

Article

Real-Time On-MCU Open-Circuit Fault Diagnosis of Electric-Vehicle Inverters Using a Lightweight Angular Sector-Energy Network

Mingxing Fang, Wenxu Yan and Wenyuan Wang *

School of Automation and Intelligent Science, Jiangnan University, Wuxi 214122, China;
1036230211@stu.jiangnan.edu.cn (M.F.); ywx01@jiangnan.edu.cn (W.Y.)

* Correspondence: wenyuanwang@jiangnan.edu.cn

Abstract

Power-switch open-circuit (OC) faults distort electric-vehicle (EV) inverter phase currents and require fast on-board diagnosis for fault-tolerant control. Trajectory-image methods encode the α - β current-vector trajectory as a binary image and classify it with a convolutional neural network (CNN); however, the baseline uses 6.46×10^5 parameters and 3.31×10^7 multiply-accumulate (MAC) operations per inference, which is costly for motor-control microcontrollers (MCUs). Here, each one-cycle trajectory is represented by a 36-dimensional normalized angular sector-energy vector and classified by a compact two-stage multilayer perceptron. Sector accumulation averages zero-mean measurement noise in the representation, without relying on noise-augmented training. The locating stage uses 1.58×10^4 parameters and 1.56×10^4 MACs per inference, 97.55% and 99.95% fewer than the baseline CNN; the complete pipeline runs on a TI F28379D in 0.52 ms. On measured resistive-load currents, both methods reach 100% accuracy from 40 to 20 dB, whereas the proposed method remains more accurate at 15 and 10 dB, including under 88% phase-current unbalance. A supplementary balanced *RL*-load experiment preserves 100% clean accuracy, confirming MCU-executable diagnosis under a lagging power-factor load for embedded EV inverter protection.

Keywords: electric vehicle; three-phase inverter; open-circuit fault diagnosis; current-vector trajectory; angular sector-energy feature; multilayer perceptron; embedded deployment; noise robustness



Academic Editor: Michael Fowler

Received: 9 June 2026

Revised: 8 July 2026

Accepted: 9 July 2026

Published: 11 July 2026

Copyright: © 2026 by the authors.

Published by MDPI on behalf of the World Electric Vehicle Association.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

In an electric vehicle (EV), the three-phase voltage-source inverter (VSI) converts the battery dc power into the ac power required by the traction motor, and power-electronics reliability is a key requirement for such drivetrain systems [1]. In some integrated on-board charger and vehicle-to-grid (V2G) topologies, the traction inverter or the motor windings can also participate in bidirectional energy transfer between the battery and the grid [2]. The traction inverter is also closely associated with the dc-link reliability chain [3]. The power switches, such as insulated-gate bipolar transistors (IGBTs), silicon-carbide (SiC) metal-oxide-semiconductor field-effect transistors (MOSFETs), and silicon MOSFETs, are among the most failure-prone components because they are repeatedly exposed to electrical and thermal stresses. Gate-driver faults, bond-wire fatigue, and thermal cycling may lead to an open-circuit (OC) fault, which distorts the phase currents, increases current harmonics, and may further degrade drivetrain stability and safety [4]. For EV powertrains that

require limp-home operation, fault-tolerant control relies on timely and reliable diagnostic information. Therefore, OC fault diagnosis should be fast, accurate, and implemented on the controller itself rather than on an external computer.

Existing OC diagnosis methods fall into two broad groups: rule-based methods and data-driven methods. Rule-based methods derive a diagnostic variable from measured currents, voltages, residuals, or vector trajectories and compare it with preset thresholds. Among current-based schemes, Rothenhagen and Fuchs [5] evaluated normalized average-current (Park's-vector) indicators for IGBT OC faults in voltage-source drives, and Guo et al. [6] extended the average-current Park's-vector approach to a six-phase fault-tolerant permanent-magnet synchronous motor (PMSM); Zhang et al. [7] instead used the phase of the current vector to detect weak OC faults that cause only small current distortion. Model- and observer-based schemes form residuals from a converter model: Wang et al. [8] estimated a voltage-vector residual by tracking the current path in a T-type inverter, Li et al. [9] used a predictive current-vector residual for a three-level active neutral-point-clamped (ANPC) inverter, and Diao et al. [10] built a hybrid-model flux observer for real-time diagnosis in sensorless vector-controlled drives. These methods are computationally light and physically interpretable, but their thresholds and model parameters must be retuned for each topology and operating point and become fragile under measurement noise, parameter drift, or severe load imbalance, where false alarms and missed detections increase.

Data-driven methods reduce this manual tuning by learning a classifier from current features. Cai et al. [11] combined principal-component analysis with a Bayesian network for a PMSM inverter, and Gou et al. [12] applied ReliefF feature selection to diagnose IGBT and current-sensor faults simultaneously in an induction-motor drive. More recent works learn the mapping with neural networks: Yan et al. [13] trained a deep neural network directly on the three-phase currents, Das et al. [14] targeted online diagnosis for a three-level six-phase PMSM, and Chu et al. [15] reported a neural-network diagnosis system for a three-phase PMSM inverter. These methods relax the dependence on hand-set thresholds, yet those that still rely on predefined feature extraction remain sensitive to the chosen feature: if it is not sufficiently discriminative, fault-related information can be weakened or lost, and under extreme current imbalance amplitude- or variance-dominated features are dominated by the heavily loaded phases, reducing sensitivity to faults in the lightly loaded phase.

To reduce the dependence on hand-crafted features, recent studies represent the fault signature by the geometry of the α - β current-vector trajectory [16–18]. Hang et al. [17] converted the normalized current-vector trajectory into an image and classified it with a wavelet convolutional neural network (CNN) that remains robust with small training sets, whereas Sun et al. [18] extracted the adjacent trend-line relationship of the trajectory as a diagnostic rule. Most relevant to this work, Lei et al. [16] rasterized the α - β current vector into a binary trajectory image and trained a two-stage CNN for fault detection and switch locating, reporting 100% accuracy over all 21 OC fault cases and a healthy condition at a phase-current unbalance rate (PCUR) of up to 88%. The binary image preserves fault-related information even when one phase current is strongly attenuated under extreme imbalance.

However, EV on-board deployment still faces three practical obstacles. First, the model size and computational cost are high: about 6.46×10^5 parameters and 3.31×10^7 multiply-accumulate (MAC) operations per inference make the method more suitable for a host personal computer or graphics processing unit (PC/GPU) than for direct microcontroller (MCU) deployment. According to the throughput measured on the target control MCU in this work, one CNN inference is estimated to require several hundred milliseconds, far beyond a single fundamental cycle. Second, the sparse binary trajectory image is sensitive

to a low signal-to-noise ratio (SNR): a noisy sample may be mapped to an otherwise empty pixel, introducing extra activations and distorting the trajectory pattern, so a single-switch fault may become closer to a corresponding double-switch pattern. Third, if diagnosis is performed on a host computer, the phase currents must be continuously streamed off-board, which introduces additional communication latency and reliability requirements in the electromagnetically harsh vehicle environment.

For EV applications, OC diagnosis is most useful when it runs in real time on the existing motor-control MCU. This setting leaves limited computation, memory, cost, and installation space. The diagnostic representation should therefore be low-dimensional, inexpensive to compute, robust to measurement noise and current imbalance, and still tied to the physical fault signature [19,20].

To meet these requirements, this paper proposes a lightweight angular sector-energy diagnosis network. The α - β plane is divided into 36 sectors of 10° ; current magnitudes within each sector are accumulated and normalized to form a 36-dimensional feature. A compact two-stage multilayer perceptron (MLP) then uses this feature for fault detection and switch locating. Because each sector aggregates multiple samples, the feature reduces the influence of zero-mean measurement noise at individual samples. Its robustness therefore comes mainly from the representation itself, rather than from noise-augmented training. The main contributions are as follows:

- An angular sector-energy feature for inverter OC diagnosis that is noise-robust and physically interpretable: compared with the sparse binary trajectory image, the per-sector magnitude accumulation attenuates the effect of measurement noise, so the model does not require noise-augmented training.
- A two-stage detect-then-locate model that is substantially lighter than the baseline CNN: the locating network has only 1.58×10^4 parameters and 1.56×10^4 MACs per inference, corresponding to 97.55% and 99.95% fewer, respectively, than the baseline CNN's 6.46×10^5 parameters and 3.31×10^7 MACs.
- Deployment and validation on the control MCU itself (TI F28379D): On clean measured experimental data, the deployed model correctly identifies all 22 modes, with a worst-case two-stage inference time of 0.52 ms, and uses 121 kB of Flash and 3.6 kB of RAM. The healthy-state and $2 \rightarrow 3$ A peak current-step tests further show that the detector does not trigger false alarms.
- Robustness on real-hardware data under both balanced and extremely unbalanced loading (PCUR up to 88%): The proposed method matches the baseline CNN from 40 to 20 dB SNR and achieves higher locating accuracy at 15 and 10 dB, where the CNN accuracy drops markedly. Supplementary tests under a balanced inductive *RL* load further examine whether the same angular feature and embedded diagnosis flow, including online and on-MCU checks, carry over when the resistor bank is replaced by a lagging power-factor load.

The distinction from earlier trajectory-image studies is therefore not another neural-network block, but the representation and where it can be executed. Instead of rasterizing the α - β trajectory into a binary image for a large CNN [16–18], this work encodes the trajectory as a low-dimensional normalized angular sector-energy vector. The per-sector aggregation is noise-averaging by construction, so robustness comes from the feature itself; the resulting classifier is small enough for real-time execution on the motor-control MCU. Table 1 positions this design against representative OC-diagnosis approaches and general lightweight networks.

General-purpose lightweight networks such as MobileNetV2 [21] and SqueezeNet [22], and the broader family of TinyML models [23], are designed mainly for two-dimensional image inputs and still contain on the order of 10^6 parameters. Applying them to

the 36-D feature used here would be unnecessarily large: a compact MLP of only 1.58×10^4 parameters already saturates the clean-data accuracy and is the model actually deployed on the MCU. These networks are therefore included for context rather than as direct baselines. The head-to-head comparison is made against the trajectory-image CNN [16], the closest prior method on the same task and measured data.

Table 1. Positioning of the proposed method against representative OC-diagnosis approaches and lightweight network families. The proposed method is highlighted in bold.

Approach	Input/Feature	Parameter Scale	On-MCU Real-Time	Noise/Imbalance Robustness
Threshold/rule-based	Currents, residuals	None	Possible but fragile	Low
Predefined-feature data-driven	Hand-crafted features	Varies	Rarely reported	Feature-sensitive
Trajectory-image CNN [16]	Binary α - β image	$\approx 6.5 \times 10^5$	No (PC/GPU)	High under reported conditions; fragile at low SNR
Lightweight CNN [21,22]	2-D image	≈ 1.2 – 3.4×10^6	Vision-oriented	Not directly applicable
Proposed	36-D sector-energy vector	1.58×10^4	Yes (0.52 ms)	High; above CNN at 10 dB

Bold text identifies the proposed method and its corresponding feature, parameter count, real-time result, and robustness summary. The CNN and proposed MLP parameter counts are computed from the implemented networks used in this study. The lightweight CNN row gives published representative model scales and is included for design context rather than as a direct OC-diagnosis baseline.

The remainder of this paper is organized as follows. Section 2 analyzes the OC-fault current-vector trajectory and the noise sensitivity of the binary-image representation. Section 3 presents the proposed sector-energy feature and the two-stage MLP. Section 4 details the embedded implementation on the F28379D MCU. Section 5 reports comparative offline, online, on-board, and supplementary *RL*-load experimental results. Section 6 discusses the limitations, and Section 7 concludes the paper.

2. Current-Vector Trajectory and Noise Fragility Analysis

2.1. Inverter Model and Open-Circuit Fault Modes

A three-phase two-level VSI comprises six controllable switches, denoted S_{a1}, S_{a2} for the upper and lower switches of phase *a*, S_{b1}, S_{b2} for phase *b*, and S_{c1}, S_{c2} for phase *c*. As shown in Figure 1, the inverter is fed by a dc bus and connected to an *LC* output filter and a three-phase resistive load. The supplementary inductive-load validation in Section 5.7 keeps the same inverter and switch labeling but replaces this bench with a direct balanced star-connected *RL* load, as shown in Figure 2.




Figure 1. Topology of the three-phase two-level VSI with an *LC* output filter.




Figure 2. Topology of the supplementary three-phase two-level VSI test bench with a direct balanced star-connected RL load.

An inverter operational mode is described by the binary label vector $\mathbf{y}$:

$$\mathbf{y} = [S_{a1}, S_{a2}, S_{b1}, S_{b2}, S_{c1}, S_{c2}],$$

where $S_j \in \{0, 1\}$, $S_j = 0$ denotes a healthy switch, and $S_j = 1$ denotes an OC fault. Because simultaneous OC faults in three or more switches are much less likely and usually correspond to more severe system failures, this paper considers the healthy mode, single-switch OC faults, and double-switch OC faults. This yields $1 + 6 + \binom{6}{2} = 22$ operational modes, as listed in Table 2.

Table 2. The 22 inverter operational modes.

#	Mode	#	Mode	#	Mode	#	Mode
1	Healthy	7	S_{c2}	13	$S_{a2} \& S_{b1}$	19	$S_{b1} \& S_{c2}$
2	S_{a1}	8	$S_{a1} \& S_{a2}$	14	$S_{a2} \& S_{b2}$	20	$S_{b2} \& S_{c1}$
3	S_{a2}	9	$S_{a1} \& S_{b1}$	15	$S_{a2} \& S_{c1}$	21	$S_{b2} \& S_{c2}$
4	S_{b1}	10	$S_{a1} \& S_{b2}$	16	$S_{a2} \& S_{c2}$	22	$S_{c1} \& S_{c2}$
5	S_{b2}	11	$S_{a1} \& S_{c1}$	17	$S_{b1} \& S_{b2}$		
6	S_{c1}	12	$S_{a1} \& S_{c2}$	18	$S_{b1} \& S_{c1}$		

2.2. Current-Vector Trajectory Under Faults

The measured three-phase currents are mapped to the stationary α - β frame by the Clarke transform:

$$\begin{bmatrix} i_\alpha \\ i_\beta \end{bmatrix} = \begin{bmatrix} \frac{2}{3} & -\frac{1}{3} & -\frac{1}{3} \\ 0 & \frac{\sqrt{3}}{3} & -\frac{\sqrt{3}}{3} \end{bmatrix} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix}. \quad (1)$$

The current vector $\mathbf{i}_{\alpha\beta}(t) = i_\alpha(t) + j i_\beta(t)$ then traces a closed trajectory over one fundamental cycle. Under balanced and nearly sinusoidal currents, the healthy trajectory approximately forms a circle. Load imbalance may make the trajectory elliptical, while measurement noise introduces local perturbations around the nominal curve, as shown in Figure 3.

When an OC fault occurs, the current path of the faulty phase is blocked for one current polarity, and the phase current can no longer follow the normal pulse-width modulation (PWM) command in that interval. In the α - β plane, this appears as a reduction in or loss of trajectory energy in the direction associated with the faulty switch. Under the current direction defined in Figure 1 and the adopted switch labeling, upper- and lower-switch faults correspond to opposite current polarities and therefore produce energy loss in different half-cycle regions. Double-switch faults combine the effects of two single-switch

faults, leading to multiple missing angular sectors, a contracted trajectory, or, in extreme cases, an approximately straight line [16]. Therefore, the discriminative information of an OC fault is mainly angular: it is reflected by which angular sectors lose energy and by the orientation of the remaining trajectory, as illustrated in Figure 3.




Figure 3. Representative current-vector trajectories in the α - β frame with the proposed 10° angular sector partition: (a) healthy, balanced; (b) S_{b1} , balanced; (c) S_{b1} and S_{c2} , balanced; (d) healthy, unbalanced; (e) S_{b1} at 10 dB SNR; (f) S_{b1} and S_{c2} at 10 dB SNR. In the two noisy panels, the blue dashed curve is the noise-free trajectory and the red dots are the 10 dB samples.

2.3. Binary Trajectory Matrix and Its Noise Fragility

The baseline trajectory-matrix scheme [16] discretizes the current-vector trajectory into a $D \times D$ binary matrix $\mathbf{X}$. For each sample t_i in one fundamental cycle ($i = 1, \dots, N_s$), the one-based pixel indices used in the reproduced baseline are first computed as

$$\tilde{m}_i = \lfloor i_\alpha^{\text{pu}}(t_i) \frac{D}{4} + \frac{D}{2} \rfloor + 1, \quad \tilde{n}_i = \lfloor i_\beta^{\text{pu}}(t_i) \frac{D}{4} + \frac{D}{2} \rfloor + 1, \quad (2)$$

where i^{pu} denotes the per-unit current and $\lfloor \cdot \rfloor$ is the floor operator. The indices are then clipped to the valid image range,

$$m_i = \text{clip}(\tilde{m}_i, 1, D), \quad n_i = \text{clip}(\tilde{n}_i, 1, D), \quad (3)$$

where $\text{clip}(x, 1, D) = \min\{\max\{x, 1\}, D\}$. The binary trajectory matrix is defined by whether any sampled point falls into each pixel,

$$\mathbf{X}(m, n) = \begin{cases} 1, & \exists i \in \{1, \dots, N_s\} : (m, n) = (m_i, n_i), \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The resulting binary matrix is then classified by a CNN.

For clean signals, this encoding preserves the main geometric shape of the trajectory. However, $\mathbf{X}$ is sparse and binary: one fundamental cycle activates only a thin set of cells. Under measurement noise, the mapped position of a sample in the α - β plane may shift to an otherwise empty cell, thereby introducing extra activations. This effect is visible

in the two noisy panels of Figure 3: at 10 dB the measured samples (red dots) scatter well beyond the noise-free trajectory (blue dashed) and fall into angular sectors that are empty in the clean case, so the sparse fault pattern is smeared. Since the CNN learns from the spatial pattern of activated cells, such isolated activations may partially refill the missing semicircle of a single-switch fault and make its pattern closer to that of a corresponding double-switch fault. This tendency is observed in Section 5, where S_{a1} and S_{a2} are increasingly misclassified as S_{a1} and S_{a2} as the SNR decreases. This noise sensitivity motivates an aggregated representation that reduces the influence of individual noisy samples.

3. Proposed Method: Angular Sector-Energy Network

3.1. Angular Sector-Energy Feature

Given one fundamental cycle of N_s samples, the current vector at instant t_i is described by its magnitude and angle:

$$r_i = \sqrt{i_\alpha^2(t_i) + i_\beta^2(t_i)}, \quad \theta_i = \text{mod}(\text{atan2}(i_\beta(t_i), i_\alpha(t_i)), 2\pi). \quad (5)$$

Here, $\text{atan2}(\cdot)$ is the four-quadrant arctangent, and the modulo operation maps the angle to $[0, 2\pi)$.

The angular range $[0, 2\pi)$ is partitioned into $K = 36$ uniform sectors of width $\Delta = 2\pi/K = 10^\circ$. The accumulated magnitude in sector k and its normalized value are

$$E_k = \sum_{i: [\theta_i/\Delta]=k} r_i, \quad f_k = \frac{E_k}{\sum_{m=0}^{K-1} E_m + \varepsilon}, \quad k = 0, \dots, K-1, \quad (6)$$

where ε is a small constant used to avoid division by zero. In this work, “energy” denotes the accumulated current-vector magnitude rather than the squared magnitude. This choice avoids excessive sensitivity to a few high-amplitude samples and keeps the feature numerically well conditioned.

The 36 normalized values form the proposed feature $\mathbf{f} = [f_0, \dots, f_{K-1}]^T \in \mathbb{R}^{36}$. The normalization in (6) reduces the influence of the overall current-amplitude scale, making the feature more comparable across different loading conditions. Since $\mathbf{f}$ describes how the trajectory magnitude is distributed over an angle, it is consistent with the physical signature of an OC fault discussed in Section 2. Figure 4 illustrates this for an S_{b1} fault, where the normalized sector energy decreases markedly in the angular region corresponding to the missing semicircle.




Figure 4. Construction of the proposed 36-D sector-energy feature for an S_{b1} fault: (a) α - β trajectory with the 10° sector partition; (b) normalized sector-energy feature.

3.2. Noise Robustness by Design

Let the measured current vector in the α - β frame be corrupted by zero-mean white Gaussian noise,

$$\mathbf{i}'_{\alpha\beta}(t_i) = \mathbf{i}_{\alpha\beta}(t_i) + \tilde{\mathbf{n}}_i, \quad \tilde{\mathbf{n}}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}).$$

For the magnitude $r_i = \|\mathbf{i}_{\alpha\beta}(t_i)\|_2$, the noisy magnitude can be approximated at moderate to high SNR by the first-order expansion:

$$r'_i \approx r_i + \mathbf{u}_i^\top \tilde{\mathbf{n}}_i, \quad \mathbf{u}_i = \frac{\mathbf{i}_{\alpha\beta}(t_i)}{r_i}.$$

The first-order perturbation $\mathbf{u}_i^\top \tilde{\mathbf{n}}_i$ is zero-mean. Therefore, when N_k samples fall into sector k , the standard deviation of the accumulated perturbation scales as $\sigma\sqrt{N_k}$, whereas the nominal accumulated magnitude scales as $\bar{r}N_k$. The relative perturbation of E_k is thus on the order of

$$\mathcal{O}\left(\frac{\sigma}{\bar{r}\sqrt{N_k}}\right),$$

which decreases as more samples are aggregated.

In contrast, the binary trajectory matrix in (4) maps each sample to a single cell. Under low SNR, a noisy sample may be mapped to an otherwise empty cell, introducing extra activations and distorting the sparse trajectory pattern. The proposed sector-energy feature reduces the influence of individual noisy samples through angular aggregation, and is therefore expected to degrade more gradually as the SNR decreases. This robustness mainly comes from the representation itself rather than from noise-augmented training: the networks are trained only on clean signals and are then evaluated under different SNR levels.

3.3. Two-Stage Lightweight Classifier

Following the detect-then-locate paradigm [16], two compact MLPs share the same 36-D input: a detection network for binary healthy/faulty classification and a locating network for 22-mode fault identification. Both networks adopt the structure $36 \rightarrow 128 \rightarrow 64 \rightarrow 32 \rightarrow C$, where $C = 2$ for detection and $C = 22$ for locating. The hidden layers use rectified linear unit (ReLU) activations, and dropout with a rate of 0.2 is applied during training.

For an input feature $\mathbf{f}$, z-score normalization is first applied as

$$\mathbf{z} = (\mathbf{f} - \boldsymbol{\mu}_f) \oslash \mathbf{s}_f,$$

where $\boldsymbol{\mu}_f$ and $\mathbf{s}_f$ are the mean and standard deviation computed from the training set, and $\oslash$ denotes element-wise division. The forward propagation is then given by

$$\mathbf{h}_l = \text{ReLU}(\mathbf{W}_l \mathbf{h}_{l-1} + \mathbf{b}_l), \quad \mathbf{p} = \text{softmax}(\mathbf{W}_L \mathbf{h}_{L-1} + \mathbf{b}_L), \quad (7)$$

where $\mathbf{h}_0 = \mathbf{z}$. The predicted class is

$$\hat{c} = \arg \max_c p_c.$$

Both networks are trained on clean data only, using the cross-entropy loss and the Adam optimizer with a mini-batch size of 128 for a fixed 60 epochs. The initial learning rate is 10^{-3} and is reduced by a factor of 0.2 after 40 epochs, and the training samples are shuffled every epoch. No early stopping is applied, so the network obtained at the final epoch is retained; regularization relies on the dropout layers introduced above (rate 0.2) together with the implicit averaging of the sector-energy feature. The locating network

contains only 1.58×10^4 learnable parameters and requires 1.56×10^4 MACs per inference, far below the baseline CNN (6.46×10^5 parameters and 3.31×10^7 MACs). The embedded implementation is detailed in Section 4.

3.4. Overall Diagnosis Scheme

Figure 5 summarizes the overall diagnosis scheme. At each sampling instant, the Clarke transform (1) and the per-sample magnitude/angle (5) are computed, and the corresponding sector energy is updated incrementally, which requires only $\mathcal{O}(1)$ operations per sample.

In the present implementation, the diagnostic feature is computed over a one-fundamental-cycle window. For steady offline evaluation the window advances by one cycle. For online operation, the same one-cycle window is maintained as a sliding buffer and refreshed every 50 samples (2.5 ms at the 20 kHz sampling rate), so the detection MLP can be evaluated at sub-cycle intervals once the first complete window is available. A first faulty indication can therefore appear within one fundamental cycle, and the experiments in Section 5 report confirmed sub-cycle detections.

To suppress isolated false alarms, a configurable confirmation rule can require N_c consecutive faulty decisions before issuing the confirmed alarm. In a sliding-window implementation this adds $(N_c - 1)$ update intervals, not necessarily a full fundamental cycle. The online fault-latency tests below use $N_c = 1$ to measure the intrinsic detection response, whereas the healthy and current-step false-alarm checks use $N_c = 2$ as an additional debounce margin. The locating MLP is then evaluated on the first complete post-fault feature window, so that its input contains only post-fault data and the faulty switch(es) can be identified. This separates a fast initial indication from a more conservative confirmed alarm, following the detect-then-locate principle.




Figure 5. Overall two-stage angular sector-energy diagnosis scheme.

4. Embedded Implementation on the Control MCU

4.1. Computational Footprint and the Infeasibility of an On-MCU CNN

The baseline CNN comprises 6.46×10^5 learnable parameters and 3.31×10^7 MACs per inference. Stored in single precision, its weights alone occupy about 2.58 MB, exceeding the 1 MB on-chip Flash of a typical motor-control MCU; even with 8 bit quantization (≈ 0.63 MB), the 3.31×10^7 MACs would require, at the single-precision throughput measured on the target device (≈ 3.4 cycles/MAC), about 1.0×10^8 clock cycles, i.e., several hundred milliseconds, far beyond one fundamental cycle. This estimate counts only the measured MAC throughput and is therefore sufficient to show the on-MCU infeasibility of the CNN. The proposed locating network needs only 1.58×10^4 parameters and 1.56×10^4 MACs (Table 3)—a per-inference cost the MCU handles in real time.

Table 3. Computational/memory footprint and on-MCU feasibility. The more favorable value is highlighted in bold.

Metric	Baseline CNN [16]	Proposed Energy-MLP
Learnable parameters	6.46×10^5	1.58×10^4
MACs per inference	3.31×10^7	1.56×10^4
Weight memory (FP32)	≈ 2.58 MB	121 kB
On-MCU inference time	~ 0.5 s (infeasible)	0.52 ms (measured)
Real time (<20 ms cycle)	No	Yes
Execution platform	Host PC/GPU	On-board MCU

Parameters and MACs are listed for the locating network (22 classes), the larger of the two stages; weight memory, inference time and footprint refer to the complete deployed detector + locator (both networks).

4.2. On-MCU Realization

The method is implemented on a Texas Instruments F28379D (dual 200 MHz C28x cores with a single-precision floating-point unit (FPU) and a trigonometric math unit (TMU)). The per-sample operations of (1) and (5) (the square root and atan2 are accelerated by the on-chip FPU and TMU when the compiler is configured with FPU32/TMU support) are folded into the existing 20 kHz analog-to-digital conversion interrupt and accumulated incrementally into the 36 sector energies, which adds only a small per-sample load to the interrupt. At each diagnostic update, the one-cycle feature window is normalized and the two-stage MLP (7) is evaluated in the main loop. In the steady offline setting, the update interval is one fundamental cycle; in the online tests, the one-cycle buffer is refreshed every 50 samples, corresponding to 2.5 ms at the 20 kHz sampling rate. The detection network runs at each update, and the locating network is triggered only after a confirmed fault. All weights are stored as read-only constants in Flash; only the feature buffer, hidden-layer activations, and diagnostic states are kept in RAM at run time. The safety-critical modulation remains on CPU1; the diagnosis may optionally be offloaded to CPU2, although a single core is sufficient here.

4.3. Measured Footprint and Inference Latency

On the target device the two networks occupy 121 kB of Flash for weights and about 150 kB of total program Flash (15% of the 1 MB available), with 3.6 kB of RAM (1.8% of 204 kB). The measured single-inference time is 219.8/407.2/516.4 μ s (minimum/mean/maximum); the worst case, which includes both the detection and locating stages, is 0.52 ms, only 2.6% of one 50 Hz fundamental cycle. Table 3 summarizes the comparison.

5. Experimental Results

5.1. Experimental Setup and Validation Roadmap

The experimental validation is organized progressively. Stage I uses a TI F28379D control board driving a three-phase two-level inverter with an LC output filter that feeds three-phase purely resistive loads. These resistive-load tests are retained as the primary comparison platform because they match the baseline trajectory-image study and allow a controlled evaluation of load imbalance, measurement noise, online response, and on-MCU execution. Two resistive loading conditions are tested: balanced and extremely unbalanced with one phase set to a much higher resistance. Following the PCUR of [16], the unbalance is quantified by

$$\text{PCUR} = \frac{|I_{\max} - I_{\text{avg}}|}{I_{\text{avg}}} \times 100\%, \quad (8)$$

where I_{avg} is the average amplitude of the three-phase currents and $I_{\max}$ is the amplitude of the phase current that deviates most from I_{avg} ; the 10/10/125 Ω load yields $\text{PCUR} \approx 88\%$.

The three-phase currents are sampled at 20 kHz ($N_s = 400$ samples per 50 Hz fundamental cycle). For each of the 22 operational modes, 1000 fundamental cycles are recorded under each loading condition, giving $22 \times 2 \times 1000 = 44,000$ labelled cycles in total. Within each class the cycles are partitioned 80/20 into a training and a test set with a fixed random seed; the balanced and unbalanced data are pooled before this split, so both loading conditions appear in training and in testing, while the test accuracy is still reported separately for each condition. For the offline robustness study, white Gaussian noise is added to the measured currents at SNRs from 40 down to 10 dB to emulate measurement noise; in all cases the networks are trained on clean data only. The 22 operational modes are injected in hardware through the ePWM trip-zone logic of the F28379D, which forces the corresponding switch(es) off: the C2000 trip-zone submodule can rapidly drive the PWM outputs to a defined safe state under a fault condition, making it well suited to emulating switch open-circuit faults. The key parameters of the resistive and supplementary *RL* validation stages are summarized in Table 4.

Table 4. Key operating conditions for the resistive and supplementary *RL* validation stages.

Parameter	R Load, Balanced	R Load, Unbalanced	RL Load, Balanced
DC-bus voltage V_{dc}	60 V	60 V	15 V
Fundamental frequency f_0	50 Hz	50 Hz	50 Hz
Switching frequency f_{sw}	20 kHz	20 kHz	20 kHz
Sampling rate f_s	20 kHz (400/cycle)	20 kHz (400/cycle)	20 kHz (400/cycle)
Output connection	LC filter + load	LC filter + load	Direct load
Filter values	$L_f = 2$ mH, $C_f = 10$ μ F	$L_f = 2$ mH, $C_f = 10$ μ F	Not used
Phase load	10/10/10 Ω	10/10/125 Ω (PCUR $\approx$ 88%)	$R = 2.5$ Ω , $L = 2$ mH per phase
Current-step check	2 $\rightarrow$ 3 A peak	2 $\rightarrow$ 3 A peak	2 $\rightarrow$ 3 A peak
Cycles per mode	1000	1000	1000

The 2 A and 3 A current levels reported in this paper denote phase-current peak amplitudes, not RMS values.

The remainder of this section follows the same validation chain. The resistive-load benchmark first quantifies noise robustness, sector-count sensitivity, false-alarm immunity, online response, and on-board MCU validation. The inductive-load extension then replaces the resistive load bank with a balanced three-phase *RL* load to test whether the same angular sector-energy representation and embedded diagnosis flow remain effective under a lagging power-factor load.



5.2. Resistive-Load Benchmark: Offline Noise Robustness

Following the protocol of [16], for each mode and SNR level 100 cycles are randomly sampled and encoded, giving 2200 test samples per SNR. For a fair comparison, the baseline CNN of [16] is reimplemented and retrained on the same measured clean data and evaluated on the same noisy cycles as the proposed network; the accuracies reported here are therefore obtained under the present experimental conditions rather than copied from [16]. Figure 6 and Table 5 report the locating accuracy. Within the SNR range reported by the baseline (40–20 dB), both methods reach 100% under both loads, faithfully reproducing the baseline. Beyond this range, at 15 and 10 dB, the CNN accuracy decreases markedly (balanced: 63.6% and 40.9%; unbalanced: 74.0% and 50.2%), whereas the proposed feature remains comparatively robust (balanced: 100% and 94.8%; unbalanced: 98.7% and 63.5%), which is consistent with the noise-averaging analysis in Section 3. Figure 7 illustrates this at 10 dB: the baseline CNN assigns many faults to a few classes, in particular confusing single-switch with double-switch faults, whereas the proposed feature stays close to the diagonal.

Table 5. Offline locating accuracy under added measurement noise.

Load	Method	40 dB	30 dB	20 dB	15 dB	10 dB
Balanced	Baseline CNN [16]	100	100	100	63.6	40.9
	Proposed	100	100	100	100	94.8
Unbalanced (PCUR 88%)	Baseline CNN [16]	100	100	100	74.0	50.2
	Proposed	100	100	100	98.7	63.5

Values are percentages on identical real-hardware test cycles. The baseline CNN [16] is reimplemented and evaluated under the same conditions. The higher accuracy at each SNR is highlighted in bold.

**Figure 6.** Offline locating accuracy versus SNR (1 dB resolution) under balanced (**left**) and extremely unbalanced (**right**, PCUR = 88%) loading.**Figure 7.** Row-normalized confusion matrices (diagonal = recall, %) at 10 dB SNR.

5.3. Resistive-Load Benchmark: Sector-Count Sensitivity

To check that $K = 36$ is a justified rather than an arbitrary choice, the sector count was varied over $K \in \{18, 24, 36, 48, 72\}$. For each K the feature is recomputed and the locating MLP is retrained from scratch with identical settings, and the held-out test set is evaluated under the same per-cycle noise protocol; the results, averaged over three training seeds, are listed in Table 6. The $K = 36$ row therefore differs slightly from the single representative run in Table 5, but the operating conclusion is the same.

Table 6. Sensitivity of the locating accuracy (%) to the sector count K ; the clean and 20 dB cases achieve 100% for all K values under both loads. The value used in this paper is highlighted in bold.

K	Learnable Params	15 dB (bal./unb.)	10 dB (bal./unb.)
18	13,494	100.0/99.8	90.9/71.5
24	14,262	99.9/97.2	95.8/71.0
36	15,798	99.9/98.8	93.5/66.4
48	17,334	99.6/97.1	91.3/61.4
72	20,406	99.6/98.2	91.9/61.2

Two trends are visible. First, within the operating range the accuracy is essentially insensitive to K : the clean and 20 dB cases stay at 100% for all values, and at 15 dB every K exceeds 97% under both loads, so K is not a critical hyper-parameter. Second, at the most demanding 10 dB unbalanced corner, the accuracy decreases monotonically with K , from 71.5% at $K = 18$ to 61.2% at $K = 72$. This directly corroborates the noise-averaging analysis of Section 3: fewer, larger sectors aggregate more samples per sector, and the standard deviation of the sector energy, which scales as $\sim 1/\sqrt{N_k}$, is correspondingly smaller. The $K = 36$ entry is consistent with the main results in Table 5; $K = 36$ (a 10° partition) is retained as a compromise between angular resolution and noise averaging. It is fine enough to preserve the sector-level shape of the fault-induced gap, but still leaves enough samples in each sector for averaging. If maximum robustness at an extremely low SNR was the only objective, a smaller K could be adopted at the cost of angular resolution. The aim of this ablation is therefore not to claim that $K = 36$ is globally optimal, but to show that the diagnosis is insensitive to K over the practical SNR range.

5.4. Resistive-Load Benchmark: False-Alarm Immunity

This paper first verifies that the detector stays silent when the inverter is healthy. Figure 8 examines this on real hardware under four conditions: healthy operation at balanced and at extremely unbalanced (PCUR = 88%) loading, and a sudden $2 \rightarrow 3$ A peak (+50%) phase-current step (an abrupt load transient) at each loading. Over the entire ≈ 2 s record (≈ 100 fundamental cycles) the detection output stays at 0 (healthy) in every case, with *zero* confirmed false alarms; the detector does not trip, even on the current step, and the $N_c = 2$ confirmation provides a further margin against isolated mis-detections. Figure 8 plots representative excerpts from these full records.




Figure 8. False-alarm immunity on real-hardware currents: (a) healthy, balanced; (b) healthy, unbalanced (PCUR = 88%); (c) a 2 → 3 A peak current step, balanced; (d) the same step, unbalanced.

5.5. Resistive-Load Benchmark: Online Real-Time Diagnosis

In the online experiment, the DSP continuously acquired the measured phase-current samples while the two-stage diagnostic pipeline processed the one-cycle features in real time. Figure 9 shows representative detection-and-locating timelines for a single-switch (S_{b1}) and a double-switch (S_{b1} and S_{c2}) fault under both loads. In every case the confirmed fault indication appears within one fundamental cycle (sub-cycle latency, 2.45–17.45 ms under the $N_c = 1$ latency protocol), and the faulty switch(es) are correctly located in the following cycle, with no false alarm. The unbalanced cases also visibly exhibit the heavily attenuated current of the lightly loaded phase, yet diagnosis remains correct.




Figure 9. Online real-time diagnosis for representative faults: (a) balanced S_{b1} ; (b) balanced S_{b1} and S_{c2} ; (c) unbalanced S_{b1} ; (d) unbalanced S_{b1} and S_{c2} . In the upper trace of each panel, the blue solid line denotes the locating-label output and the red dashed line denotes the binary detection output. The black solid, red dashed, and blue dash-dotted vertical lines mark the fault instant t_1 , confirmed detection t_2 , and locating decision t_3 , respectively; the panel titles and red in-panel annotations list the detection and locating delays relative to t_1 .

5.6. Resistive-Load Benchmark: On-Board MCU Validation

With the model deployed on the F28379D, all 22 modes were injected on the LC-filtered resistive bench under both balanced and unbalanced loading. The on-board diagnosis correctly identified every mode (22/22) under both loads, at the worst-case inference time of 0.52 ms reported in Section 4. The diagnosis accuracy that the baseline CNN attains only on a host computer is thus obtained here on the control MCU itself for the original resistive-load platform.

5.7. Inductive-Load Extension: Supplementary RL Validation

To examine whether the angular sector-energy feature carries over beyond the purely resistive LC-filtered bench, a supplementary inductive-load experiment was added. The output side of the same inverter was connected directly to a balanced three-phase RL load (Figure 2; $R = 2.5 \, \Omega$ and $L = 2 \, \text{mH}$ per phase, star-connected), whose power-factor angle is $\arctan(2\pi fL/R) \approx 14^\circ$ at $f = 50 \, \text{Hz}$. The dc bus was lowered to 15 V so that the phase-current peak stayed at $\approx 2 \, \text{A}$, i.e., the same amplitude scale as the resistive datasets. All 22 modes were captured with the identical runtime protocol (1000 fundamental cycles per mode), and the two-stage network was retrained and evaluated on this RL data using the same 80/20 stratified split and training settings as used for the resistive-load experiments.

Figure 10 shows the measured α - β trajectories under the inductive load. The healthy trajectory remains a closed circle. A single-switch OC fault (S_{b1}) blocks one current polarity and leaves the characteristic angular gap, while a double-switch fault (S_{b1} and S_{c2}) distorts the loop further. This geometry is consistent with the resistive-load observations: the inductive power-factor angle rotates and reshapes the trajectory, but it does not fill the fault-induced gap.




Figure 10. Measured α - β current-vector trajectories under a balanced inductive RL load ($R = 2.5 \, \Omega$, $L = 2 \, \text{mH}$ per phase): (a) healthy; (b) single-switch fault S_{b1} ; (c) double-switch fault S_{b1} and S_{c2} .

On this RL data the proposed network reached 100% detection and 100% locating accuracy on the clean held-out set. The locating accuracy stayed at $\geq 99.98\%$ down to 10 dB SNR when noise was added during evaluation, consistent with the balanced resistive-load behavior reported above. This supplementary result indicates that the fault-induced angular energy gap is not an artifact of the original resistive load or the LC filter. The bench load is, however, balanced and only mildly inductive. A strongly inductive, variable-speed traction motor introduces back-EMF and additional dynamics; this boundary is discussed in Section 6.

Beyond this offline evaluation, the two-stage diagnosis was also run online on the same inductive bench. The energy-MLP processed the streamed currents with the same one-cycle diagnostic window and update protocol used in the resistive-load online tests. Across all 21 fault modes, it detected and located every fault correctly (21/21 each), raised no pre-fault false alarm, and detected every fault within one fundamental cycle. The mean and maximum detection latencies were 9.95 and 17.45 ms, respectively. Figure 11 shows four representative timelines in which the detection flag and the located label settle within one cycle of the injection instant t_1 .




Figure 11. Online real-time diagnosis under the inductive RL load for four representative faults: (a) S_{a1} ; (b) S_{b1} ; (c) S_{c1} ; (d) S_{b1} and S_{c2} . In the upper trace of each panel, the blue solid line denotes the locating-label output and the red dashed line denotes the binary detection output. The black solid, red dashed, and blue dash-dotted vertical lines mark the fault instant t_1 , confirmed detection t_2 , and locating decision t_3 , respectively; the panel titles and red in-panel annotations list the detection and locating delays relative to t_1 .

The healthy/current-step check was also repeated on the RL bench. During the $2 \rightarrow 3$ A peak phase-current step, the $N_c = 2$ confirmation logic produced zero confirmed false alarms over the recorded ≈ 2 s interval, consistent with the resistive-load false-alarm result in Figure 8.

The retrained RL -load model was then exported to the F28379D project and checked on the target MCU. For this supplementary RL firmware, the embedded diagnostic log indicated correct classification for all 22 clean RL operating modes, including the healthy mode and the 21 fault modes, and the measured minimum/mean/maximum two-stage inference time was 220.24/424.92/516.48 μ s. The corresponding linker map reports about 150.6 kB total program Flash and 3.7 kB RAM, close to the resistive-load firmware footprint in Table 3. These results were obtained without changing the one-cycle feature extraction or the two-stage inference pipeline. The inductive-load evidence is therefore not limited to an off-board MATLAB (2026b) evaluation: the same lightweight diagnosis flow also runs on the control MCU under the supplementary RL load.

5.8. Overall Comparison and Experimental Takeaway

Table 7 compares the two methods. They achieve comparable accuracy on clean data and within the baseline's reported SNR range, but differ substantially in model size, in low-SNR robustness, and, most importantly for EV deployment, in on-board real-time capability.

Representative predefined-feature, data-driven methods [11,12] also tend to degrade under extreme imbalance, as analyzed in [16], and their on-MCU real-time execution is outside the evidence reported in those studies. The comparison table is based on the original LC-filtered resistive benchmark; the supplementary *RL* tests above further show that the proposed feature and embedded pipeline remain effective when the resistor bank is replaced by a balanced inductive load.

Table 7. Overall comparison between the baseline CNN and the proposed method. Bold text indicates the more favorable value where the two methods differ.

Aspect	Baseline CNN [16]	Proposed
Parameters	6.46×10^5	1.58×10^4
MACs per inference	3.31×10^7	1.56×10^4
Clean accuracy (22 modes)	100%	100%
Acc. @ 20 dB (bal./unbal.)	100/100%	100/100%
Acc. @ 10 dB (bal./unbal.)	40.9/50.2%	94.8/63.5%
Noise augmentation in training	not required	not required
On-board MCU real-time performance	No	Yes (0.52 ms)

6. Discussion

The results can be read against the gap identified in Section 1. Prior trajectory-based diagnosis attains high accuracy but classifies a binary image with a large CNN [16–18], which is difficult to run on motor-control MCUs, whereas earlier threshold- and predefined-feature methods avoid that cost yet become fragile under measurement noise and extreme load imbalance. Within the SNR range reported by the baseline, the proposed network matches its locating accuracy while reducing the model size by more than an order of magnitude and running in real time on the control MCU, a combination that those earlier methods did not provide. Two contributions underlie this result. The representation contribution is the angular sector-energy feature: accumulating the current magnitude within each 10° sector averages zero-mean measurement noise, so robustness follows from the representation rather than from noise-augmented training. This sector-wise averaging, together with the sector-count ablation (Table 6), explains why the compact feature is more robust under low SNRs than the sparse binary image, without claiming that feature compression creates new class information. The practical contribution is that the resulting classifier is small enough (1.58×10^4 parameters, 0.52 ms worst case) to be deployed and validated on the control MCU itself, which closes the on-board real-time gap left open by the image-based CNN.

Information-processing view: Reducing the raw α – β trajectory to a 36-D feature is not expected to increase the information available to an ideal classifier. Instead, the proposed mapping imposes an inductive bias matched to noisy embedded diagnosis: isolated sample perturbations are averaged within each angular sector, while the fault-induced missing angular regions are preserved. On clean data both the binary image and sector-energy representations reach 100%, so the benefit is not additional clean-data separability. The two representations differ once measurement noise enters the realized input [16]. In the $D \times D$ binary image, a single perturbed sample can activate a previously empty cell, so isolated noise directly alters the classified pattern. Sector-energy aggregation instead averages each perturbation over the N_k samples of a sector, as quantified in Section 3, and is correspondingly less affected. The benefit thus comes from an inductive bias suited to low-SNR embedded operation.

Limitations: At the most demanding operating point, 10 dB SNR with PCUR = 88%, the locating accuracy falls to 63.5%. This remains above the CNN's 50.2% but is far from ideal. Here the lightly loaded phase carries a current close to the noise level, so the sector

energies along its direction become unreliable. The same regime is identified in [16]: as PCUR approaches 100%, the healthy and faulty waveforms of the affected leg become nearly indistinguishable, which is a limitation of the measurement rather than of the proposed feature. For a safety-critical drive, however, the first-line protective response is triggered by fault detection rather than by exact switch localization: binary detection is considerably easier than 22-way locating, and the N_c -decision confirmation can reject isolated noisy decisions. A persistent fault is therefore still flagged at this operating point. The part that degrades at this extreme corner, which combines 10 dB noise with PCUR = 88% beyond the baseline's reported range, is the precise localization that subsequently refines the fault-tolerant reconfiguration. Thus, the 63.5% locating accuracy should be interpreted as an offline stress-test result, not as a stand-alone basis for an irreversible safety action without persistent fault confirmation and a conservative fallback or degraded operating mode. At the more typical ≥ 20 dB, the locating accuracy is 100%, even under 88% imbalance. Two further constraints apply to the present study. The initial detection latency, although within one fundamental cycle, varies with the fault-injection phase. The confirmed alarm also depends on the update interval and the confirmation count N_c . The test bench is restricted to a 50 Hz fundamental and a single power level.

Fault scope and current-sensor faults: The proposed network is a closed-set classifier trained for the 22 switch states under healthy current sensors, and it targets switch OC faults rather than sensor faults. The two fault types leave distinct signatures: a switch OC fault removes only one current polarity of the affected phase, so the α - β trajectory stays two-dimensional but becomes asymmetric, whereas a complete current-sensor fault (open or disconnected) drives the measured phase toward zero in both half-cycles and collapses the trajectory along that phase axis; the two are therefore separable in principle. Because the present classifier is not trained on sensor faults, such a fault would be forced into one of the 22 classes. A plausibility check on the measured currents, a joint switch/sensor diagnosis [12], or an open-set rejection stage [24] would be needed to label it explicitly, which we leave to future work.

Generalization to inductive loads and traction-motor operation: The experiments here use resistive loads and a balanced inductive RL load at a fixed 50 Hz, whereas an EV traction drive is strongly inductive, runs at a variable speed, and is subject to motor back-EMF. The proposed feature is expected to be less sensitive to load-amplitude changes than amplitude-dominated features, because the per-cycle normalization makes it describe the angular distribution of the α - β current trajectory. An inductive lagging load primarily rotates and reshapes the trajectory through the power-factor angle and the load impedance. It does not, in principle, remove the angular energy gap created by the blocked current polarity of an OC fault. This expectation is supported within the present evidence: across the balanced and strongly unbalanced (PCUR = 88%) resistive conditions and the inductive RL load of Section 5.7, the method retains 100% clean accuracy (and $\geq 99.98\%$ at 10 dB on the RL load). A strongly inductive, variable-speed traction motor nonetheless introduces back-EMF, magnetic dynamics, current-control dynamics, speed variation, and mechanical transients, which may cause a distribution shift relative to the bench data used here. Variable-speed operation could be handled by a frequency-synchronous window over one electrical period and, if necessary, by resampling the trajectory or the sector energies onto the same 36-sector angular grid. A full validation on a motor-driven bench is therefore the immediate next step; if a residual distribution shift remains, the network can be retrained on motor data without changing the feature definition or the on-MCU inference pipeline.

Implications and future work: Because the diagnosis runs on the existing control MCU, a localized fault can feed directly into on-board fault-tolerant control without additional hardware. Once the faulty switch is identified, the controller can reconfigure the mod-

ulation or compensate the faulty-phase current to sustain limp-home operation [4,25]. The computational cost is modest, below 3% of one fundamental cycle, and the memory footprint is about 15% of Flash. This leaves a timing margin for the motor-control task on CPU1, and the diagnosis can also be moved to the second C28x core. Three extensions follow from the present setup. Operation at a fixed 50 Hz and a single power level should be generalized to the variable-speed, variable-load conditions of traction drives. Multilevel and other topologies, whose current-vector trajectories are more similar across modes, may need richer angular features than the 36 sectors used here. Combining the diagnosis with fault-tolerant control [25] and with an uncertainty- or unknown-fault-rejection stage [24] would extend it toward open-world EV operation.

7. Conclusions

This paper represented the α - β current vector of a three-phase inverter as a 36-dimensional normalized angular sector-energy feature and classified it with a compact two-stage MLP for open-circuit fault detection and switch locating. Aggregating the current magnitude within each angular sector attenuates zero-mean measurement noise. This explains why the method does not need noise-augmented training and why it is less sensitive to low-SNR perturbations than the sparse binary trajectory image. On identical real-hardware resistive-load data, the method reproduced the performance of the baseline over the 40–20 dB range (100% under both balanced and PCUR = 88% loading) and retained higher locating accuracy at 15 and 10 dB, where the CNN accuracy decreased markedly. It also used 97.55% fewer parameters than the baseline CNN and about three orders of magnitude fewer MACs. On the TI F28379D control MCU, the deployed model identified all 22 operational modes on clean measured data with a worst-case two-stage inference time of 0.52 ms. The detector raised no false alarms under healthy operation or a 2 → 3 A peak (+50%) current-step test (Figure 8 and Section 5.7). The supplementary *RL*-load experiment further showed that the same one-cycle feature and two-stage pipeline retain 100% clean accuracy and can be executed on the MCU under a balanced lagging power-factor load. The contribution is therefore twofold: a noise-averaging angular feature whose robustness is explained by a sector-wise perturbation analysis and corroborated by a sector-count ablation, and a diagnosis pipeline light enough to run on the motor-control MCU itself. The present study also has several limitations. The main validation is based on a 60 V, 50 Hz inverter bench with resistive loads, supplemented by a balanced *RL*-load test; therefore, the results do not yet constitute full traction-motor validation under variable speed, variable torque, motor back-EMF, and closed-loop drive transients. In addition, the classifier is trained as a closed-set switch OC diagnostic model under healthy current sensors, so explicit current-sensor faults and unknown fault modes require additional plausibility checks, joint diagnosis, or open-set rejection. Future studies will therefore extend the validation to motor-driven benches and variable operating conditions, examine multilevel and other inverter topologies, and integrate the diagnosis with fault-tolerant control and uncertainty-aware decision logic.

Author Contributions: Conceptualization, M.F. and W.W.; methodology, M.F.; software, M.F.; validation, M.F. and W.Y.; formal analysis, M.F.; investigation, M.F.; resources, W.W. and W.Y.; data curation, M.F.; writing—original draft preparation, M.F.; writing—review and editing, W.W. and W.Y.; supervision, W.W.; project administration, W.W.; funding acquisition, W.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China, grant number 62101215.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

OC	Open-circuit
EV	Electric vehicle
VSI	Voltage-source inverter
CNN	Convolutional neural network
MLP	Multilayer perceptron
SNR	Signal-to-noise ratio
PCUR	Phase-current unbalance rate
MAC	Multiply–accumulate
MCU	Microcontroller unit
TMU	Trigonometric math unit
ADC	Analog-to-digital converter
DSP	Digital signal processor
FPU	Floating-point unit
IGBT	Insulated-gate bipolar transistor
PWM	Pulse-width modulation
V2G	Vehicle-to-grid
WGN	White Gaussian noise

References

1. Wang, H.; Liserre, M.; Blaabjerg, F. Toward Reliable Power Electronics: Challenges, Design Tools, and Opportunities. *IEEE Ind. Electron. Mag.* **2013**, *7*, 17–26. <https://doi.org/10.1109/MIE.2013.2252958>.
2. Amry, Y.; Elbouchikhi, E.; Le Gall, F.; Ghogho, M.; El Hani, S. Electric Vehicle Traction Drives and Charging Station Power Electronics: Current Status and Challenges. *Energies* **2022**, *15*, 6037. <https://doi.org/10.3390/en15166037>.
3. Wei, X.; Yao, B.; Peng, Y.; Sun, Y.; Wang, K.; Wang, H. An Improved Discharge Profile-Based DC-Link Capacitance Estimation for Traction Inverter in Electric Vehicle Applications. *IEEE Trans. Power Electron.* **2024**, *39*, 8696–8708. <https://doi.org/10.1109/TPEL.2024.3383153>.
4. Choi, U.-M.; Blaabjerg, F.; Lee, K.-B. Study and Handling Methods of Power IGBT Module Failures in Power Electronic Converter Systems. *IEEE Trans. Power Electron.* **2015**, *30*, 2517–2533. <https://doi.org/10.1109/TPEL.2014.2373390>.
5. Rothenhagen, K.; Fuchs, F.W. Performance of diagnosis methods for IGBT open-circuit faults in three-phase voltage source inverters for AC variable speed drives. In Proceedings of the European Conference on Power Electronics and Applications (EPE), Dresden, Germany, 11–14 September 2005; pp. 1–10. <https://doi.org/10.1109/EPE.2005.219426>.
6. Guo, H.; Guo, S.; Xu, J.; Tian, X. Power Switch Open-Circuit Fault Diagnosis of Six-Phase Fault-Tolerant PMSM System Using the Average Current Park's Vector Approach. *IEEE Trans. Power Electron.* **2021**, *36*, 2641–2660. <https://doi.org/10.1109/TPEL.2020.3017637>.
7. Zhang, Y.; Liu, Y.; Kang, S.; Wang, P. Current Vector Phase Based Weak Open-Circuit Fault Diagnosis of Voltage-Source Inverters. *IEEE Trans. Power Electron.* **2026**, *41*, 9502–9512. <https://doi.org/10.1109/TPEL.2024.3478766>.
8. Wang, B.; Li, Z.; Bai, Z.; Krein, P.T.; Ma, H. A Voltage Vector Residual Estimation Method Based on Current Path Tracking for T-Type Inverter Open-Circuit Fault Diagnosis. *IEEE Trans. Power Electron.* **2021**, *36*, 13460–13477. <https://doi.org/10.1109/TPEL.2021.3087488>.
9. Li, G.; Xu, S.; Sun, Z.; Yao, C.; Ren, G.; Ma, G. Open-Circuit Fault Diagnosis for Three-Level ANPC Inverter Based on Predictive Current Vector Residual. *IEEE Trans. Ind. Appl.* **2023**, *59*, 6837–6851. <https://doi.org/10.1109/TIA.2023.3310491>.

10. Diao, N.; Zhang, Y.; Sun, X.; Song, C.; Wang, W.; Zhang, H. A Real-Time Open-Circuit Fault Diagnosis Method Based on Hybrid Model Flux Observer for Voltage Source Inverter Fed Sensorless Vector Controlled Drives. *IEEE Trans. Power Electron.* **2023**, *38*, 2539–2551. <https://doi.org/10.1109/TPEL.2022.3216870>.
11. Cai, B.; Zhao, Y.; Liu, H.; Xie, M. A Data-Driven Fault Diagnosis Methodology in Three-Phase Inverters for PMSM Drive Systems. *IEEE Trans. Power Electron.* **2017**, *32*, 5590–5600. <https://doi.org/10.1109/TPEL.2016.2608842>.
12. Gou, B.; Xu, Y.; Xia, Y.; Deng, Q.; Ge, X. An Online Data-Driven Method for Simultaneous Diagnosis of IGBT and Current Sensor Fault of Three-Phase PWM Inverter in Induction Motor Drives. *IEEE Trans. Power Electron.* **2020**, *35*, 13281–13294. <https://doi.org/10.1109/TPEL.2020.2994351>.
13. Yan, H.; Peng, Y.; Shang, W.; Kong, D. Open-Circuit Fault Diagnosis in Voltage Source Inverter for Motor Drive by Using Deep Neural Network. *Eng. Appl. Artif. Intell.* **2023**, *120*, 105866. <https://doi.org/10.1016/j.engappai.2023.105866>.
14. Das, P.P.; Satpathy, S.; Bhattacharya, S. An Online Open-Circuit Fault Diagnosis Technique for Three-Level Inverter-Fed Six-Phase PMSM Drives. *IEEE Trans. Power Electron.* **2024**, *39*, 14974–14987. <https://doi.org/10.1109/TPEL.2024.3440641>.
15. Chu, K.S.K.; Chew, K.W.; Chang, Y.C.; Morris, S. An Open-Circuit Fault Diagnosis System Based on Neural Networks in the Inverter of Three-Phase Permanent Magnet Synchronous Motor (PMSM). *World Electr. Veh. J.* **2024**, *15*, 71. <https://doi.org/10.3390/wevj15020071>.
16. Lei, X.; Wu, F.; Liu, Y. An Online Convolutional Neural Network Based Method for Open-Circuit Fault Diagnosis in Three-Phase Inverters Under Extremely Unbalanced Loading Condition. *IEEE Trans. Power Electron.* **2026**, 1–14, early access. <https://doi.org/10.1109/TPEL.2026.3669457>.
17. Hang, J.; Shu, X.; Ding, S.; Huang, Y. Robust Open-Circuit Fault Diagnosis for PMSM Drives Using Wavelet Convolutional Neural Network with Small Samples of Normalized Current Vector Trajectory Graph. *IEEE Trans. Ind. Electron.* **2023**, *70*, 7653–7663. <https://doi.org/10.1109/TIE.2022.3231304>.
18. Sun, X.; Diao, N.; Song, C.; Qiu, Y.; Zhao, X. An Open-Circuit Fault Diagnosis Method Based on Adjacent Trend Line Relationship of Current Vector Trajectory for Motor Drive Inverter. *Machines* **2023**, *11*, 928. <https://doi.org/10.3390/machines11100928>.
19. Chen, L.; An, K.; Huang, D.; Wang, X.; Xia, M.; Lu, S. Noise-Boosted Convolutional Neural Network for Edge-Based Motor Fault Diagnosis with Limited Samples. *IEEE Trans. Ind. Inform.* **2023**, *19*, 9491–9502. <https://doi.org/10.1109/TII.2022.3228902>.
20. Han, S.; Zhang, S.; Lv, W. Fault Diagnosis of Grid-Connected NPC Inverter Based on Information Fusion and Lightweight Neural Network. *Electr. Power Syst. Res.* **2025**, *247*, 111835. <https://doi.org/10.1016/j.epsr.2025.111835>.
21. Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; Chen, L.-C. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Salt Lake City, UT, USA, 18–23 June 2018; pp. 4510–4520. <https://doi.org/10.1109/CVPR.2018.00474>.
22. Iandola, F.N.; Han, S.; Moskewicz, M.W.; Ashraf, K.; Dally, W.J.; Keutzer, K. SqueezeNet: AlexNet-Level Accuracy with 50× Fewer Parameters and <0.5 MB Model Size. *arXiv* **2016**, arXiv:1602.07360. <https://doi.org/10.48550/arXiv.1602.07360>.
23. Abadade, Y.; Temouden, A.; Bamoumen, H.; Benamar, N.; Chtouki, Y.; Hafid, A.S. A Comprehensive Survey on TinyML. *IEEE Access* **2023**, *11*, 96892–96922. <https://doi.org/10.1109/ACCESS.2023.3294111>.
24. Wei, D.; Zuo, M.; Tian, Z. Open-Set Fault Diagnosis for Industrial Rotating Machines Based on Trustworthy Deep Learning. *IEEE Trans. Ind. Cyber-Phys. Syst.* **2025**, *3*, 181–189. <https://doi.org/10.1109/TICPS.2025.3539997>.
25. Chen, C.; Sun, X.; Chen, Z.; Liu, X.; Zhu, Z.Q.; Liao, X. Fault-Tolerant Control with Faulty Phase Winding Currents Compensation for Current Source Inverter-Fed Five-Phase PMSM Under Adjacent Double-Phase Open-Circuit Fault. *IEEE Trans. Transp. Electr.* **2025**, *11*, 7697–7709. <https://doi.org/10.1109/TTE.2025.3530522>.

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.