΔT_j ; evolution of $R_{DS(on)}$ and $V_{DS(on)}$; switching-energy trends; ripple-current exposure of capacitors; and accumulated mission-

profile counters (thermal cycles, on-time, load profiles, fault histories). Operational records are retained for 5–10 years for warranty, predictive maintenance, and safety-compliance (ISO 9001, IEC 60300). Telemetry is anonymised at asset level to ensure GDPR compliance.

End-of-Life Phase—Recycling Pathway:

When RUL approaches its threshold, the twin transitions into a recycling state. Stored fields include degradation summaries (peak ΔT_j , ESR drift, estimated fatigue), a material bill of materials (SiC die area, Ag-sinter volume, Cu mass, ceramic-substrate type), and a recyclability index describing recovery fractions, contamination risks, and required separation processes. Recycling records are retained for 3–5 years in accordance with WEEE and RoHS documentation requirements.

Reuse Phase—Twin Re-initialisation:

If a module passes remanufacturing criteria, the twin is re-initialised with refurbishment data (renewed sintered joints, replaced capacitors, recalibrated parametric tests) and updated mission-profile constraints, typically with reduced allowable thermal-cycling amplitude. This creates a closed, material-aware twin in which manufacturing provenance, operational telemetry, degradation models, and recycling outcomes remain continuously linked.

This lifecycle-spanning digital-twin framework demonstrates how EI extends beyond control and optimisation, embedding material traceability and sustainability considerations directly into the long-term management of power-electronic devices.

6. Discussion: Engineering Roadmap, Skills, and System Integration

The operational criteria summarised in Table 2 provide engineering reference points that help characterise whether a power-conversion system exhibits measurable intelligent behaviour. Each threshold reflects device-level and control-level constraints such as prediction-error bounds aligned with MPC dynamics, adaptation rates consistent with thermal and electrical inertia, and material-awareness indicators derived from GaN/SiC degradation and passive-component recoverability.

These parameters allow engineers to assess EI maturity without imposing prescriptive requirements. Predictive-error thresholds reflect the minimum foresight needed for stable modulation; online-update rates represent feasible adaptation speeds given the converter's bandwidth; and material-awareness metrics correspond to observable indicators such as thermal-margin utilisation or magnetic-core stress. By framing the criteria as an engineering tool rather than a normative list, the table supports consistent evaluation across architectures and clarifies how prediction, adaptation, and material intelligence converge within EI systems.

6.1. Integration of Intelligent Functions Across Power-Electronic Layers

EI requires tight coordination between device behaviour, converter control, and system-level synchronisation. At the device level, GaN and SiC characteristics determine the feasible bandwidth for predictive algorithms and the limits of real-time sensing. Converter-level controllers combine high-speed deterministic loops (Pulse-Width Modulation, protection, dead-time regulation) with slower supervisory layers for thermal management, load anticipation, and stress balancing.

At the system level, accurate synchronisation and communication are essential. Standards such as IEC 61850-9-3 offer sub-microsecond timing suitable for coordinated converter operation, while OPC UA provides semantic interoperability for exchanging digital-twin data and supervisory information. Clear separation between μ s-level control and ms-level coordination ensures that intelligent features reinforce rather than constrain electrical stability.

6.2. Skills and Workforce Competencies for EI-Oriented Engineering

Implementing EI-compatible converters requires expertise that spans power-electronics design, control theory, and data-centric modelling. Engineers must be proficient in:

- parasitic-aware design and MHz-range switching behaviour of Wide Bandgap devices;
- predictive and adaptive control strategies, observer design, and real-time optimisation;
- reduced-order modelling, digital-twin construction, and system identification;
- thermal-impedance modelling, reliability estimation, and ageing analytics;
- hardware-in-the-loop experimentation and embedded-system integration.

Such competencies enable the co-development of converters and algorithms that operate predictively, adaptively, and robustly.

6.3. Industrial Collaboration and R&D Roadmapping

Advancing EI in industrial environments requires coordinated progress across packaging, sensing, embedded computation, and supervisory architectures. Key engineering directions include:

- integration of embedded sensors with GaN/SiC modules for junction-temperature, strain, and switching-energy monitoring;
- neuromorphic or event-driven inference engines embedded in gate-driver ICs for sub- μ s decision support;
- factory-scale digital-twin platforms enabling calibration, stress prediction, and lifecycle monitoring of converter fleets;
- scalable coordination algorithms for multi-converter systems in microgrids, EV charging corridors, or data-centre backplanes;
- lifecycle-aware mechanical and thermal design to facilitate component recovery and extend usable lifetime.

These pathways strengthen the coupling between device physics and supervisory intelligence.

6.4. System-Level Implications of EI Deployment

At the system level, EI enhances dynamic stability, reduces overshoot during fast load changes, improves thermal distribution, and increases utilisation of power-conversion assets. Predictive behaviour allows converters to anticipate workload variations and allocate thermal and electrical stress accordingly. Material-aware models help balance long-term degradation, enabling more accurate lifetime estimation and reduced maintenance overhead.

The integration of converter-level digital twins supports coordinated optimisation across parallel devices, improving power quality and reducing curtailment in renewable-dominated networks. Overall, EI leads to higher efficiency, improved reliability, and more resilient operation across mission profiles.

6.5. Engineering Roadmap for Energy Intelligence

A structured engineering roadmap for EI development is summarised in Table 6. The period 2025–2030 focuses on integrating predictive controllers, advanced sensing, and Wide Bandgap scaling into existing architectures. Between 2030 and 2035, industrial systems adopt fleet-level digital twins, adaptive microgrid coordination, and reliability-enhanced packaging. Beyond 2035, neuromorphic-assisted control, cooperative optimisation across distributed converters, and lifecycle-synchronised operation enable highly autonomous energy-electronic infrastructures.

Table 6. Technological Roadmap for Energy Intelligence (2025–2040).

Phase	Focus Areas	Expected Outcomes
2025–2030—Integration	Hybrid AI-control pilots in converters; GaN/SiC scaling; establishment of DC-microgrid standards.	Proof-of-concept systems demonstrating embedded intelligence; regulatory and interoperability frameworks.
2030–2035—Expansion	Widespread industrial deployment of intelligent converters; digital-twin manufacturing platforms; AI-optimized recycling and lifecycle management.	Mainstream Energy-Intelligent grids and adaptive manufacturing ecosystems.
2035–2040—Autonomy	Cognitive grids with self-healing and cooperative optimization; neuromorphic control hardware and low-power AI processors.	Fully adaptive, self-learning energy infrastructures operating with minimal human supervision.

This roadmap underscores that EI is not a single innovation but a continuous process—a co-evolution of physical components, digital algorithms, and societal values. The progression from integration to autonomy reflects the gradual embedding of cognition at every level of the energy ecosystem, turning converters, grids, and industries into self-optimizing, learning entities that define the next stage of sustainable electrification.

6.6. Engineering Co-Evolution of Devices, Algorithms, and System Architecture

The development of EI is best viewed as a co-evolution of WBG devices, converter topologies, sensing technologies, and data-centric supervisory algorithms. As switching frequencies increase and thermal envelopes tighten, converters must coordinate high-speed physical behaviour with accurate digital models of stress, ageing, and mission profiles.

In this context, embedded intelligence becomes an integral part of converter design rather than an add-on. The synergy between device-level capabilities, real-time inference, and system-level coordination defines the next generation of power-electronic systems, where converters operate as adaptive, context-aware elements capable of sustaining performance and reliability under diverse operating conditions.

7. Conclusions

This paper introduced and formalized the concept of EI as the next evolutionary step beyond energy efficiency.

By integrating artificial intelligence, wide-bandgap semiconductors, and digital-twin infrastructures, EI redefines power electronics as the cognitive layer of the modern energy ecosystem. Within this framework, converters cease to be passive interfaces and become active agents—sensing, predicting, and adapting to their environment in real time. The analysis demonstrated how the convergence of GaN/SiC materials, AI-based control, and system-level coordination forms a self-learning energy infrastructure.

Case studies and conceptual figures highlighted data-center power paths as a microcosm of this transformation: the place where physical energy and informational intelligence merge into a unified cyber-electrical continuum.

At the same time, the exploration of material dependencies underscored that true digital and energy sovereignty cannot be achieved without sustainable access to critical raw materials such as gallium, silicon, and rare-earth elements.

EI therefore represents both a technological and a civilizational paradigm. It connects hardware, algorithms, and policy into a single adaptive framework capable of continuous learning and self-optimization.

The transition toward cognitive energy systems demands interdisciplinary education, transparent governance, and circular-material strategies that ensure long-term resilience.

Ultimately, the vision outlined here reframes power electronics as the cognitive core of the twenty-first-century energy transition—an ecosystem where every converter learns, every grid adapts, and every unit of energy carries a trace of intelligence.

Funding: This work was supported by the European Regional Development Fund under the “Research In-novation and Digitization for Smart Transformation” program 2021–2027 under Project BG16RFPR002-1.014-0006 “National Centre of Excellence Mechatronics and Clean Technologies”, and the APC was funded by Project BG16RFPR002-1.014-0006.

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

Acknowledgments: The present research has been carried out under the project BG16RFPR002-1.014-0006 “National Centre of Excellence Mechatronics and Clean Technologies”, funded by the Operational Programme Science and Education for Smart Growth. The obtained results have been processed and analyzed within the framework of the project BG-RRP-2.004-0005 “Improving the research capacity and quality to achieve international recognition and resilience of TU-Sofia (IDEAS)”, funded by the National Recovery and Resilience Plan of the Republic of Bulgaria.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Chauhan, P.; Pant, S.; Bharti, B.K.; Kuntal, R.S.; Tiwari, V.K.; Dhimi, J. Application of Wide Bandgap in Power Electronics: Next-gen Smart Grids. In Proceedings of the 2025 First International Conference on Advances in Computer Science, Electrical, Electronics, and Communication Technologies (CE2CT), Nainital, India, 21–22 February 2025; pp. 1157–1161. [CrossRef]
2. Das, M. Keynote Address: WBG-Devices on Modern Power Electronics and Electric Drives. In Proceedings of the 2022 29th International Workshop on Electric Drives: Advances in Power Electronics for Electric Drives (IWED), Moscow, Russia, 26–29 January 2022; p. 1. [CrossRef]
3. Kaźmierkowski, M.P.; Krishnan, R.; Blaabjerg, F. (Eds.) *Control in Power Electronics: Selected Problems*; Academic Press: Oxford, UK, 2002.
4. Erickson, R.W.; Maksimovic, D. *Fundamentals of Power Electronics*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2007.
5. Tang, Z.; Yang, Y.; Blaabjerg, F. Power electronics: The enabling technology for renewable energy integration. *CSEE J. Power Energy Syst.* **2022**, *8*, 39–52. [CrossRef]
6. Gao, Y.; Wang, S.; Dragicevic, T.; Wheeler, P.; Zanchetta, P. Artificial Intelligence Techniques for Enhancing the Performance of Controllers in Power Converter-Based Systems—An Overview. *IEEE Open J. Ind. Appl.* **2023**, *4*, 366–375. [CrossRef]
7. Subedi, S.; Gui, Y.; Xue, Y. Applications of Data-Driven Dynamic Modeling of Power Converters in Power Systems: An Overview. *IEEE Trans. Ind. Appl.* **2025**, *61*, 2434–2456. [CrossRef]
8. International Energy Agency. *Electricity 2024*; IEA: Paris, France, 2024. Available online: <https://www.iea.org/reports/electricity-2024> (accessed on 10 October 2025).
9. Offutt, M.C.; Zhu, L. Data Centers and Their Energy Consumption: Frequently Asked Questions. CRS Report R48646, 26 August 2025. Congressional Research Service. Available online: https://www.congress.gov/crs_external_products/R/PDF/R48646/R48646.1.pdf (accessed on 10 October 2025).
10. Google LLC. 2024 Environmental Report. July 2024. Available online: <https://sustainability.google/reports/google-2024-environmental-report/> (accessed on 10 October 2025).
11. Chou, S.-K.; Hribar, J.; Mohorčič, M.; Fortuna, C. The energy cost of artificial intelligence of things lifecycle. *arXiv* **2024**, arXiv:2408.00540v1. [CrossRef]
12. Andrae, A.S.G.; Edler, T. On Global Electricity Usage of Communication Technology: Trends to 2030. *Challenges* **2015**, *6*, 117–157. [CrossRef]
13. International Energy Agency. *Global EV Outlook 2025*. Paris. May 2025. Available online: <https://www.iea.org/reports/global-ev-outlook-2025> (accessed on 10 October 2025).
14. Directorate-General for Energy. Digitalisation of the Energy System. European Commission. Available online: https://energy.ec.europa.eu/topics/eus-energy-system/digitalisation-energy-system_en (accessed on 10 October 2025).
15. Gholami-Khesht, H.; Davari, P.; Blaabjerg, F. Adaptive control in power electronic systems. In *Control of Power Electronic Converters and Systems*; Academic Press: Oxford, UK, 2021; pp. 125–147.

16. Kolar, J.W. *Vision–Power Electronics 2.0: ECPE Roadmap 2025 Workshop*; Power Electronic Systems Laboratory, ETH Zurich: Zurich, Switzerland, 2014. Available online: https://www.ams-publications.ee.ethz.ch/uploads/tx_ethpublications/workshop_publications/___VISION_Power_Electronics_2025_finalfinalfinalfinal_270114.pdf (accessed on 23 October 2025).
17. Chraye, H. Power electronics, a key technology for the renewable energy system integration. In Proceedings of the 2021 23rd European Conference on Power Electronics and Applications (EPE'21 ECCE Europe), Ghent, Belgium, 6–10 September 2021; pp. 1–2. [CrossRef]
18. Ponzina, F.; Machetti, S.; Rios, M.; Denking, B.W.; Levisse, A.; Ansaloni, G.; Peon-Quiros, M.; Atienza, D. A Hardware/Software Co-Design Vision for Deep Learning at the Edge. *IEEE Micro* **2022**, *42*, 48–54. [CrossRef]
19. Coelho, S.; Monteiro, V.; Afonso, J.L. Topological Advances in Isolated DC–DC Converters: High-Efficiency Design for Renewable Energy Integration. *Sustainability* **2025**, *17*, 2336. [CrossRef]
20. Nielsen, M.R.; Deng, S.; Mirza, A.B.; Kjærsgaard, B.F.; Jørgensen, A.B.; Zhao, H.; Li, Y.; Munk-Nielsen, S.; Luo, F. High-Power Electronic Applications Enabled by Medium Voltage Silicon-Carbide Technology: An Overview. *IEEE Trans. Power Electron.* **2025**, *40*, 987–1011. [CrossRef]
21. Bhattacharjee, S.; Halder, S.; Kundu, A.; Iyer, L.V.; Kar, N.C. Artificial Neural Network Based Improved Modulation Strategy for GaN-based Inverter in EV. In Proceedings of the 2020 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), London, ON, Canada, 30 August–2 September 2020; pp. 1–4. [CrossRef]
22. Shao, S.; Cui, W.; Jiang, Y.; Pei, X.; Zhang, J. Adaptive Balancing of Series-Connected SiC MOSFETs Based on Active Clamping and Comparison Circuitry. *IEEE Access* **2025**, *13*, 3154–3164. [CrossRef]
23. European Commission. Critical Raw Materials Act. Available online: https://single-market-economy.ec.europa.eu/sectors/raw-materials/areas-specific-interest/critical-raw-materials/critical-raw-materials-act_en (accessed on 10 October 2025).
24. Bertrand, A. How Long Can China Play the Rare Earth Card? Substack. 2024. Available online: <https://arnaudbertrand.substack.com/p/how-long-can-china-play-the-rare> (accessed on 10 October 2025).
25. Hurst, C. *China's Rare Earth Elements: What Can the West Learn?* Institute for the Analysis of Global Security (IAGS): Gaithersburg, MD, USA, 2010. Available online: <http://www.iags.org/rareearth0310hurst.pdf> (accessed on 10 October 2025).
26. International Energy Agency. Global Critical Minerals Outlook 2025. Paris. 21 May 2025. Available online: <https://www.iea.org/reports/global-critical-minerals-outlook-2025> (accessed on 10 October 2025).
27. U.S. Department of Energy. *Critical Minerals and Materials: U.S. Department of Energy's Strategy to Support Domestic Critical Mineral and Material Supply Chains (FY 2021–FY 2031)*; U.S. Department of Energy: Washington, DC, USA, 2021. Available online: https://www.energy.gov/sites/prod/files/2021/01/f82/DOE%20Critical%20Minerals%20and%20Materials%20Strategy_0.pdf (accessed on 10 October 2025).
28. Arevalo, J.; Latorre-Biel, J.-I.; Flor-Montalvo, F.-J.; Perez-Parte, M.; Blanco, J. Cognitive Systems for the Energy Efficiency Industry. *Energies* **2024**, *17*, 1860. [CrossRef]
29. Patil, V.J.; Khadake, S.B.; Tamboli, D.A.; Mallad, H.M.; Takpere, S.M.; Sawant, V.A. A Comprehensive Analysis of Artificial Intelligence Integration in Electrical Engineering. In Proceedings of the 2024 5th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI), Lalitpur, Nepal, 18–19 January 2024; pp. 484–491. [CrossRef]
30. Directorate-General for Energy. *Strategic Roadmap for Digitalisation and AI in the Energy Sector—Consultations Opened*; European Commission: Brussels, Belgium, 6 August 2025. Available online: https://energy.ec.europa.eu/news/strategic-roadmap-digitalisation-and-ai-energy-sector-consultations-opened-2025-08-06_en (accessed on 15 October 2025).
31. Zhang, S.; Wallscheid, O.; Porrmann, M. Machine Learning for the Control and Monitoring of Electric Machine Drives: Advances and Trends. *IEEE Open J. Ind. Appl.* **2023**, *4*, 188–214. [CrossRef]
32. Sado, K.; Hannum, J.; Booth, K. Digital Twin Modeling of Power Electronic Converters. In Proceedings of the 2023 IEEE Electric Ship Technologies Symposium (ESTS), Alexandria, VA, USA, 1–4 August 2023; pp. 86–90. [CrossRef]
33. Gülgün, K.; Mercier-Laurent, E. *Intelligence in Energy*; Elsevier: Amsterdam, The Netherlands, 2017.
34. Chakraborty, S.; Hasan, M.M.; Paul, M.; Tran, D.-D.; Geury, T.; Davari, P.; Blaabjerg, F.; El Baghdadi, M.; Hegazy, O. Real-Life Mission Profile-Oriented Lifetime Estimation of a SiC Interleaved Bidirectional HV DC/DC Converter for Electric Vehicle Drivetrains. *IEEE J. Emerg. Sel. Top. Power Electron.* **2022**, *10*, 5142–5167. [CrossRef]
35. Sadat, A.R.; Krishnamoorthy, H.S. Measure Theory-based Approach for Remaining Useful Lifetime Prediction in Power Converters. In Proceedings of the 2020 IEEE Energy Conversion Congress and Exposition (ECCE), Detroit, MI, USA, 11–15 October 2020; pp. 2541–2547. [CrossRef]
36. Yan, L.; Deng, X.; Li, J. Integrated energy hub optimization in microgrids: Uncertainty-aware modeling and efficient operation. *Energy* **2024**, *291*, 130391. [CrossRef]
37. Remus, T.; Liserre, M.; Rodriguez, P. *Grid Converters for Photovoltaic and Wind Power Systems*; John Wiley & Sons: Hoboken, NJ, USA, 2011.
38. Dragičević, T. Model predictive control of power converters for robust and fast operation of AC microgrids. *IEEE Trans. Power Electron.* **2017**, *33*, 6304–6317. [CrossRef]

39. International Energy Agency (IEA). *Energy Demand from Artificial Intelligence*; IEA: Paris, France, 2024. Available online: <https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai>, (accessed on 14 November 2025).
40. Isazadeh, A.; Ziviani, D.; Claridge, D.E. Global trends, performance metrics, and energy reduction measures in datacom facilities. *Renew. Sustain. Energy Rev.* **2023**, *174*, 113149. [CrossRef]
41. Kim, J.H.; Shin, D.U.; Kim, H. Data Center Energy Evaluation Tool Development and Analysis of Power Usage Effectiveness with Different Economizer Types in Various Climate Zones. *Buildings* **2024**, *14*, 299. [CrossRef]
42. CISPR 11; Industrial, Scientific and Medical Equipment—Radio-Frequency Disturbance Characteristics—Limits and Methods of Measurement. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2015.
43. CISPR 32; Electromagnetic Compatibility of Multimedia Equipment—Emission Requirements. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2015.
44. IEC 61850-9-3; Communication Networks and Systems for Power Utility Automation—Part 9-3: Precision Time Protocol Profile for Power Utility Automation. International Electrotechnical Commission: Geneva, Switzerland, 2016.
45. IEC 62541; OPC Unified Architecture. International Electrotechnical Commission: Geneva, Switzerland, 2016.
46. IEC 62430; Environmentally Conscious Design (ECD)—Principles, Requirements and Guidance. International Electrotechnical Commission: Geneva, Switzerland, 2019.
47. IEC 62402; Obsolescence Management, 2nd ed. IEC: Geneva, Switzerland, 2019.
48. ISO 9001; Quality Management Systems—Requirements. ISO: Geneva, Switzerland, 2015.
49. IEC 60300; Dependability Management. IEC: Geneva, Switzerland, 2003.
50. Google DeepMind. *DeepMind AI Reduces Google Data Centre Cooling Bill by 40%*; Google: Mountain View, CA, USA, 2018.
51. Microsoft. *RECOM: Resource-Efficient Cloud Management*; Microsoft: Redmond, WA, USA, 2023.
52. Navitas. *GaN Fast Power Bench*; Navitas Semiconductor: El Segundo, CA, USA, 2022.
53. Texas Instruments (TI). *GaN Evaluation Modules*; Texas Instruments: Dallas, TX, USA, 2021.
54. Statista. *Energy Use of Facebook (Meta) Data Centers*; Statista: Hamburg, Germany, 2023. Available online: <https://www.statista.com/statistics/580087/energy-use-of-facebook-meta/> (accessed on 5 November 2025).
55. Google. *Data-Center Freecooling Study*; Google: Mountain View, CA, USA, 2020; Available online: <https://www.google.com/about/datacenters/efficiency/> (accessed on 5 November 2025).
56. ANSI/ASHRAE Standard 90.4; Energy Standard for Data Centers. ASHRAE: Atlanta, GA, USA, 2022.
57. NXP Semiconductors. *GaN Reliability Paper*; NXP Semiconductors: Eindhoven, The Netherlands, 2022.
58. Fraunhofer ISE. *GaN Power Device Reliability Study*; Fraunhofer Institute for Solar Energy Systems (ISE): Freiburg, Germany, 2020.
59. Navitas. *GaN Whitepapers*; Navitas Semiconductor: El Segundo, CA, USA, 2021–2023.
60. Infineon Technologies. *GaN Whitepapers*; Infineon Technologies: Neubiberg, Germany, 2021–2023.
61. EPC. *GaN Whitepapers*; Efficient Power Conversion (EPC): El Segundo, CA, USA, 2021–2023.
62. Rajendran, G.; Raute, R.; Caruana, C. A Comprehensive Review of Solar PV Integration with Smart-Grids: Challenges, Standards, and Grid Codes. *Energies* **2025**, *18*, 2221. [CrossRef]
63. Zavodovski, A.; Gebremehdin, A.; Matta, A.; Kung, A.; Krukowski, A.; Hovsto, A.; Mukherjee, A.; Zachäus, C.; Filipovic, D.; Pongracz, E.; et al. *AI in Energy*; AIOTI: Brussels, Belgium, 2024.
64. Nafeez, A. “Planetary phase shift” as a new systems framework to navigate the evolutionary transformation of human civilization. *Foresight* **2024**, *27*, 240–266.
65. International Energy Agency. *AI Is Set to Drive Surging Electricity Demand from Data Centers While Offering the Potential to Transform How the Energy Sector Works*. *IEA News*. 10 April 2025. Available online: <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works> (accessed on 10 October 2025).
66. IEC 62740; Application of Reliability, Availability, Maintainability and Safety (RAMS) for Railway Systems. IEC: Geneva, Switzerland, 2021.
67. IEC 62402; Application of Reliability, Availability, Maintainability and Safety (RAMS)—Guidance on Achieving RAMS. IEC: Geneva, Switzerland, 2019.

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.