FE-Based Data Generation for Structural Integrity Verification with Neural Networks

Maximilian Posner and Martin Dazer

Institute for Machine Components, University of Stuttgart, Stuttgart, 70569, Germany
maximilian.posner@ima.uni-stuttgart.de
martin.dazer@ima.uni-stuttgart.de

ABSTRACT

The application of deep learning within Prognostics and Health Management (PHM) frameworks offers immense potential for real-time structural health monitoring. However, in high-consequence applications such as high-velocity collisions, algorithmic deployment is severely bottlenecked by a fundamental scarcity of training data. Because destructive physical testing is prohibitively expensive and real-world incident telemetry is rare, acquiring the massive, diverse datasets required to prevent neural network overfitting remains a critical operational challenge.

To overcome this data scarcity and mitigate the resulting curse of dimensionality, this paper presents a scalable methodology that transforms high-fidelity Finite Element (FE) analysis into an automated Physics-of-Failure "data engine". Using the safety-critical underrun guard of a high-speed train as a primary case study, a complete Digital Thread is constructed. Complex impact kinematics are mathematically distilled, and six non-linear material parameters are algorithmically condensed into a singular "strength scaling factor" SSF variable. Utilizing a low-discrepancy Sobol sequence, 256 transient-dynamic simulations are sampled uniformly, ensuring highly sensitive boundary transitions of structural failure are densely captured without statistical clustering.

To simulate real-world hardware constraints, a virtual sensor network is deployed, with multi-axial tensors mathematically reduced to scalar invariants to optimize the Machine Learning (ML) input layer. Concurrently, standardized visual renderings enable domain experts to label physical damage into discrete operational directives. Crucially, to further combat data scarcity, the framework expands standard ML training by introducing an asymmetric training paradigm that exploits "hidden" FE simulation data, such as continuous plastic strain measurements. While physically unmeasurable in real-world scenarios, these metrics are utilized exclusively during training as auxiliary targets to structurally regularize the network, before being completely discarded for real-world inference.

Following the experimental validation of the digital twin against physical destructive laboratory tests, the resulting synthetic dataset was utilized to train a baseline classification neural network. Despite the severely constrained sample size ($n = 256$), the diagnostic ML achieved an outstanding predictive accuracy of 82%. This confirms that mathematically optimized data generation and dimensionality reduction, combined with the strategic extraction of auxiliary simulation data, can successfully bridge the data gap in modern PHM frameworks, providing a robust pathway for real-time diagnostic deployment.

1. INTRODUCTION

Assessing structural integrity following high-energy impacts is a critical operational challenge across modern transportation and aerospace sectors (Wu, Cantero-Chinchilla, Prescott, Remenyte-Prescott, & Chiachío, 2024). High-velocity collisions induce large-strain plastic deformations, characterized by highly non-linear material behavior, strain-rate dependencies, and complex multi-body contact physics (Pore, Thorat, & Nema, 2021). Because traditional analytical models and linear structural health monitoring (SHM) methodologies struggle to accurately capture these extreme transient dynamics (Avevor, Adeniyi, Enyejo, & Aikins, 2024), PHM frameworks are increasingly pivoting toward Industry ML (Ochella, Shafiee, & Dinmohammadi, 2022). Deep neural networks excel at learning complex mapping functions, correlating abstract, high-frequency sensor telemetry (e.g., from strain gauges and accelerometers) with specific, expert-validated damage states (Kashiparekh, Narwariya, Malhotra, Vig, & Shroff, 2019). Once trained, an AI-driven diagnostic system can evaluate an impact event in real time, translating raw telemetry into actionable operational instructions, such as authorizing continued service or mandating an immediate emergency stop.

However, in modern PHM applications involving extreme physical environments, the algorithm is not the only bottleneck. The other bottleneck is the acquisition of high-quality, physics-grounded training data (Akrim, Gogu, Vingerhoeds, & Salaün, 2023). Deep learning architectures require massive, highly diverse datasets across a broad parameter space to generalize effectively and avoid overfitting (Peng, Gui, & Wu, 2024). Collecting this data through physical means is universally unfeasible. Across all heavy industries, destructive physical testing of full-scale assemblies is prohibitively expensive, time-consuming, and strictly limited to a handful of design-validation scenarios (Fang, Otero, Greenberg, & Neumann, 2011). Furthermore, relying on historical real-world incident data is insufficient: severe collisions are fortunately rare, and the requisite high-frequency sensor telemetry is often not installed on current or past vehicles (Zhang et al., 2024).

To overcome this data scarcity, the primary objective of this paper is to outline a generalized framework that transitions automated FE simulation from a traditional design tool into an active "data engine" for model-based prediction (Zhang et al., 2024). By automating high-fidelity, non-linear FE pipelines, a Physics-of-Failure Digital Twin can be constructed to virtually execute hundreds of collision scenarios without destroying physical hardware (Zhang et al., 2024). To ensure this automated data generation is mathematically efficient for downstream machine learning, standard random sampling methods (such as Monte Carlo) are insufficient due to uneven distribution of data points and clustering in high-dimensional spaces (Chopin, 2004). Instead, low-discrepancy sampling techniques, such as the Sobol sequence (Sobol', 1967), must be implemented to uniformly cover the whole parameter space. This algorithmic optimization guarantees that the Digital Twin efficiently captures the highly sensitive boundary transitions between minor elastic deformation and localized, catastrophic structural failure (Saltelli et al., 2010).

To demonstrate and validate this scalable methodology, this paper presents a comprehensive case study on a highly complex, safety-critical component: the underrun guard of a German high-speed intercity train (ICE). Engineered to absorb immense kinetic energy and prevent derailments, this structure serves as an ideal subject for establishing a "Digital Thread." A fundamental challenge in post-crash diagnostics is that complex, large-scale plastic deformations are notoriously difficult to describe using purely quantitative numerical metrics. In practice, rating the severity of a collision and determining the necessary operational or maintenance response relies heavily on qualitative expert knowledge and visual inspection (DIN EN 15227, 2020).

Therefore, a critical challenge addressed in this work is translating this qualitative human expertise into a structured format comprehensible to Machine Learning algorithms. By integrating validated, strain-rate-dependent virtual sensors with an expert-driven visual labeling interface, this methodology successfully transforms raw, high-frequency time-series simulation data into a rich, labeled dataset for supervised learning. Ultimately, the overarching goal of this paper is to outline a generalized framework that bridges the gap between quantitative, automated FE data generation and human-centric operational decision-making by solving two distinct challenges: the curse of dimensionality and the underutilization of simulation data.

While recent model-based prediction frameworks successfully automate structural dynamics simulations, they traditionally treat the finite element (FE) solver merely as a virtual replica of the field environment, entirely discarding the vast volume of internal, continuous physical data calculated during an event (such as localized plastic strain distributions, stress fields, and element-level damage counts) simply because they cannot be measured by real-world sensors. To significantly amplify data efficiency in sparse-label regimes, the pipeline proposed in this paper diverges from this conventional methodology by actively exploiting this "privileged" or "hidden" simulation data. Drawing on asymmetric training concepts, these unmeasurable internal metrics are explicitly extracted to serve as secondary regression targets exclusively during the neural network's training phase. By forcing the model's latent representation to concurrently learn these underlying continuum mechanics alongside the primary classification task, the network's hypothesis space receives deep, physically grounded regularization that prevents overfitting to statistical noise. During real-world deployment, these auxiliary data requirements and regression heads are entirely discarded, leaving a lightweight, highly accurate diagnostic model capable of performing robust predictions using only standard, sparse sensor telemetry.

The following sections detail the construction of the Physics-of-Failure data engine, the algorithmic reduction of the sampling space, the virtual telemetry and auxiliary data framework, and the expert classification pipeline required to generate datasets ready for modern foundation models.

2. STATE OF THE ART

Finite Element (FE) analysis is firmly established as the standard methodology for evaluating the structural dynamics and crashworthiness of safety-critical components under extreme loading (Pore et al., 2021). In domains such as aerospace and high-speed rolling stock (Morris, 2009), accurately modeling the dissipation of massive kinetic energy requires capturing severe non-linearities, such as strain-rate-dependent plasticity (e.g., the Johnson-Cook material model (Mareau, 2020)) and complex self-contact friction algorithms. While historically restricted to deterministic design-validation cases due to computational burdens, these FE analyses are increasingly utilized to produce virtual sensor telemetry, such as strain and acceleration time-series data, reproducing real-world physical anomalies within a controlled virtual space (Avevor et al., 2024). To translate these complex, abstract sensor readings into reliable operational diagnostics, modern Prognostics and Health Management (PHM) frameworks increasingly rely on Deep Learning (DL) (Ochella et al., 2022). As universal function approximators, artificial neural networks possess the unique capability to learn high-dimensional mappings between chaotic time-series inputs and discrete damage states without requiring hand-crafted features (Kashiparekh et al., 2019).

Despite this architectural flexibility, deep learning introduces a severe vulnerability in safety-critical applications: overfitting (Peng et al., 2024). Overfitting occurs when a neural network possesses excessive model capacity relative to the quantity or diversity of its training data. Instead of learning the underlying physical relationships and structural invariant laws governed by continuum mechanics, the network "memorizes" random noise, specific geometric configurations, or statistical artifacts present within a sparse dataset (Akrim et al., 2023). Consequently, when deployed against unseen collision scenarios, the model's predictive accuracy collapses. Mathematically, ensuring generalizability requires a massive, diverse dataset that scales exponentially with the number of input parameters, a phenomenon known as the curse of dimensionality (Peng et al., 2024). Because empirical physical data collection for extreme impact events is economically and ethically unfeasible, automated data generation via simulation serves as the only viable pathway to achieve the necessary data density (Fang et al., 2011).

Finally, the computational feasibility of any advanced FE dataset generation campaign is strictly dictated by the strategy used to sample the high-dimensional collision parameter space (e.g., velocity, impact coordinates, material variations) (Fang et al., 2011). Traditional dataset generation methods frequently employ standard Monte Carlo (MC) random sampling. However, pseudo-random MC sampling is prone to clustering effects, which create redundant data points in certain regions of the parameter space while leaving other critical zones completely unsampled (Zhang et al., 2024). Conversely, deterministic grid-based sampling scales poorly as dimensionality increases, causing an immediate computational bottleneck (Zhang et al., 2024). To resolve this tradeoff, recent model-based prediction frameworks utilize Quasi-Monte Carlo methods, specifically Sobol sequences (Sobol', 1967). A Sobol sequence is a deterministic, low-discrepancy sequence designed to construct a highly uniform distribution of points across a multi-dimensional hypercube (Sobol', 1967). Unlike pseudo-random distributions, a Sobol sequence inherently "remembers" where previous points were placed, positioning successive points as far away from existing ones as possible (Bratley & Fox, 1988). For non-linear structural failure regimes, this uniformity is indispensable. It guarantees that the highly sensitive, narrow boundary transitions separating minor elastic deformation from catastrophic material fracturing are sampled densely and uniformly, ensuring rapid ML convergence without wasting computational resources on redundant simulations (Saltelli et al., 2010).

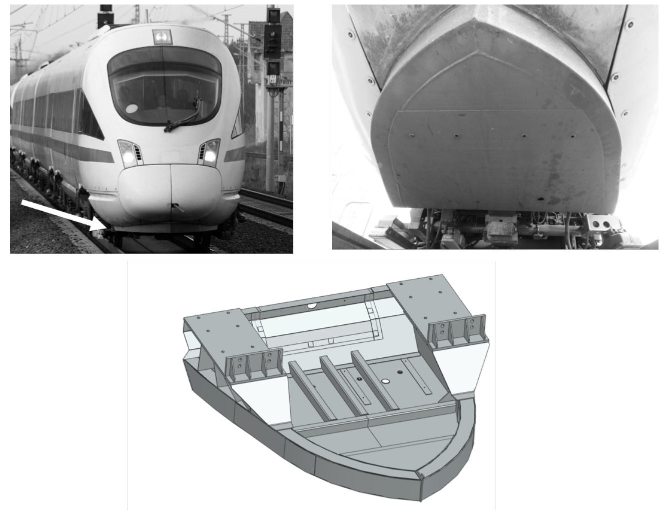

Figure 1. Underrun Guard of a German High Speed Train (ICE) in three perspectives, including the CAD Model (Terfloth, 2008)

3. CONSTRUCTING THE PHYSICS-OF-FAILURE DATA ENGINE

To construct a high-fidelity Physics-of-Failure data engine capable of training a diagnostic neural network, the physical design space must be strictly bounded. A boundless simulation environment would lead to an infinite parameter space, instantly triggering the curse of dimensionality during machine learning training. This section details the strategic bounding of the Digital Twin through the physical simplification of the collision event, the algorithmic reduction of the material property space, and the definition of the operational boundaries used to initialize the quasi-random Sobol sequence generator.

The physical domain digitized in this study is the underrun guard of a German ICE, a highly optimized, complex steel structure designed to protect the vehicle against track-bound foreign objects. Mounted directly beneath the front nose of the train in close proximity to the tracks, it serves the critical safety functions of clearing obstacles, absorbing catastrophic impact energy, and preventing train derailment. Figure 1 illustrates the physical deployment and the corresponding Computer-Aided Design model of the structure.

To further accelerate the dataset generation and isolate the localized structural degradation, the complete vehicle geometry is excluded from the Digital Twin. Instead, only the CAD model of the underrun guard is explicitly meshed for the finite element analysis. To strictly enforce realistic boundary conditions and compensate for the physical absence of the train chassis, an equivalent inertial mass of 230 metric tons is mathematically coupled directly to the structural bolt mounting points at the rear of the underun guard, to simulate the train's dynamics. This kinematic abstraction guarantees accurate macroscopic kinetic energy transfer during the impact while drastically reducing the total computational node count of the solver.

3.1. Bounding the Digital Twin - Collision Simplification

The operational velocity of a high-speed train ($v_{train} \leq 90$ m/s) heavily outweighs the cross-track velocity component of any typical moving obstacle. The physical response of the guard is therefore dominated by the train's longitudinal momentum. Relying on Galilean invariance and classical Newtonian mechanics, the complex multi-body kinematics can be mathematically distilled into a single relative state without losing high-fidelity impact dynamics (Goldsmith, 2001). The impact depends solely on the relative velocity vector $\vec{v}_{rel}$ at the exact instant of contact:

$$\vec{v}_{\text{rel}} = \vec{v}_{\text{train}} - \vec{v}_{\text{object}} \tag{1}$$

Because this relative kinematic framework entirely governs the momentum exchange, it is mathematically irrelevant whether the train or the obstacle is designated as the primary moving body. However, optimizing a Finite Element pipeline for high-throughput dataset generation requires ruthlessly minimizing computational overhead. Therefore, the simulation architecture reverses the real-world kinematic state: the geometrically massive, highly meshed train assembly is constrained as a stationary boundary, while the smaller collision object is propelled toward it at $\vec{v}_{rel}$. This kinematic inversion significantly reduces the spatial coordinate updates required by the solver at every time step, further accelerating the data engine.

By defining the longitudinal axis (x-axis) as the primary vector of kinetic energy exchange, any minor lateral or diagonal velocity from a crossing obstacle can be mathematically substituted by scaling the nominal train velocity variable within its operational envelope. Furthermore, an axial collision with a moving obstacle (such as a shunting wagon) is mechanically equivalent to a static impact configuration operating at an adjusted relative velocity, governed by the conservation of linear momentum:

$$m_1\vec{v}_1 + m_2\vec{v}_2 = (m_1 + m_2)\vec{v}_{\text{post}} \tag{2}$$

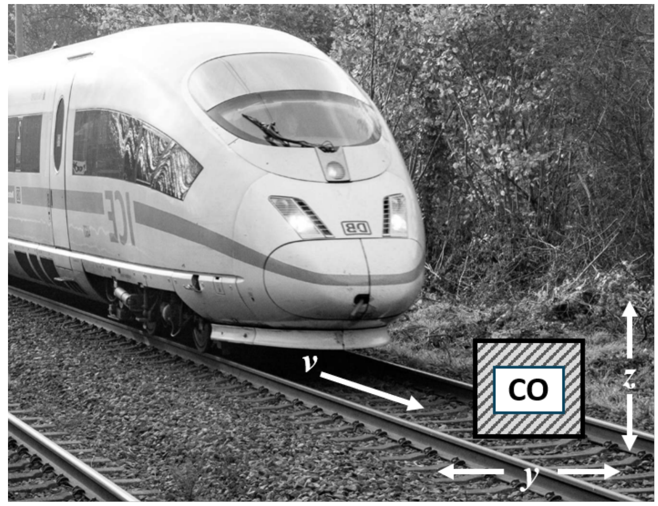

Figure 2. Collision Event Parameters: Relative speed and the 2D Posotioning of the Collision Object (CO) (24RHEIN, 2024)

Crucially for the downstream ML architecture, these physics-based equivalencies allow the FE simulation to be collapsed into a pure frontal crash scenario. This completely eliminates the longitudinal distance vector (x) as an independent input variable, as it merely dictates the pre-contact time-of-flight.

Consequently, the spatial variability of the entire collision dataset is rigorously bounded to a two-dimensional plane perpendicular to the direction of travel, defined exclusively by the lateral (y) and vertical (z) impact coordinates on the structure. By reducing the physical event to three continuous parameters, lateral position, vertical position, and relative impact speed (v), the input feature space is optimally constrained for dense, uniform sampling. These simplified boundary parameters are visualized in Figure 2.

3.2. Material Parameter Dimensionality Reduction

Accurately capturing high-energy impact physics up to structural failure requires modeling both the initial elastic-plastic transition and the progressive macro-void shearing that initiates structural fracture. A standard high-fidelity material formulation for structural steel requires at least six independent parameters to describe this transient behavior: Young's Modulus (E), density (ρ), Poisson's ratio (ν), yield strength (σ_y), ultimate tensile strength (σ_{uts}) and the maximum failure plastic strain ($\varepsilon_{p,\text{max}}$) (Johnson & Cook, 1983; Altair Engineering Inc., 2023).

In traditional Monte Carlo simulation, each of these variables might be sampled independently. However, within an automated data engine designed to train deep neural networks, increasing the dimensionality of the parameter space imposes a severe computational penalty. A high-dimensional space causes a geometric explosion in the number of required finite element runs, fundamentally degrading the space-filling efficiency of low-discrepancy sampling.

To prevent this disadvantage of high dimensionality, the material property space must be algorithmically reduced. From a PHM and diagnostic perspective, forcing the downstream ML to isolate independent micro-metallurgical variables is unnecessary. The diagnostic neural network does not need to deduce the exact microstructural composition of the steel, it only needs to classify whether the macroscopic underrun guard behaves as a structurally "weak" or "tough" variant within known operational bounds. Consequently, this methodology condenses the six-dimensional material space into a singular, dimensionless material parameter SSF (ξ) that represents the strength-ductility trade-off. As detailed in the subsequent scaling equations, this single parameter deterministically controls all six material properties simultaneously. Conceptually, ξ is defined over the bounded domain:

$$\xi \in [-1, 1] \tag{3}$$

To establish how this parameter leads to the final material states, a baseline mechanical reference configuration must first be anchored to the nominal values for standard S235JR structural steel (EN 10025-2, 2019). The physical variability applied to this baseline is governed by empirical deviation ranges, representing the statistically expected variance under standard industrial manufacturing and processing tolerances (Agnew & Duygulu, n.d.). These parameters are detailed in Table 1.

Table 1. Baseline Mechanical Properties (European Committee for Standardization (CEN), 2019) and Statistical Deviation Boundaries (Joint Committee on Structural Safety (JCSS), 2001; Hess et al., 2002) for S235JR Structural Steel

Parameter	**Symbol**	**Unit**	**Baseline Value**	**Deviation Range**
Young's Modulus	E	kN/mm^2	210	0.0
Density	ρ	kg/mm^3	7.85×10^{-6}	0.0
Poisson's Ratio	ν	—	0.3	0.0
Yield strength	σ_y	kN/mm^2	0.235	0.1
Ultimate Strength	σ_{uts}	kN/mm^2	0.360	0.08
Maximum Plastic Strain	$\varepsilon_{p,\max}$	—	0.260	0.2

To simulate this physical scatter within a unified algorithmic framework, the parameters are scaled deterministically according to their sensitivity to manufacturing deviations. Because the elastic constants (E, ρ, ν) exhibit negligible statistical scatter in standardized steel, they remain constant (range = 0). Conversely, the strength and failure parameters are dynamically coupled to the SSF ξ. During the initialization of each finite element run, the structural parameters are updated according to the following linear scaling equations:

$$\sigma_y^* = \sigma_y \cdot (1 + \text{range}_{\sigma_y} \cdot \xi) \tag{4}$$

$$\sigma_{\text{uts}}^* = \sigma_{\text{uts}} \cdot (1 + \text{range}_{\sigma_{\text{uts}}} \cdot \xi) \tag{5}$$

$$\varepsilon_{p,\max}^* = \varepsilon_{p,\max} \cdot (1 - \text{range}_{\varepsilon} \cdot \xi) \tag{6}$$

Crucially, the mathematical formulation of these equations reflects a fundamental physical invariant in metallurgy: the inverse correlation between a material's structural strength and its ultimate ductility (Morris, 2009). When microstructural mechanisms, such as strain hardening or localized grain refinement, shift a steel alloy toward higher yield and ultimate strengths, the material inherently suffers a reduction in its uniform elongation limit before fracturing. To honor this physical trade-off, the SSF ξ acts symmetrically but in opposing directions for strength versus strain.

Through this operational principle, the single parameter ξ continuously maps the physical trade-off between material strength and ductility. A value of $\xi = -1$ configures the solver at the lower statistical boundary of strength, representing a highly ductile but mechanically weak component susceptible to excessive plastic deformation and early yielding. Conversely, a value of $\xi = 1$ shifts the assembly to its upper strength boundary, modeling a highly rigid component with low failure-strain tolerance that is prone to abrupt brittle fracture. Consequently, the nominal state near $\xi = 0$ represents the optimally balanced material profile, offering the highest functional energy absorption before either extreme failure mode occurs. By enforcing this deterministic physical coupling, the framework successfully collapses the material space from six independent dimensions down to one, permanently bounding the sampling space and ensuring rapid, artifact-free convergence for the diagnostic machine learning architecture.

As illustrated in Figure 3, the relationship between the SSF and the normalized accumulated plastic strain demonstrates two distinct regimes of material behavior. Each data point represents a single simulation event initialized with a unique SSF and otherwise identical simulation parameters. The vertical axis represents the total plastic strain summed across all finite elements in the model. As the SSF increases from -1.0, the accumulated strain initially exhibits a steady, linear decline, reflecting the material's increasing yield strength and enhanced resistance to ductile deformation. However, beyond a SSF threshold of approximately 0.4, the strain abruptly jumps to a maximum plateau, capturing the transi-

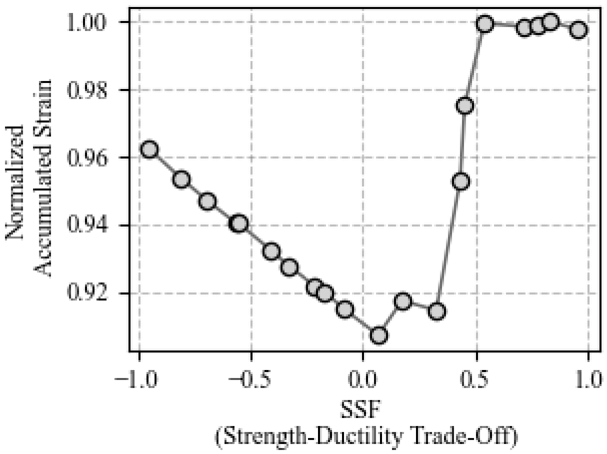

Figure 3. Algorithmic Reduction of Material Deviations into a Unified Trade-off Factor

tion into brittle failure as the maximum plastic strain capacity is severely constrained. This clear, predictable bifurcation, capturing both gradual yielding and sudden fracture, demonstrates that coupling the underlying structural parameters to a single scalar successfully captures the physical spectrum of material deviation. This validates the approach of reducing a complex six-parameter space into one unified variable to efficiently simulate variations in overall material quality.

3.3. Definition of the 4-Dimensional Sampling Space

Through the physics-based simplifications and material property reductions outlined in the preceding sections, the structural data generation problem is rigorously compacted into a concise four-dimensional parameter space ($s = 4$). The upper and lower bounds for each design variable, summarized in 2, are strictly selected based on structural geometry, rolling-stock operational envelopes, and the typical spatial profiles of standard track obstacles:

- Lateral Position (y): Varies from 0 mm (the exact geometric centerline of the underrun guard) to -700 mm (the outer structural edge). This range accounts for the spatial width and offset of typical obstacle profiles.
- Vertical Position (z): Varies from 0 mm (the lower clearance limit of the shield) to 210 mm (the upper geometric boundary where the guard interfaces with the main chassis).
- Relative Impact Velocity (v): Bounded by a minimum speed greater than zero ($v_{\text{min}} > 0$) to capture low-energy plastic deformation thresholds, up to the maximum operational limits of high-speed rolling stock ($v_{\text{max}} = 90$ m/s or 324 km/h).
- SSF (ξ): Bounded between -1 and 1, driving the underlying material parameter combinations.

To populate this continuous, four-dimensional hypercube for numerical analysis, the design parameter sets are constructed utilizing a quasi-random Sobol sequence framework. Unlike standard pseudo-random Monte Carlo generators, which are prone to statistical clustering and sparse zones, the Sobol sequence guarantees an optimally uniform space-filling distribution (Sobol', 1967). As detailed by Sobol' and Bratley, a sequence of s-dimensional quasi-random vectors $q^1, q^2, \ldots, q^n$ achieves these enhanced, low-discrepancy uniformity properties whenever the total sample size satisfies:

$$
n = 2^k \quad \text{for } k \geq \max(2s, \tau_s + s - 1)
$$

$$
\begin{array}{c|cccccccc}
s & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\
\tau_s & 0 & 0 & 1 & 3 & 5 & 8 & 11 & 15
\end{array} \tag{7}
$$

where τ_s represents a deterministic mathematical parameter governing the initialization blocks for a given dimension s. For the specific four-dimensional space ($s = 4$) modeled in this digital twin, the corresponding sequence initialization value is mathematically established as $\tau_4 = 3$ (Bratley & Fox, 1988).

To satisfy the baseline Property-A uniformity conditions required to accurately map highly non-linear structural transitions, a minimum sampling volume threshold of $n = 256$ (2^8) automated, transient-dynamic simulations was implemented. This systematic configuration ensures that the resulting discrete data points are spaced uniformly across the entire operational envelope.

From a PHM perspective, prioritizing this low-discrepancy uniformity is vital for downstream ML training. The transition between safe elastic bounce-back, localized plastic buckling, and catastrophic material fracturing in high-energy collisions is often separated by an incredibly thin boundary in the parameter space (Chen, Qiu, Gong, Shen, & Zhang, 2026). A standard random distribution might completely miss this critical transition due to clustering. By utilizing the Sobol sequence, the data engine densely and uniformly maps these boundary thresholds, providing the diagnostic neural network with an idealized, spatially uniform training environment that inherently resists local generalization errors.

The mathematical uniformity of these 256 generated collision

Table 2. Sampling Parameters and Operational Boundary Envelopes

Parameter	**Symbol**	**Unit**	**Lower**	**Upper**
Lateral Position	(y)	mm	0	-700
Vertical Position	(z)	mm	0	210
Impact Velocity	(v)	m/s	> 0	90
SSF	(ξ)	—	-1	1

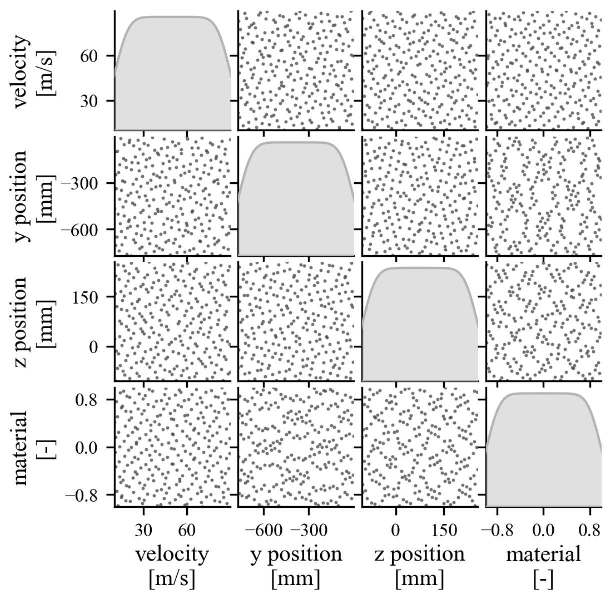

Figure 4. Pairwise Scatter Plots of Sobol Sequence of Simulation Parameters

scenarios is visualized in Figure 4, which presents pairwise scatter plots.

To illustrate the variance captured within the generated dataset, Figure 5 presents two representative examples of distinct structural outcomes. Because the Sobol sequence modifies all spatial, kinematic, and material variables simultaneously, these two instances are not intended for isolated parametric comparison, rather they demonstrate the variation of the simulation.

The left image displays minor deformation. This outcome is the result of a combined moderate impact velocity ($v = 51.9m/s$) and a statistically tougher material state ($xi \approx +0.42$), allowing the structure to safely absorb the kinetic energy. Conversely, the right image exhibits severe structural damage. This damage is driven by a much higher, near-maximum impact velocity ($v = 89.5m/s$) coupled with a statistically weaker material profile ($xi \approx -0.44$). Ultimately, isolating the exact influence of each independent variable is impossible from isolated pairs. It requires the mathematical interpolation of the entire high-dimensional dataset, which the downstream diagnostic neural network is designed to perform.

4. THE DIGITAL THREAD - VIRTUAL TELEMETRY AND OPERATIONAL LABELING

Bridging the gap between the continuous physics simulated by the Digital Twin and the discrete classification outputs required by the Industry ML necessitates a "Digital Thread." This thread begins with the extraction of virtual telemetry and concludes with the systematic assignment of human-expert operational labels.

4.1. Virtual Sensor Networks and Hardware Feasibility

Optimizing the sensor layout is a critical yet highly constrained challenge in the design of SHM systems for high-speed rail. For an AI-driven predictive maintenance framework, the selection and placement of sensors directly dictate both the dimensional complexity of the neural network's input layer and the physical feasibility of real-world deployment. This study utilizes two distinct modalities of virtual telemetry: strain gauges and accelerometers (visualized in Figure 6).

While a high-density sensor network theoretically yields a more granular representation of structural behavior, it introduces significant technical trade-offs that hinder both algorithmic convergence and practical implementation:

- **Data Scarcity and Dimensionality Challenge:** In machine learning, every additional sensor channel introduces a new high-frequency time-series vector into the network's input layer. As the input dimensionality increases, the network's parameter space expands rapidly, exponentially increasing the volume of training data required to achieve convergence and prevent overfitting. Given that high-fidelity non-linear finite element (FE) simulations are computationally intensive, rigorously minimizing the number of input channels is essential to maintain a mathematically manageable training dataset.
- **Economic and Practical Retrofitting Feasibility:** To transition from a virtual framework to real-world rolling stock applications, the sensor layout must account for the strict physical constraints of active train fleets. Retrofitting existing underrun guards with extensive wiring harnesses, complex cable routing, and massive data-acquisition nodes is economically prohibitive and frequently mechanically unfeasible due to space limitations. A minimalist sensor footprint is therefore mandatory for viable cross-fleet deployment.

The underrun guard features a highly complex, geometrically

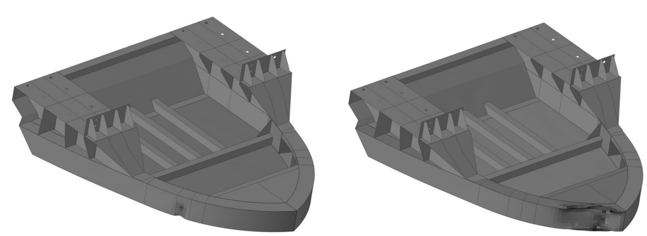

(a) Minor Deformation (b) Severe Failure

Figure 5. Visual Comparison of Simulated Damage States for Different Initial Parameters

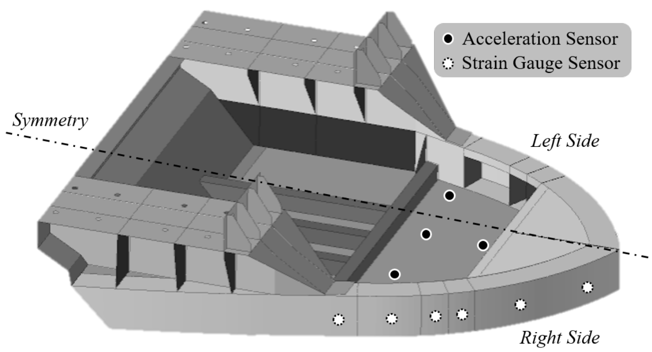

Figure 6. Acceleration and Strain Gauge Sensor Positioning

segmented layout divided into a series of internal chambers and box-section compartments. This chambered architecture effectively absorbs impact energy but significantly alters how mechanical loads propagate through the assembly. To capture these localized dynamics without overcrowding the ML feature space, the monitoring framework decouples localized structural deformation from global rigid-body kinematics through a optimized, hybrid placement strategy.

To monitor localized plastic yielding, a network of 12 virtual strain gauges is distributed strategically across the segmented internal chambers. Rather than restricting instrumentation to a single half of the structure, the network utilizes a symmetric configuration, with 6 sensors allocated to the left side and 6 to the right.

This specific configuration explicitly exploits the structural symmetry along the longitudinal centerline ($y = 0$) to optimize dataset diversity. During the automated data generation, collision impacts are restricted exclusively to a single lateral half of the component (as defined in Section 3.1). However, because a high-energy localized impact triggers global structural twisting, bending, and energy-absorption propagation, the mechanical response of the entire underrun guard is highly asymmetric. For instance, an off-center collision causes the structural frame to deflect laterally, plunging the impacted side into compression and localized buckling, while simultaneously placing the opposite side under tensile stretching. Placing sensors on both sides ensures that these decoupled, non-linear global bending behaviors and cross-compartment load transfers are seamlessly captured. Each individual internal chamber is instrumented with either one or two sensors based on its geometric volume, ensuring that topological damage indicators are fully mapped into the dataset.

To capture the global intensity, impulse duration, and rigid-body deceleration of the collision event, a network of 4 accelerometers is deployed. These sensors are positioned further inside the rigid structural backing of the component, safely isolating them from the severe localized material tearing occurring at the impact front. The sensors are arranged in a multi-axial cross-configuration on the floor of the underrun guard:

- Two sensors are aligned along the longitudinal centerline (close to the front interface and the rear chassis) to capture axial deceleration and pitching moments.
- Two lateral sensors are mirrored on the left and right extremities of the internal floor to measure asymmetric torsional twisting and lateral shock wave propagation.

Ultimately, this distributed, symmetric strain and acceleration layout ensures maximum information density with a minimal hardware footprint. It provides a streamlined, informative input layer optimized for modern diagnostic neural networks.

5. DATA PROCESSING AND EXPERT LABELING PIPELINE

To transform raw finite element simulation outputs into a supervised machine learning dataset, the transient time-series telemetry and structural deformation states must be processed, consolidated, and mapped to qualitative operational decisions. This section details the signal processing pipeline, the visual data export, and the expert-driven labeling workflow used to bridge the gap between quantitative engineering metrics and field-maintenance logic.

5.1. Sensor Signal Consolidation and Reduction

To transform raw Finite Element (FE) simulation outputs into a supervised machine learning dataset, the transient time-series telemetry and structural deformation states must be strictly processed and consolidated. The raw data generated by the high-fidelity simulation contains multi-axial components that introduce high-frequency noise and highly redundant feature dimensions. If fed directly into a neural network, this bloated feature space would severely degrade algorithmic convergence. To optimize the neural network's input layer, the multi-component sensor signals are algorithmically reduced to omnidirectional scalar metrics at each time step.

Acceleration Vector Fusion Each of the 4 deployed accelerometers captures raw acceleration data across three orthogonal spatial axes: $a_x(t)$, $a_y(t)$, and $a_z(t)$. To isolate the absolute magnitude of the collision impulse and eliminate directional dependencies that could lead to ML overfitting, the triaxial signals are mathematically fused into a singular, time-dependent overall acceleration magnitude ($a_{\text{total}}(t)$) using the Euclidean norm:

$$a_{\text{total}}(t) = \sqrt{a_x(t)^2 + a_y(t)^2 + a_z(t)^2} \tag{8}$$

Strain Gauge Tensor Reduction To capture the complex multi-axial deformation states without expanding the input feature space, a rigorous dimensional reduction strategy is applied to the virtual strain gauge data. For each localized sensor position, the finite element solver outputs eight distinct strain components and shell element integration metrics. These consist of membrane strains ($\varepsilon_1, \varepsilon_2, \varepsilon_{12}$), transverse shear strains ($\mathrm{SH}_1, \mathrm{SH}_2$), and structural curvatures ($\mathrm{K}_1, \mathrm{K}_2, \mathrm{K}_{12}$).

Feeding all eight highly correlated channels per sensor into the neural network would expand the strain feature space to 96 continuous time-series channels. Instead, these values are synthesized into a single, comprehensive scalar indicator per time step: the Von Mises equivalent strain ($\varepsilon_{\text{total}}(t)$). This specific scalar invariant provides the best mechanical representation of the overall strain state and distortion energy under complex, large-strain conditions. Furthermore, a significant practical advantage of the Von Mises formulation is its direct compatibility with real-world, field-deployable Industry ML hardware. It is easily calculable using standard physical tri-axial strain gauge rosettes (Shrivastava, Ghosh, & Jonas, 2012).

Because the localized deformation intensity is accurately captured by this equivalent strain metric, the remaining shear strain and curvature components can be safely discarded. This reduction condenses the strain input space from 96 channels down to just 12 integrated inputs, significantly decreasing computational overhead during training and mitigating the model's vulnerability to the disadvantages of high dimensionality (Peng et al., 2024).

The scalar reduction relies on plane-stress shell mechanics (Shrivastava et al., 2012). First, the unmeasured through-thickness principal strain component (ε_3) is derived dynamically by enforcing the elastic-plastic transverse boundary conditions as a function of the material's Poisson's ratio (ν):

$$\varepsilon_3 = -\frac{\nu}{1.0-\nu}(\varepsilon_1 + \varepsilon_2) \tag{9}$$

Once the triaxial strain configuration is established, the final omnidirectional scalar representation of the localized deformation intensity is computed via the continuous Von Mises equivalent strain criterion (Shrivastava et al., 2012):

$$\varepsilon_{\text{total}}(t) = \frac{1}{1.0+\nu}\sqrt{\begin{matrix} 0.5\Big[(\varepsilon_1-\varepsilon_2)^2 + (\varepsilon_2-\varepsilon_3)^2 + \\ (\varepsilon_3-\varepsilon_1)^2 + 1.5\varepsilon_{12}^2\Big]\end{matrix}} \tag{10}$$

This mathematical compression retains the mechanical invariants of the localized stress-strain fields, ensuring that the critical energy-absorption transitions and plastic distortions within the structural compartments are fully preserved (Shrivastava et al., 2012). Consequently, the downstream diagnostic network is supplied with a highly descriptive, mathematically streamlined input vector that reflects the true underlying continuum mechanics of the collision event.

5.2. Translating Structural Deformation into Operational Logic

In the context of SHM and PHM, determining the operational safety of a vehicle following a high-energy collision is incredibly complex. Purely numerical analysis of severe, chaotic plastic deformation is often insufficient to automatically dictate definitive operational and maintenance actions. Currently, in real-world railway operations, post-collision damage assessment relies mainly on manual inspections, resulting in a qualitative evaluation by a domain expert.

To construct a robust, AI-driven predictive operations and maintenance framework, it is imperative to fuse this qualitative expert knowledge with the quantitative, high-frequency sensor time-series data generated by the finite element simulations. We achieve this synthesis through a structured "expert-in-the-loop" methodology. By translating the complex numerical FE data into intuitive visual formats, we can conduct expert evaluation sessions where domain specialists rate the simulated damage images exactly as they would physical hardware.

For every simulated collision, the deformed state of the underrun guard is rendered and exported across a standardized image array. To ensure the visual information is robust and independent of visualization artifacts, the pipeline generates images from four distinct viewing angles as shown in Figure 7:

- Front View (y-z plane)
- Top View (x-y plane)
- Side View (x-z plane)
- Isometric View (3D perspective)

Additionally, each perspective is exported in two distinct color modes: a realistic metallic texture rendering (reproducing exactly what a field technician would observe during a physical inspection) and a localized plastic strain contour plot. By identifying peak deformation zones that remain undetected during visual inspection, this approach transcends standard expert methodologies and improves the overall precision of the expert analysis.

Utilizing this standardized visual array, human domain experts, specifically railway safety engineers and maintenance supervisors, review the 4-angle, deformation images. Assisted by the supporting numerical severity indices extracted from the mesh, the experts classify each post-collision state

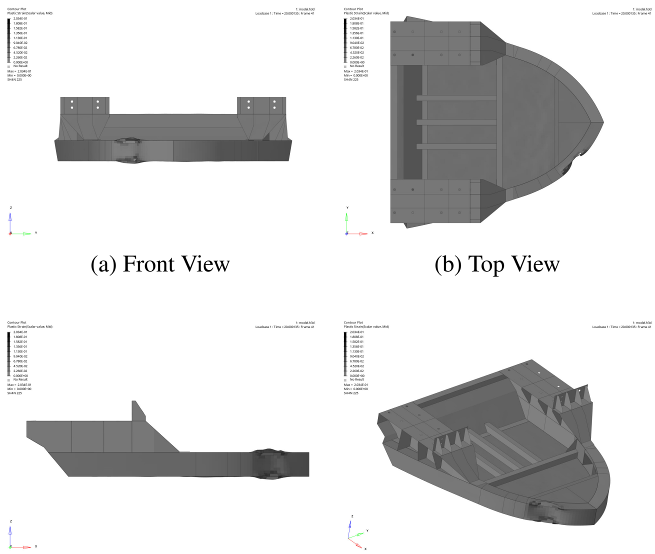

(a) Front View

(b) Top View

(c) Side View

(d) Isometric View

Figure 7. Standardized Multi-Angle Damage Visualization for Expert Labeling

into one of three predefined, actionable categories that dictate both immediate operations and subsequent maintenance protocols:

- **Continued Service and Part Reuse:** Minor elastic or negligible localized plastic deformation; the train can safely complete its scheduled route, and the underrun guard can be authorized for continued use without structural intervention.
- **Reduced Speed and Part Repair:** Moderate structural damage or partial chamber buckling detected. The train may proceed safely to the nearest maintenance depot under velocity restrictions, where the component will require targeted repair procedures.
- **Immediate Stop and Part Replacement:** Severe structural failure, cracking, or catastrophic energy-absorption exhaustion. The vehicle must be halted immediately to prevent derailment risks, and the compromised structure mandates a complete replacement.

Through this labeling process, the qualitative knowledge of the expert is permanently captured and fused with the corresponding virtual sensor time-series data. This directly results in a highly specialized, supervised learning dataset that intrinsically links raw crash telemetry with authoritative structural health assessments.

6. DIGITAL TWIN VALIDATION AND DATASET APPLICATION

Before the synthetic dataset can be reliably deployed to train a diagnostic neural network, the foundational Physics-of-Failure Digital Twin must be validated against real-world physical mechanics. This section details the experimental validation of the simulation framework and the subsequent performance of the resulting machine learning dataset.

6.1. Experimental Validation of the Data Engine

To ensure that the automated finite element (FE) pipeline serves as a reliable surrogate for real-world impact physics, destructive laboratory experiments and quasi-static clamp tests were conducted. Physical specimens of the underrun guard were subjected to impact testing to establish an empirical ground truth. The transient sensor telemetry, peak-to-peak acceleration values, and localized strain gauge responses generated by the virtual digital twin were cross-referenced against these physical tests to quantify the ”sim-to-reality” gap.

As illustrated in Figure 8, the simulation exhibits a high degree of qualitative and quantitative correlation with the physical laboratory experiments. The digital twin reproduces the complex macroscopic folding, localized buckling, and plastic energy dissipation mechanisms observed in the destructive physical tests. For the acceleration telemetry, filtering the raw signals via a CFC1000 filter brings the simulated and experimental peaks closely together: The filtered experimental peak reaches 337.2 g, while the filtered simulation peak reaches 343.4 g, resulting in a deviation of only 1.8 %. Furthermore, a quantitative comparison of the strain gauge (DMS) signals indicates that the simulation predicts the correct deformation states (capturing compression where compression occurred and tension where tension was measured), with quantitative deviations typically ranging between 10.9% and 44.2 %. This validation concludes, that the automated FE data engine produces field-equivalent collision telemetry.

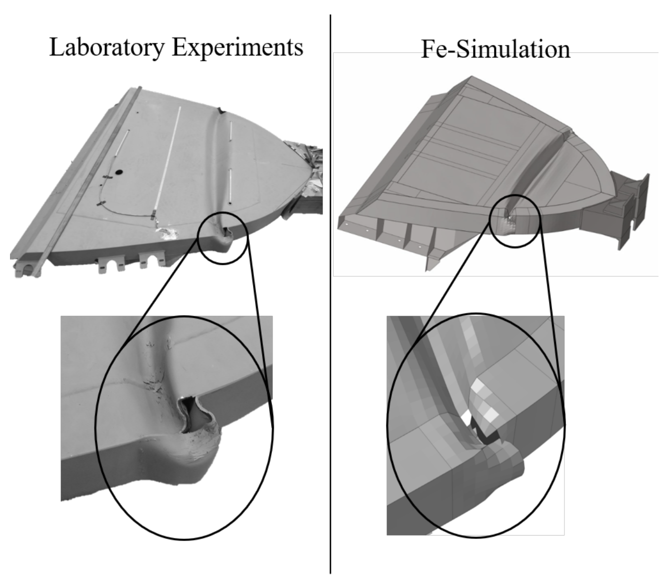

Figure 8. Qualitative Validation: Structural Deformation in the FE Simulation vs. Destructive Laboratory Experiments

6.2. Machine Learning Application and Diagnostic Performance

Following the physical validation of the digital twin, the pipeline was executed to generate the final supervised learning dataset consisting of 256 distinct, fully resolved collision events. Which were split into an 80/20 training and evaluation set to perform cross Validation. Each event fundamentally links the reduced time-series sensor telemetry (12 strain metrics and 4 acceleration profiles) with the specific, qualitative operational and maintenance labels provided by the human domain experts. The model is optimized using standard multi-task loss weighting across its classification and regression heads, while regularization is implemented via spatial batch normalization, ReLU activation blocks, and selective dropout rates of 0.1 within the fusion and decision dense layers.

To evaluate the efficacy of this data-generation pipeline, the dataset was utilized to train a multi-modal Siamese neural network architecture integrating CNN feature extraction, a centralized fusion layer, and a multi-modal-objective output head, as shown in Figure 9. To combat data scarcity and prevent the model from memorizing noise rather than learning underlying physical laws, the architecture incorporates structural inductive biases (via temporal invariance and sensor symmetry) and auxiliary physical regression tasks predicting continuous deformation and plastic strain metrics. While these auxiliary regression heads are essential during training for backpropagating physical constraints and regularizing the encoder, their outputs are discarded at inference, leaving the primary classification heads to operate on standard telemetry.

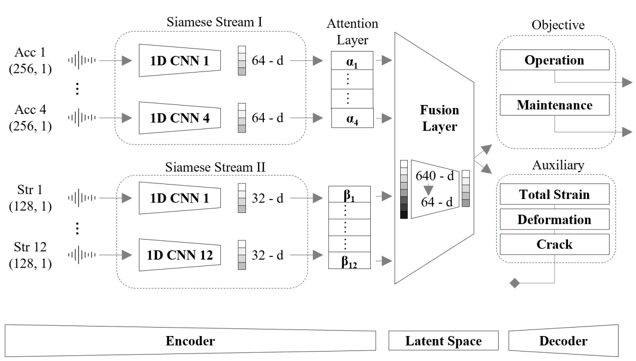

Figure 9. ML-Architecture: Multi-Modal Multi-Dimensional Siamese Network with Multi-Classification-Objective and Multi-Regression-Auxiliary Heads

Despite the severely constrained size of the training dataset ($n = 205$), the resulting diagnostic machine learning model achieves good classification performance on the unseen validation set when incorporating auxiliary physical regression. As detailed in the error analysis via confusion matrices in Figure 10, the integration of physical regression significantly re-

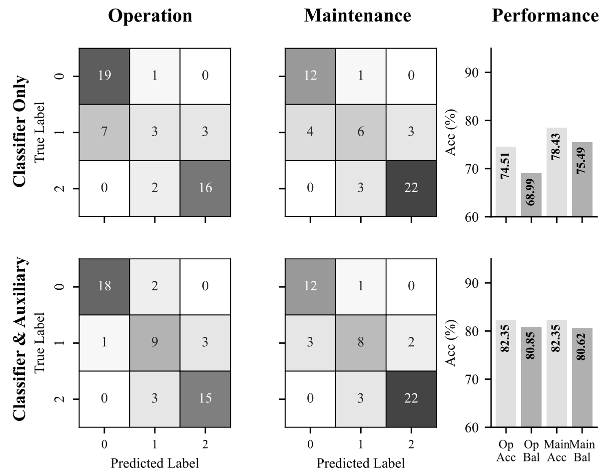

Figure 10. Confusion matrix and accuracy comparison: The full architecture yields higher accuracy across the different classes

duces high-risk misclassifications by building robust decision boundaries that respect structural safety limits.

Comparing the baseline model (Classifier Only) with the full architecture (Classifier and Auxiliary) across both the operation and maintenance domains reveals clear performance gains. For the operation task, both overall and balanced accuracies improve substantially. Specifically, the confusion matrices demonstrate that severe cross-class errors—such as true intermediate operation states being incorrectly categorized as safe continued service—drop dramatically from 7 instances down to just 1 instance. Similarly, in the maintenance domain, overall and balanced accuracies increase, and intermediate maintenance states exhibit improved classification fidelity with fewer misclassifications bleeding into adjacent severity bins. These results illustrate that the auxiliary regression tasks successfully guide the latent feature space toward physically consistent representations, eliminating dangerous misclassification events and ensuring more reliable decision-making across all operational categories.

Furthermore, to assess sample efficiency, a data ablation study was conducted by systematically reducing the training set size for the full architecture until its performance matched the dual-classifier baseline. The baseline represents the accuracy of a model utilizing only the two classifiers. For the full architecture with both classifiers and the three regressions, the number of training samples was reduced stepwise by 10 % until the accuracy reached the previously set baseline accuracy. The results show that at 40 % of the available training data, the full architecture drops to the same level as the classifier-only architecture. By repeating this training sample reduction with the classifier-only architecture, it is visible that the accuracy drops at a significantly higher rate. As pre-

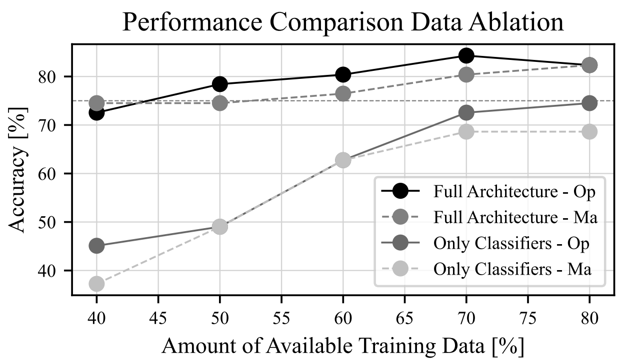

Figure 11. Data ablation results: The full architecture has a lower accuracy drop for less training data then the only classifiers architecture

sented in Figure 11, the ”Only Classifier” algorithm exhibits a much higher sensitivity to training data volume compared to the proposed full architecture, demonstrating that standard supervised methods fail to generalize at this critical threshold while the auxiliary regression targets maintain high performance.

7. CONCLUSION

The integration of Machine Learning (ML) into Prognostics and Health Management (PHM) frameworks is fundamentally bottlenecked by data scarcity in extreme, high-consequence environments. While employing Finite Element (FE) simulations to synthetically generate standard sensor telemetry is an established practice, this conventional approach severely underutilizes the true analytical power of the solver. The core contribution of this research is a paradigm shift in data utilization: redefining the FE simulation not merely as a virtual replica of sparse field sensors, but as a rich, physical teacher. By explicitly extracting ”hidden” continuum mechanics, data that is inherently calculated by the simulation but not available to measure in reality, this study establishes an asymmetric ”Digital Thread” that improves training efficiency and algorithm robustness.

Crucially, this framework pioneers an asymmetric training strategy to combat label scarcity. Rather than discarding unmeasurable physical data, this methodology actively extracts continuous variables, such as localized plastic strain distributions, to serve as auxiliary regression targets. By forcing a multi-modal Siamese network to learn these ”hidden” physical laws alongside the primary qualitative expert labels, the model receives a deep, physically grounded regularization. Because these auxiliary targets are utilized exclusively during the training phase and entirely discarded for real-world inference, the architecture remains lightweight, highly deployable, and inherently resists the overfitting typically associated with sparse-label regimes.

To make this advanced data generation computationally feasible and mitigate the curse of dimensionality, this methodology simultaneously introduced rigorous physical and algorithmic boundaries. By mathematically simplifying collision kinematics, condensing multi-dimensional material properties into a singular Strength Scaling Factor (SSF), and employing a low-discrepancy Sobol sequence, the automated data engine densely mapped the highly sensitive boundary transitions of structural failure. Furthermore, the algorithmic reduction of raw multi-axial FE tensors into comprehensive scalar invariants (Von Mises equivalent strain and Euclidean acceleration) successfully compressed the virtual telemetry, preventing input-layer bloat.

The efficacy of this combined approach was confirmed through experimental validation and subsequent ML evaluation. Despite a severely constrained dataset size of only $n = 256$, the full architecture achieved an outstanding classification accuracy of over 82 %, significantly outperforming traditional baseline models. The ablation studies conclusively proved that leveraging auxiliary ”hidden” simulation data dramatically reduces the algorithm’s sensitivity to training volume, maintaining robust predictive performance even when data is heavily restricted. This confirms the core thesis of the study: strategically exploiting unmeasurable FE data exclusively for training, combined with mathematically optimized parameter reduction, fundamentally bridges the data scarcity gap.

While this framework successfully demonstrates the viability of automated data engines, certain methodological constraints must be acknowledged to direct future investigations. First, the algorithmic reduction of multi-dimensional material properties into the singular SSF variable currently operates as an asserted engineering heuristic. Although this scaling is strictly bounded by statistical manufacturing tolerances and governed by the physically validated strength-ductility tradeoff, it requires independent empirical validation through targeted material testing. Furthermore, while existing literature confirms the superiority of low-discrepancy Sobol sequences over standard Monte Carlo sampling in minimizing clustering within high-dimensional spaces, a direct empirical comparison of their respective impacts on downstream classification accuracy was precluded by the computational scope of this study. Establishing this full-scale Monte Carlo baseline remains a priority for subsequent evaluations.

Beyond addressing these sampling and heuristic validations, the next phases of this research will focus on advancing both the physical and algorithmic components of the diagnostic framework. This includes introducing multi-object collision scenarios and evaluating the deployed model’s robustness against real-world sensor degradation by injecting synthetic white noise and simulated channel blackouts into the telemetry array. These advancements will ensure that the resulting PHM network maintains high classification integrity, ultimately delivering a highly reliable, autonomous diagnostic asset for modern high-speed operations.

NOMENCLATURE

Acronyms

CNN	Convolutional Neural Network
DL	Deep Learning
FE	Finite Element
ICE	Intercity-Express
ML	Machine Learning
NN	Neural Network
PHM	Prognostics and Health Management
SHM	Structural Health Monitoring
SSF	Strength Scaling Factor

Roman Symbols

a_x, a_y, a_z	Orthogonal acceleration components
a_{total}	Overall Euclidean acceleration magnitude
E	Young's Modulus
$\mathrm{K}_1, \mathrm{K}_2, \mathrm{K}_{12}$	Structural curvatures
m_1, m_2	Masses of the colliding bodies
n	Total sample size (number of simulations)
s	Dimensionality of the parameter space
$\mathrm{SH}_1, \mathrm{SH}_2$	Transverse shear strains
v_{train}	Operational velocity of the high-speed train
v_{object}	Velocity vector of the collision object
v_{rel}	Relative impact velocity vector
x, y, z	Spatial coordinate axes (longitudinal, lateral, vertical)

Greek Symbols

$\varepsilon_1, \varepsilon_2$	In-plane membrane strains
ε_3	Through-thickness principal strain
ε_{12}	In-plane shear strain
$\varepsilon_{p,\text{max}}$	Maximum allowable failure plastic strain
$\varepsilon_{\text{total}}$	Von Mises equivalent strain
ν	Poisson's ratio
ξ	Dimensionless material strength scaling factor
ρ	Material density
σ_y	Initial yield strength
σ_{uts}	Ultimate tensile strength
τ_s	Sobol sequence initialization parameter

REFERENCES

24RHEIN. (2024). *ICE-Zug (Image regarding DB railway blockage Oberhausen-Emmerich).* https://www.24rhein.de/assets/images/34/539/34539698-ice-zug-2sVrJNjiNM73.jpg. (Accessed: 2026-05-28. Image from article: https://www.24rhein.de/leben-im-westen/verkehr/ausfaelle-busse-ersatz-oepnv-bahnsperrung-deutsche-bahn-oberhausen-emmerich-nrw-ice-re-93062986.html)

Agnew, S. R., & Duygulu, O. (n.d.). Statistical evaluation of mechanical property variability in structural alloys. *Materials Science and Engineering: A*, *407*(1-2), 29–37.

Akrim, A., Gogu, C., Vingerhoeds, R., & Salaün, M. (2023). Self-supervised learning for data scarcity in a fatigue damage prognostic problem. *Engineering Applications of Artificial Intelligence*, *120*, 105837. Retrieved from https://www.sciencedirect.com/science/article/pii/S0952197623000210 doi: https://doi.org/10.1016/j.engappai.2023.105837

Altair Engineering Inc. (2023). Altair radioss theory manual [Computer software manual]. Troy, Michigan, USA. (Material Laws: Elasto-plasticity and Failure Criteria)

Avevor, J., Adeniyi, M., Enyejo, L., & Aikins, S. (2024, 08). Machine learning-driven predictive modeling for frp strengthened structural elements: A review of ai-based damage detection, fatigue prediction, and structural health monitoring. *International Journal of Scientific Research and Modern Technology*, *3*, 1-20. doi: 10.38124/ijsrmt.v3i8.420

Bratley, P., & Fox, B. L. (1988, March). Algorithm 659: Implementing sobol's quasirandom sequence generator. *ACM Trans. Math. Softw.*, *14*(1), 88–100. Retrieved from https://doi.org/10.1145/42288.214372 doi: 10.1145/42288.214372

Chen, X., Qiu, G., Gong, J., Shen, S., & Zhang, Y. (2026). A physics-informed design generator for long-span reticulated domes: Replacing iterative finite element analysis for optimal solutions. *Buildings*, *16*(9). Retrieved from https://www.mdpi.com/2075-5309/16/9/1663 doi: 10.3390/buildings16091663

Chopin, N. (2004). Central limit theorem for sequential monte carlo methods and its application to bayesian inference. *The Annals of Statistics*, *32*(6), 2385–2411. doi: 10.1214/009053604000000698

DIN EN 15227: Railway applications - Crashworthiness requirements for railway vehicle bodies (Bahnanwendungen - Anforderungen an die Kollisionssicherheit von Schienenfahrzeugkästen) (Standard No. DIN EN 15227:2020-06). (2020, 6). Berlin, Germany: Deutsches Institut für Normung (DIN).

EN 10025-2: Hot rolled products of structural steels - Part 2: Technical delivery conditions for non-alloy structural steels (Standard No. EN 10025-2:2019). (2019). Brussels, Belgium: European Committee for Standardization (CEN). (Details baseline properties for S235 steel including 235 MPa yield strength, 360 MPa ultimate tensile strength (UTS), and 26% elongation)

European Committee for Standardization (CEN). (2019). *En*

10025-2: Hot rolled products of structural steels - part 2: Technical delivery conditions for non-alloy structural steels (Tech. Rep.). Brussels: CEN.

Fang, C., Otero, H. J., Greenberg, D., & Neumann, P. J. (2011). Cost-utility analyses of diagnostic laboratory tests: A systematic review. *Value in Health*, *14*(8), 1010-1018. Retrieved from `https://www.sciencedirect.com/science/article/pii/S109830151101494X` doi: https://doi.org/10.1016/j.jval.2011.05.044

Goldsmith, W. (2001). *Impact: The theory and physical behaviour of colliding solids*. Dover Publications.

Hess, P. E., Bruchman, D., Assakkaf, I. A., & Ayyub, B. M. (2002). Uncertainties in material and geometric strength and load variables. *Naval Engineers Journal*, *114*(2), 139–166. doi: 10.1111/j.1559-3584.2002.tb00128.x

Johnson, G. R., & Cook, W. H. (1983). A constitutive model and data for metals subjected to large strains, high strain rates and high temperatures. *Proceedings of the 7th International Symposium on Ballistics*, 541–547.

Joint Committee on Structural Safety (JCSS). (2001). *Probabilistic model code: Part 3 - resistance models (structural steel)* (Tech. Rep.). JCSS. (Provides the foundational probabilistic distributions and expected coefficients of variation (COV) for structural steel yield and ultimate strengths.)

Kashiparekh, K., Narwariya, J., Malhotra, P., Vig, L., & Shroff, G. (2019). Convtimenet: A pre-trained deep convolutional neural network for time series classification. *2019 International Joint Conference on Neural Networks (IJCNN)*. doi: 10.1109/ijcnn.2019.8852105

Mareau, C. (2020). A thermodynamically consistent formulation of the johnson–cook model. *Mechanics of Materials*, *143*, 103340. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0167663619307720` doi: https://doi.org/10.1016/j.mechmat.2020.103340

Morris, J. W. (2009). The strength and ductility of structural alloys. *Progress in Materials Science*, *54*(3), 312–345.

Ochella, S., Shafiee, M., & Dinmohammadi, F. (2022). Artificial intelligence in prognostics and health management of engineering systems. *Engineering Applications of Artificial Intelligence*, *108*, 104552. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0952197621003961` doi: https://doi.org/10.1016/j.engappai.2021.104552

Peng, D., Gui, Z., & Wu, H. (2024). *Interpreting the curse of dimensionality from distance concentration and manifold effect.* Retrieved from `https://arxiv.org/abs/2401.00422` doi: 10.48550/arXiv.2401.00422

Pore, T., Thorat, S. G., & Nema, A. A. (2021). Review of contact modelling in nonlinear finite element analysis. *Materials Today: Proceedings*, *47*, 2436-2440. Retrieved from `https://www.sciencedirect.com/science/article/pii/S2214785321034295` (International Conference on Materials and System Engineering) doi: https://doi.org/10.1016/j.matpr.2021.04.504

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., & Tarantola, S. (2010). Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer Physics Communications*, *181*(2), 259-270. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0010465509003087` doi: https://doi.org/10.1016/j.cpc.2009.09.018

Shrivastava, S., Ghosh, C., & Jonas, J. (2012, 03). A comparison of the von mises and hencky equivalent strains for use in simple shear experiments. *Philosophical Magazine*, *92*, 779-786. doi: 10.1080/14786435.2011.634848

Sobol', I. M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *Zh. Vychisl. Mat. Mat. Fiz.*, *7*(4), 784–802.

Terfloth, S. (2008). *ICE-TD Neustadt(Dosse).* `https://de.wikipedia.org/wiki/Datei:ICE-TD_Neustadt(Dosse).jpg`. (Accessed: 2026-05-28)

Wu, W., Cantero-Chinchilla, S., Prescott, D., Remenyte-Prescott, R., & Chiachío, M. (2024). A general approach to assessing shm reliability considering sensor failures based on information theory. *Reliability Engineering System Safety*, *250*, 110267. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0951832024003399` doi: https://doi.org/10.1016/j.ress.2024.110267

Zhang, Y., Hao, R., Niiranen, J., Yang, Y., Brühwiler, E., Su, D., & Nagayama, T. (2024). Computational engine for finite element digital twins of structural dynamics via motion data. *Engineering Structures*, *316*, 118630. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0141029624011921` doi: https://doi.org/10.1016/j.engstruct.2024.118630

BIOGRAPHIES

Maximilian Posner received his Bachelor of Science and Master of Science degrees in Mechatronics from the University of Stuttgart in Stuttgart, Germany. He is currently working as a scientific researcher and pursuing his Ph.D. in Mechanical Engineering at the University of Stuttgart. Alongside his doctoral studies, he serves as the

”Bereichsorganisator” (Division Coordinator) for the Department of Schienenfahrzeugtechnik (Railway Vehicle Technology). His current research focuses on the intersection of autonomous vehicles and machine learning, specifically applying advanced artificial intelligence algorithms to complex engineering and transportation systems.

Martin Dazer received his Bachelor of Engineering in Mechanical Engineering from the Baden-Wuerttemberg Cooperative State University (DHBW). He subsequently earned his Master of Science in Mechanical Engineering, his Dr.-Ing. in Reliability Engineering, and completed his Habilitation in Reliability Engineering, all at the University of Stuttgart in Stuttgart, Germany. He serves as a Head of Department at the Institute for Machine Components at the University of Stuttgart. His research and professional expertise center on reliability engineering, advanced machine components and Prognostics and Health Management.