Predictive Modeling of Ductile-to-Brittle Transition in Hierarchical RC Frames: A Comparative Study of Deep Learning and Machine Learning Surrogate Models

Alessandro Calvi[1]

[1] alessandro.calvi84@gmail.com , ENEL Green Power S.p.A, Torino, Italy

Abstract

This research addresses the critical challenge of progressive collapse in Reinforced Concrete (RC) frames by establishing a rigorous multi-scale link between material fracture mechanics and macro-scale structural hierarchy. The study investigates the highly non-linear "Ductile-to-Brittle" transition, identifying two dimensionless ratios as the fundamental governing parameters: the Geometric Ratio (a/H), which marks the onset of brittle behavior at a threshold of approximately 0.5, and the Energetic Ratio (L_{crit}/L), which triggers the loss of energy dissipation capacity at approximately 0.75.
To overcome the computational intensity of traditional Finite Element Method (FEM) simulations, this work implements and compares various Artificial Intelligence (AI) and Machine Learning (ML) surrogate models. Specifically, a Deep Learning (DL) architecture with multiple hidden layers and ReLU activation was developed to map the structural fragility landscape. The DL model successfully captured the physics of the transition, with the loss landscape confirming convergence to a low, well-defined minimum. The resulting 3D prediction surface effectively identifies the "non-linear basin" where maximum gradient shifts occur exactly at the predicted physical thresholds, validating the AI's ability to learn complex fracture mechanics without explicit programming.
Furthermore, the study evaluates the performance of other ML paradigms, such as Random Forest (RF) for discrete decision-tree logic and Support Vector Machines (SVM) for defining clear-cut safety boundaries in forensic engineering. The comparison highlights that while RF is suitable for rapid inspection protocols, DL provides a superior, continuous representation of the structural tipping point.
The results demonstrate that AI-driven surrogate modeling can predict structural robustness and identify catastrophic vulnerabilities far more efficiently than conventional methods. This approach provides a scalable framework for the next generation of performance-based design and structural health monitoring, ensuring safer and more resilient infrastructure against accidental actions and progressive collapse.

Keywords: Progressive Collapse, Fracture Mechanics, Deep Learning, Neural Networks, Random Forest, RC Frames, Ductile-to-Brittle Transition, Structural Robustness

1. Introduction: Bridging Fracture Mechanics and Structural Collapse

1.1 The Challenge of Progressive Collapse in Modern Infrastructure

The phenomenon of progressive collapse represents one of the most critical and least understood failure modes in structural engineering. Unlike conventional failure mechanisms that are typically localized and predictable, progressive collapse is characterized by a disproportionate spread of damage following an initial localized failure. A single column removal, a localized explosion, or even a vehicle impact can trigger a cascade of failures that propagates throughout the structure, ultimately leading to partial or total collapse. This type of failure has been tragically demonstrated in several high-profile incidents throughout modern history, from the partial collapse of Ronan Point in 1968 to the catastrophic failure of the Alfred P. Murrah Federal Building in Oklahoma City in 1995, and most notably, the collapse of the World Trade Center towers in 2001.

In civil engineering practice, structural resistance has traditionally been assessed through two fundamental design criteria that have formed the backbone of structural codes and standards for decades. The first is the Serviceability Limit State (SLS), which ensures that structures perform correctly under normal operational conditions without experiencing excessive deformation, vibration, or cracking that would compromise their intended function or cause discomfort to occupants. The second is the Ultimate Limit State (ULS), which defines the maximum load-carrying capacity a structure can sustain before experiencing failure. These criteria have served the engineering community well for routine design scenarios, providing clear benchmarks for structural adequacy and safety margins.

However, the increasing frequency and severity of extreme events—whether natural disasters such as earthquakes and hurricanes, or man-made threats including terrorist attacks and accidental explosions—have repeatedly exposed the limitations of these conventional approaches. Real structural failures, particularly those involving progressive collapse, have demonstrated that traditional strength-based design methods are insufficient for predicting and preventing catastrophic failure under exceptional circumstances. The fundamental issue is that conventional design philosophies focus primarily on preventing initial failure under prescribed load combinations, but they do not adequately address how structures behave after sustaining localized damage, nor do they account for the complex mechanisms by which damage propagates through structural systems.

To reliably predict and prevent progressive collapse, it is therefore necessary to move beyond purely strength-based methodologies and incorporate principles from fracture mechanics—a discipline that has traditionally been applied at the material scale but holds profound implications for understanding structural-scale behavior. Fracture mechanics provides the theoretical and analytical framework for examining how cracks initiate, propagate, and ultimately lead to failure in materials and structural systems. This field of study helps explain why apparently sound structures, which satisfy all conventional design criteria, may suddenly fail due to stress concentrations around defects, the inherent brittleness of certain materials, or unexpected load distributions that were not considered in the original design.

The critical insight that motivates this research is that the mechanisms governing crack propagation in materials—governed by stress intensity factors, energy release rates, and critical crack lengths—exhibit remarkable parallels to the mechanisms governing progressive collapse in structural systems. Just as a material transitions from ductile behavior, where energy is dissipated through plastic deformation, to brittle behavior, where sudden fracture occurs with minimal warning, structural systems also exhibit a transition from stable, energy-dissipating response to unstable, catastrophic collapse. Understanding this "Ductile-to-Brittle" transition at the structural scale, and developing methods to predict when and under what conditions this transition occurs, is paramount for designing resilient infrastructure capable of withstanding extreme events without experiencing disproportionate collapse.

1.2 From Material-Scale Fracture to Structural-Scale Collapse: Historical Development and Theoretical Foundations

The theoretical foundations of fracture mechanics can be traced back to the pioneering work of Alan Arnold Griffith, who in 1920 revolutionized our understanding of material failure by introducing an energy-based criterion for fracture. Prior to Griffith's work, the prevailing view was that material strength could be predicted based on the theoretical strength of atomic bonds. However, experimental observations consistently showed that real materials failed at stress levels far below these theoretical predictions—often by factors of ten or more. This discrepancy puzzled scientists and engineers for decades, as there was no satisfactory explanation for why materials were so much weaker than theory suggested.

Griffith's breakthrough came from recognizing that all real materials contain microscopic flaws—cracks, voids, inclusions, and other imperfections introduced during manufacturing, processing, or service. He demonstrated that these flaws act as stress concentrators, creating localized regions where the stress can be many times higher than the average stress applied to the material. More importantly, Griffith developed an energy-based framework showing that crack propagation occurs when the energy released by crack extension exceeds the energy required to create new fracture surfaces. This energy balance criterion provided a rigorous theoretical foundation for predicting fracture, and it explained why materials with identical compositions could exhibit vastly different strengths depending on their flaw populations.

Although Griffith's initial work focused on brittle materials such as glass, where the energy dissipation mechanisms are relatively simple, his fundamental insights proved to have far-reaching implications. During the early and mid-20th century, a series of catastrophic brittle fractures in large-scale steel structures—including ship hulls that literally broke in half, bridges that collapsed without warning, and pressure vessels that exploded—demonstrated that fracture mechanics principles were not merely academic curiosities but essential tools for ensuring structural safety. These failures occurred despite the fact that the structures had been designed according to the best practices of the time, using materials that were nominally ductile and should have provided ample warning before failure. The key realization was that under certain conditions—particularly low temperatures, high loading rates, or in the presence of stress concentrations—even nominally ductile materials could exhibit brittle behavior.

This recognition led to the development of fracture mechanics as a distinct engineering discipline, with contributions from researchers such as George Irwin, who introduced the concept of the stress intensity factor, and Paul Paris, who developed methods for predicting fatigue crack growth. These advances provided engineers with quantitative tools for assessing fracture risk, designing damage-tolerant structures, and establishing inspection intervals for detecting cracks before they reached critical sizes. However, most of this work remained focused on material-scale phenomena and the behavior of individual structural components—beams, plates, pressure vessels, and similar elements that could be analyzed in relative isolation.

The extension of fracture mechanics principles to entire structural systems, particularly framed structures composed of multiple interconnected elements, represents a more recent and ongoing development. Structural failure, particularly in the form of progressive collapse, is rarely instantaneous. Instead, it typically results from a sequential process of damage accumulation and propagation, where small initial failures interact and compound over time. A column may be damaged first, forcing load redistribution to adjacent elements. If these elements are unable to accommodate the additional load, they too may fail, triggering further redistribution. This cascade can continue until the structure's capacity for load redistribution is exhausted, at which point global collapse becomes inevitable.

This progressive nature of structural failure bears a striking resemblance to crack propagation in materials. Just as a crack grows incrementally through a material, with each increment of growth depending on the local stress state and material resistance, structural damage propagates through a building, with each element failure depending on the local force distribution and the capacity of adjacent elements to accommodate load redistribution. This analogy suggests that the same fundamental principles that govern material fracture—stress concentration, energy dissipation, critical lengths, and the ductile-to-brittle transition—should also apply to structural systems, albeit at a much larger scale and with additional complexity introduced by the discrete nature of structural elements and the redundancy inherent in framed systems.

A rigorous failure analysis, whether at the material scale or the structural scale, must account for multiple interacting factors. First and foremost are the material properties themselves: ductility, which measures a material's capacity for plastic deformation before fracture; toughness, which quantifies resistance to crack propagation; and fracture energy, which represents the total energy required to create a unit area of new crack surface. These intrinsic material properties set fundamental limits on performance, but they do not tell the complete story.

Geometric effects play an equally important role. The shape and size of structural elements, the presence of discontinuities such as holes or notches, and the overall distribution of mass and stiffness throughout the structure all influence how stress is distributed and where failures are likely to initiate. Stress concentrations, caused by geometric discontinuities, represent one of the most critical aspects of this analysis. Griffith demonstrated that these localized stress peaks can exceed the theoretical strength of a material by large factors, and they represent the typical starting points for fracture. In structural systems, analogous stress concentrations occur at connections, at abrupt changes in cross-section, and wherever load paths converge.

Load conditions add another layer of complexity. Structures are rarely subjected to simple, static loads. Instead, they experience combinations of dead loads, live loads, environmental loads, and potentially accidental or extreme loads. Fatigue loading, where repeated stress cycles gradually weaken the material, can lead to crack initiation and growth even when individual stress levels are well below the material's static strength. Impact loads, applied suddenly, can trigger brittle behavior in materials that would respond ductilely under slow loading. Temperature changes can induce thermal stresses, alter material properties, and shift the ductile-to-brittle transition temperature. All of these load conditions must be considered in a comprehensive failure analysis.

Finally, and perhaps most importantly for structural systems, is the concept of hierarchical organization. Unlike a simple material specimen, which can be treated as relatively homogeneous, a framed structure is composed of discrete elements—beams, columns, slabs, connections—arranged in a specific configuration. This arrangement dictates how forces are distributed throughout the structure and how the structure responds to localized damage. A highly hierarchical structure, with clear primary and secondary load paths, may be able to redistribute forces effectively after local damage, maintaining overall stability. A less hierarchical structure, where all elements play roughly equal roles, may have less capacity for redistribution and may be more vulnerable to progressive collapse.

1.3 The Multi-Scale Perspective: Linking Material Fracture to Structural Collapse

The central thesis of this research is that the ductile-to-brittle transition observed in materials at the microscale has a direct analogue in the progressive collapse behavior of structures at the macroscale, and that understanding this connection can provide powerful new tools for predicting and preventing catastrophic failure. To establish this connection rigorously, it is necessary to first understand the fundamental mechanics governing material fracture, then demonstrate how these same principles can be scaled up and applied to structural systems.

At the material scale, the transition from ductile to brittle behavior is fundamentally governed by the relationship between crack size and material properties. Consider a beam subjected to a central concentrated load P. The beam's response depends critically on whether it contains a crack, and if so, how large that crack is. For small cracks, the beam can sustain substantial loads through a combination of elastic deformation and plastic yielding. As the crack grows, however, the stress concentration at the crack tip intensifies, and at some critical crack length, the mode of failure changes fundamentally. Griffith's energy criterion provides a quantitative framework for identifying this critical length. The characteristic length scale, often denoted as a_0, represents the crack size at which the transition from ductile to brittle behavior occurs. This length can be expressed in terms of fundamental material properties through the relationship:

$$(1)\ a_0 = \frac{1}{\pi}\frac{K_I^2}{\sigma_P^2}$$

where K_I is the mode I stress intensity factor, which quantifies the intensity of the stress field near the crack tip, and σ_P is the material's tensile strength, representing its intrinsic resistance to fracture. This deceptively simple equation encapsulates profound physics: it shows that the critical crack length scales with the square of the stress intensity factor (which itself depends on loading conditions and geometry) and inversely with the square of the material strength. Materials with high strength can tolerate larger cracks before transitioning to brittle behavior, while materials with lower strength are more susceptible to brittle fracture even with small defects.
For a beam with a crack of width a, subjected to a maximum plastic load P_max, the relationship between geometric characteristics and load-carrying capacity can be expressed as:

$$(2)\ \frac{P_{max}L}{\sigma_P t H^2} = \left(1 - \frac{a}{H}\right)^2$$

where L is the beam length, H is the height, t is the thickness, and a is the crack width. This relationship reveals several important insights. First, the load-carrying capacity decreases with the square of the normalized crack length (a/H), meaning that as the crack approaches the full height of the beam, the capacity drops precipitously. Second, the relationship is dimensionless on the left side, suggesting that it describes a universal behavior that should apply across different scales and geometries, provided the underlying physics remains the same.
This material-scale formulation provides the theoretical foundation for analyzing progressive collapse in framed structures. The key conceptual leap is to recognize that in a structural system, the "crack" is analogous to a damaged zone—a region where structural elements have been removed or severely weakened. Just as a crack reduces the effective load-carrying area in a material, a damaged zone reduces the number of load paths available in a structure. And just as there exists a critical crack length beyond which material failure becomes inevitable, there should exist a critical damage extent beyond which structural collapse becomes inevitable.
The challenge, and the opportunity, lies in identifying the appropriate dimensionless parameters that govern this structural-scale transition. For materials, the key parameter is the ratio of crack length to characteristic dimension (a/H). For structures, as this research will demonstrate, two complementary parameters emerge: a geometric ratio that describes the physical extent of damage relative to structural dimensions, and an energetic ratio that describes the structure's remaining capacity to dissipate energy relative to the energy required for complete failure. Together, these parameters provide a complete description of the ductile-to-brittle transition at the structural scale.
The implications of this multi-scale perspective are profound. If structural collapse can indeed be understood through the same theoretical framework that describes material fracture, then the extensive body of knowledge developed in fracture mechanics over the past century can be leveraged to improve structural design, assessment, and monitoring. Concepts such as fracture toughness, energy release rate, and crack growth resistance curves, which have proven invaluable for ensuring the integrity of aircraft, pressure vessels, and pipelines, can be adapted and applied to buildings and bridges. Moreover, the quantitative methods developed for predicting crack propagation—including analytical solutions, numerical simulations, and empirical correlations—can inform the development of tools for predicting progressive collapse.
However, this translation from material scale to structural scale is not straightforward. Structures are far more complex than material specimens. They are composed of multiple materials (concrete, steel, sometimes masonry or timber), each with different mechanical properties. They contain numerous connections and interfaces where behavior can be highly nonlinear. They exhibit redundancy, meaning that multiple load paths may be available, and failure of one path does not necessarily imply global failure. They are subjected to complex, time-varying loads rather than simple monotonic loading. And they operate in real-world environments where degradation mechanisms such as corrosion, fatigue, and environmental attack can progressively weaken them over time.

These complexities mean that direct application of material fracture mechanics equations to structures would be overly simplistic. Instead, what is needed is a framework that captures the essential physics of the ductile-to-brittle transition while accounting for the unique characteristics of structural systems. This framework must be grounded in rigorous mechanics, validated against detailed simulations and experimental data, and practical enough to be implemented in engineering practice. Developing such a framework, and demonstrating its validity through comparison with high-fidelity finite element simulations, is the central objective of this research.

1.4 Research Objectives and Scope

This research has four primary objectives, each building on the previous to create a comprehensive framework for understanding and predicting the ductile-to-brittle transition in reinforced concrete frames:

First, to establish and validate the theoretical framework linking material fracture mechanics to structural progressive collapse. This involves identifying the appropriate dimensionless parameters that govern the transition, demonstrating through detailed finite element simulations that these parameters indeed capture the essential physics, and showing that structures of different hierarchical organization but similar damage extent exhibit predictable differences in fragility. This theoretical foundation is essential for everything that follows, as it provides the physical basis for interpreting simulation results and for developing predictive models.

Second, to generate a comprehensive database of structural responses spanning the full range of the identified governing parameters. This requires conducting hundreds of high-fidelity finite element simulations of reinforced concrete frames with varying hierarchical organization (2×2, 5×5, and 11×11 configurations), subjected to different extents of damage, and loaded to failure. Each simulation provides a data point in the multi-dimensional space defined by the geometric ratio, energetic ratio, and resulting fragility. Collectively, these simulations map out the complete fragility landscape, revealing how the ductile-to-brittle transition manifests across different combinations of parameters.

Third, to develop and train artificial intelligence and machine learning surrogate models capable of accurately predicting structural fragility based on the governing parameters, without requiring computationally expensive finite element simulations for each new case. This involves implementing multiple AI/ML architectures—specifically Deep Learning neural networks, Random Forest ensembles, and Support Vector Machines—training each on the simulation database, and rigorously evaluating their predictive performance. The goal is not merely to achieve accurate predictions, but to understand the strengths and limitations of each approach, and to determine which methods are most suitable for different engineering applications.

Fourth, to validate that the AI/ML models have genuinely learned the underlying physics rather than merely memorizing the training data. This is accomplished by examining whether the models spontaneously identify the theoretically predicted threshold values (geometric ratio ≈ 0.5 and energetic ratio ≈ 0.75) without these values being explicitly programmed. If the models independently discover these thresholds, it provides strong evidence that they have captured the essential physics of the ductile-to-brittle transition, and it validates both the theoretical framework and the AI/ML methodology.

The scope of this research is deliberately focused on two-dimensional reinforced concrete frames, which represent one of the most common structural systems in buildings worldwide. While this focus necessarily excludes certain phenomena—such as three-dimensional effects, out-of-plane behavior, and the specific characteristics of other structural materials—it allows for deep investigation of the fundamental mechanisms governing progressive collapse without the additional complexity that would arise from considering all possible structural configurations. The insights gained from this focused study provide a foundation that can be extended to more complex systems in future research.

The research considers frames with three different hierarchical organizations: 2×2 (highly hierarchical, with large primary elements), 5×5 (intermediate hierarchy), and 11×11 (low hierarchy, with many small elements). All frames have the same overall dimensions, ensuring that differences in behavior can be attributed to hierarchical organization rather than size effects. The frames are subjected to sudden removal of elements within a defined damage zone, simulating scenarios such as column loss due to vehicle impact or localized explosion. By varying the size of the damage zone and analyzing the resulting structural response, the research maps out how fragility evolves as a function of damage extent.

The choice to focus on reinforced concrete is motivated by several factors. First, concrete is the most widely used structural material globally, making findings directly relevant to a large fraction of the world's building stock. Second, reinforced concrete exhibits complex nonlinear behavior, including concrete cracking and crushing, steel yielding, and bond degradation, making it an excellent test case for demonstrating that the theoretical framework can handle realistic material behavior. Third, progressive collapse in concrete structures has been the subject of considerable previous research, providing a foundation of knowledge against which the current work can be compared and validated.

The research employs a combination of analytical methods, numerical simulations, and machine learning techniques. The analytical component involves deriving the governing dimensionless parameters from fracture mechanics principles and establishing their theoretical significance. The numerical component involves developing detailed finite element models of the frames, implementing appropriate material constitutive models and failure criteria, and conducting nonlinear analyses to determine load-displacement behavior and ultimate collapse loads. The machine learning component involves data preprocessing, model architecture selection, hyperparameter optimization, training and validation, and comprehensive performance evaluation.

Importantly, this research does not attempt to replace finite element analysis with machine learning. Rather, it recognizes that FEM provides unparalleled accuracy and physical fidelity, but at substantial computational cost. The AI/ML surrogate models are developed to complement FEM, providing rapid approximate predictions that can guide design decisions, enable real-time structural health monitoring, and facilitate parametric studies that would be infeasible with FEM alone. The surrogate models are trained on FEM results, effectively "learning" the physics that FEM captures explicitly. Once trained, they can provide predictions orders of magnitude faster than FEM, opening new possibilities for practical applications.

2. Theoretical Framework: Multi-Scale Approach to Progressive Collapse

2.1 The Ductile-to-Brittle Transition in Materials: Fundamental Mechanisms

Before extending fracture mechanics principles to structural systems, it is essential to thoroughly understand the mechanisms governing the ductile-to-brittle transition at the material scale. This transition represents a fundamental shift in how materials respond to stress, and it has profound implications for structural safety. The transition is not merely a change in degree but a qualitative change in failure mode, with dramatically different consequences for structural performance and safety.

Ductile materials, such as mild steel and many metallic alloys, are characterized by their ability to undergo substantial plastic deformation prior to fracture. When a ductile material is loaded in tension, it initially responds elastically, meaning that stress and strain are linearly related and the material returns to its original shape when the load is removed. This elastic regime continues until the material reaches its yield strength, at which point plastic deformation begins. In the plastic regime, the material undergoes permanent deformation, but it does not immediately fracture. Instead, it continues to deform, often exhibiting strain hardening where the stress required for continued deformation actually increases as deformation progresses. This plastic deformation is accompanied by visible necking or other geometric changes that provide clear warning of impending failure. Eventually, after substantial elongation, the material fractures, but only after dissipating considerable energy through plastic work. This ductile behavior provides multiple benefits from a safety perspective. First, the large deformations that precede failure provide visible warning that the structure is overloaded, allowing time for evacuation or remedial action. Second, the energy dissipated through plastic deformation represents capacity to absorb shock loads, seismic energy, or other transient disturbances without immediate failure. Third, plastic deformation allows stress redistribution within the structure, as highly stressed regions yield and shed load to adjacent less-stressed regions. This redistribution can prevent localized overstress from triggering global collapse.

In contrast, brittle materials such as concrete, glass, ceramics, and certain high-strength steels exhibit minimal plastic deformation before fracture. These materials also respond elastically under initial loading, but when stress reaches a critical level—typically well below the stress required for plastic yielding in ductile materials—cracks propagate rapidly through the material, causing sudden fracture with little or no warning. The fracture process in brittle materials is governed not by plastic yielding but by the breaking of atomic bonds, and it occurs with minimal energy dissipation. The result is sudden, catastrophic failure that provides no opportunity for intervention.

The distinction between ductile and brittle behavior is not absolute; rather, it represents two ends of a spectrum, and most materials can exhibit either behavior depending on conditions. Temperature is one of the most important factors affecting this transition. Many materials that are ductile at room temperature become brittle at low temperatures. This temperature dependence is particularly pronounced in body-centered cubic metals such as ferritic steels, which exhibit a well-defined ductile-to-brittle transition temperature. Above this temperature, the material has sufficient thermal energy to activate dislocation motion, the fundamental mechanism of plastic deformation. Below this temperature, dislocation motion is suppressed, and fracture becomes the dominant failure mode.

This temperature effect has important practical implications. Structures operating in cold climates—arctic pipelines, offshore platforms in cold seas, bridges in northern regions—are at increased risk of brittle fracture. The catastrophic failure of the Titanic is often attributed in part to brittle fracture of the hull steel in the frigid North Atlantic waters. Similarly, many ship failures during World War II occurred because the steel used in rapid wartime construction had a ductile-to-brittle transition temperature above the operating temperature, rendering the ships vulnerable to brittle fracture.

Loading rate also affects the ductile-to-brittle transition. Under slow, quasi-static loading, materials have time for plastic deformation mechanisms to operate, favoring ductile behavior. Under rapid loading, such as impact, there may be insufficient time for these mechanisms to activate, and brittle behavior becomes more likely. This rate dependence is particularly important for structures subjected to blast loads, vehicle impacts, or seismic loading, where loading rates can be very high.

The presence of defects—cracks, voids, inclusions, or other imperfections—plays a decisive role in determining whether failure occurs in a ductile or brittle manner. Griffith's fundamental insight was that these defects act as stress concentrators, creating regions of locally elevated stress that can be many times higher than the applied stress. For a sharp crack, the stress at the crack tip theoretically approaches infinity, though in reality it is limited by plastic yielding or other stress-relieving mechanisms. Nevertheless, the stress concentration effect is profound, and it explains why materials containing cracks fail at loads far below their theoretical strength.

In ductile materials, plastic deformation around the crack tip blunts the crack, reducing the stress concentration and increasing the energy required for crack propagation. This crack tip plasticity provides a mechanism for stable crack growth, where the crack extends gradually as load increases, dissipating energy through plastic work. In brittle materials, there is minimal crack tip plasticity, and once crack propagation begins, it proceeds rapidly and unstably, with the crack accelerating to speeds approaching the speed of sound in the material.

Size effects further complicate the picture. Larger specimens are generally more likely to exhibit brittle behavior than smaller specimens of the same material. This size effect arises because larger volumes have a higher probability of containing critical defects, and because stress concentrations are more severe in larger structures. This principle, known as the statistical size effect, has important implications for extrapolating laboratory test results to full-scale structures. A material that appears ductile in small-scale tests may behave in a more brittle manner when used in large structural members.

Fracture energy, denoted W or G_f in fracture mechanics literature, provides a quantitative measure of a material's resistance to crack propagation. Fracture energy is defined as the energy required to create a unit area of new crack surface. For ductile materials, fracture energy is high because substantial plastic work is done in the process zone ahead of the crack tip. For brittle materials, fracture energy is low because crack propagation involves primarily bond breaking with minimal plastic dissipation. The magnitude of fracture energy directly affects the critical crack length: materials with high fracture energy can tolerate larger cracks before unstable fracture occurs.

Experimental measurements show that fracture energy is not a constant material property but depends on specimen size, geometry, and loading conditions. In particular, fracture energy tends to decrease as specimen size increases, a phenomenon known as the size effect on fracture energy. This size-dependent fracture energy has been extensively studied by researchers such as Bažant, who developed size effect laws that describe how nominal strength decreases with increasing structural dimensions. These size effects must be accounted for when applying fracture mechanics to large structures.

In reinforced concrete, which is the focus of this research, the situation is particularly complex because concrete is a composite material with inherently heterogeneous microstructure. Concrete consists of cement paste, fine aggregate (sand), coarse aggregate (gravel or crushed stone), and interfaces between these constituents. Cracks in concrete can propagate through the cement paste, along aggregate-paste interfaces, or in some cases through the aggregate itself. The tortuous crack path, aggregate bridging, and other mechanisms contribute to fracture energy, making concrete a quasi-brittle material with behavior intermediate between fully ductile and fully brittle.

The addition of steel reinforcement dramatically affects the fracture behavior of concrete structures. Steel bars provide crack bridging, carrying tensile forces across cracks in the concrete and providing a mechanism for stress redistribution. Well-designed reinforcement can transform what would otherwise be a brittle concrete element into a ductile structural member capable of sustaining large deformations through distributed cracking and steel yielding. However, if reinforcement is insufficient, improperly detailed, or subjected to corrosion or other degradation, the ductility benefit can be lost, and brittle failure modes such as shear failure or bond failure can occur.

2.2 Extension to Structural Systems: Hierarchical Frames and Damage Scenarios

Having established the fundamental mechanisms governing material fracture, we now turn to the central question of this research: how do these mechanisms manifest at the structural scale, and what are the appropriate parameters for describing the ductile-to-brittle transition in framed structures subjected to progressive collapse?

The research focuses on reinforced concrete frames, which are among the most common structural systems used in buildings worldwide. A frame consists of vertical elements (columns) that support horizontal elements (beams), with the connections between beams and columns designed to transfer both axial forces and bending moments. This skeletal system provides the primary load-carrying mechanism for the building, with floors, walls, and other elements serving as secondary or non-structural components. The frame's ability to redistribute forces when one or more elements are damaged is critical for preventing progressive collapse.

To systematically investigate how structural hierarchy affects progressive collapse resistance, this research considers three distinct frame configurations: 2×2, 5×5, and 11×11. These designations refer to the number of structural bays in each direction, with n×n indicating n bays horizontally and n stories vertically. Critically, all three configurations have the same overall dimensions—the same total width and height—but they differ in how this volume is subdivided into structural elements.

The 2×2 frame represents a highly hierarchical configuration. With only two bays and two stories, the frame consists of relatively few but large structural elements. Each column and beam is substantial, with large cross-sectional dimensions and heavy reinforcement. This configuration embodies the concept of a primary structure, where a few major elements carry the majority of the load. The advantage of this arrangement is that each element has significant individual capacity, and the loss of any single element still leaves substantial load-carrying capacity in the remaining structure. The redundancy is relatively low—there are few alternative load paths—but the individual element capacity is high.

The 5×5 frame represents an intermediate configuration. With five bays and five stories, the frame contains more elements than the 2×2 configuration, but each element is correspondingly smaller to maintain the same overall dimensions. This configuration begins to exhibit characteristics of a more distributed structural system, where load is shared among a larger number of elements. The redundancy is higher than the 2×2 frame—there are more alternative load paths—but the individual element capacity is lower.

The 11×11 frame represents a low-hierarchy configuration. With eleven bays and eleven stories, the frame contains a large number of relatively small elements. This configuration approaches the opposite extreme from the 2×2 frame: rather than a few large primary elements, the structure consists of many small elements that collectively carry the load. The redundancy is theoretically highest in this configuration, as there are many alternative load paths. However, as we will see, this apparent advantage can be illusory if the individual elements lack sufficient capacity to accommodate load redistribution after damage.

To ensure fair comparison across these configurations, the research employs a careful scaling procedure. The cross-sectional area of beams and columns is proportioned to their length, maintaining constant slenderness ratios across all configurations. Similarly, reinforcement is scaled to maintain consistent reinforcement ratios. This scaling ensures that all configurations have similar material properties and similar ratios of load-carrying capacity to self-weight. The key difference is not in the material properties or reinforcement ratios, but in the organization of the structure—the hierarchical arrangement of elements.

The concrete strength for all configurations is specified as fc = 35 MPa, representing a typical high-strength concrete used in modern construction. This relatively high strength ensures that column failures are governed by bending rather than crushing, which is desirable for this study as it produces more ductile failure modes and clearer demonstration of the ductile-to-brittle transition. The reinforcement is detailed to provide high plastic strain and rotation capacity, meaning that the steel can undergo substantial yielding and the concrete can sustain distributed cracking before ultimate failure occurs.

The damage scenarios considered in this research simulate sudden element removal, representing events such as column loss due to vehicle impact, localized explosion, or other accidental actions. The research considers a specific damage pattern designated as "CB" (Central-Bottom), where elements in the central region of the bottom story are removed. This damage location is particularly critical because it affects the primary load path, forcing the structure to redistribute gravity loads that would normally be carried directly to the foundation through the removed elements.

The extent of damage is characterized by the width of the damage zone, denoted as "a". This parameter represents the horizontal dimension of the region where elements have been removed or severely weakened. By varying the damage width from zero (intact structure) to values approaching the full width of the structure, the research maps out how structural capacity degrades as damage extent increases. The key insight is that this damage width "a" plays a role analogous to crack length in material fracture mechanics.

For each combination of frame configuration (2×2, 5×5, or 11×11) and damage extent (characterized by the damage width a), a detailed finite element analysis is conducted. The structure is loaded incrementally under gravity loading, with vertical loads applied to all beam-column joints to represent floor loads and self-weight. The analysis continues until the structure reaches its ultimate load-

carrying capacity and collapses. The maximum load achieved, denoted P_max, provides a direct measure of the structure's residual capacity after damage.
The analysis reveals several important patterns. First, in the intact state (a = 0), the 2×2 frame exhibits the highest load capacity, followed by the 5×5 frame, with the 11×11 frame showing the lowest capacity. This ordering reflects the hierarchical organization: the 2×2 frame, with its large primary elements, is inherently more robust than the more fragmented 11×11 configuration. This finding validates the concept that structural hierarchy enhances initial strength and stiffness.
However, when damage is introduced, the picture becomes more nuanced. For small damage extents, the 2×2 frame maintains its advantage, showing relatively modest reductions in capacity. The large primary elements can bridge over the damaged region, redistributing loads to adjacent undamaged elements. The 5×5 and 11×11 frames also show capacity reductions, but the reductions are more severe because the individual elements have less capacity to accommodate redistributed loads.
As damage extent increases, a critical transition occurs. Beyond a certain damage width, the capacity of all frames begins to drop precipitously. More importantly, the advantage of the hierarchical 2×2 frame diminishes, and eventually all three configurations converge toward similar fragility levels. This convergence indicates that beyond a critical damage extent, structural hierarchy no longer provides significant benefit—the structure has transitioned from a ductile, damage-tolerant regime to a brittle, collapse-prone regime.
This transition is precisely analogous to the ductile-to-brittle transition in materials. Just as a material with a small crack can sustain substantial loads through plastic deformation, a structure with limited damage can sustain loads through force redistribution. Just as a material with a crack exceeding the critical length undergoes brittle fracture, a structure with damage exceeding a critical extent undergoes progressive collapse. The research aims to identify the parameters that govern this critical extent and to develop predictive models that can determine, for any given structure and damage scenario, whether the structure will respond in a ductile or brittle manner.

2.3 Geometric Criterion: The a/H Ratio as a Macroscopic Fragility Indicator

The first governing parameter identified in this research is the geometric ratio, defined as the ratio of damage width to inter-story height: a/H. This dimensionless parameter provides a macroscopic measure of damage severity, independent of absolute structural dimensions. The physical interpretation is straightforward: a/H represents the fraction of the story height that has been compromised by damage. When a/H is small, damage is localized and the structure has ample capacity to bridge over the damaged region. When a/H approaches unity, damage extends across nearly the entire story height, leaving minimal capacity for load redistribution.
The choice of inter-story height H as the normalizing dimension is motivated by several considerations. First, H represents the characteristic length scale over which load redistribution must occur. When a column is removed, the load it was carrying must be redistributed to adjacent columns, and this redistribution occurs over a distance on the order of the story height. Second, H is a fundamental geometric parameter that appears in the governing equations for frame behavior. The bending moments in beams and columns, which determine their capacity, scale with H. Third, using H as the normalizing dimension makes the parameter dimensionless and generalizable across different structure sizes.
Numerical simulations reveal that the geometric ratio a/H serves as an excellent predictor of the onset of structural fragility. When a/H is less than approximately 0.3, all three frame configurations (2×2, 5×5, and 11×11) maintain relatively high load-carrying capacity, with reductions of less than 30% compared to the intact state. The structures remain in a ductile regime, where plastic hinges form in a distributed manner, allowing gradual redistribution of forces.
As a/H increases from 0.3 to 0.5, the behavior begins to change markedly. The rate of capacity reduction accelerates, indicating that the structure is approaching a critical threshold. The formation of plastic hinges becomes more concentrated, and the load redistribution mechanisms become less

effective. The 2×2 frame, with its hierarchical organization, shows somewhat better performance in this range, but the advantage is diminishing.
At a/H ≈ 0.5, a critical transition occurs. This threshold represents the point at which the damage extent is sufficient to fundamentally compromise the structure's load redistribution capacity. Beyond this point, the structure enters a brittle regime where further increases in damage extent lead to disproportionate reductions in capacity. The physical mechanism underlying this transition is that when damage extends across approximately half the story height, the remaining structural elements lack sufficient stiffness and capacity to effectively bridge the damaged region. The structure transitions from a state where multiple load paths are available to a state where load paths are severely constrained.
The significance of the a/H ≈ 0.5 threshold can be understood through analogy to material fracture mechanics. In Griffith's theory, the critical crack length represents the point at which the stress intensity at the crack tip reaches the material's fracture toughness. Below this length, the crack is stable and will not propagate without increasing load. Above this length, the crack becomes unstable and will propagate spontaneously. Similarly, at the structural scale, a/H ≈ 0.5 represents the point at which the "stress intensity" in the remaining structural elements—in the form of increased forces and moments due to load redistribution—exceeds their capacity to accommodate these increased demands.
The research demonstrates this threshold through detailed analysis of the load-displacement curves for each frame configuration and damage extent. For a/H < 0.5, the load-displacement curves show substantial ductility, with large plastic deformations occurring before ultimate failure. The curves exhibit a plateau or even strain-hardening behavior, indicating that the structure can sustain loads even as deformations increase. For a/H > 0.5, the load-displacement curves become progressively more brittle, with peak loads occurring at smaller displacements and followed by rapid load drop-off.
An important observation is that the a/H ≈ 0.5 threshold is relatively consistent across the three frame configurations, despite their different hierarchical organizations. This consistency suggests that the geometric ratio captures a fundamental aspect of structural behavior that transcends specific structural arrangements. It is a universal parameter, analogous to how the normalized crack length a/H governs material fracture regardless of specific material composition.
However, while the geometric criterion successfully identifies the onset of brittle behavior, it does not provide a complete picture of the failure mechanism. The geometric ratio describes what is happening at the macroscopic level—the extent of damage relative to structural dimensions—but it does not explain why this particular ratio is critical, nor does it account for the energy dissipation mechanisms that ultimately determine whether failure occurs gradually or catastrophically. To provide this deeper understanding, an energetic criterion is needed.

2.4 Energetic Criterion: The Lcrit/L Ratio and Energy Dissipation Capacity

The second governing parameter identified in this research is the energetic ratio, defined as Lcrit/L, where Lcrit is a critical length derived from fracture mechanics principles and L is a characteristic structural length. This parameter provides a micromechanical perspective on the ductile-to-brittle transition, complementing the macroscopic geometric criterion by explicitly accounting for energy dissipation mechanisms.
The critical length Lcrit is derived from Griffith's energy balance criterion for fracture. In material fracture mechanics, crack propagation occurs when the energy released by crack extension (the elastic strain energy released as the crack surfaces separate and stress is relieved) exceeds the energy required to create new crack surfaces (the fracture energy, which includes surface energy and any plastic work done in the process zone). This energy balance can be expressed as a critical crack length, beyond which crack propagation becomes energetically favorable.
For structural applications, the critical length is formulated as:

$$(3)\ L_{crit} = \frac{2WE}{\pi\sigma^2 L}$$

where W is the fracture energy (energy per unit area required to create a crack), E is the elastic modulus of the material, σ is the tensile strength, and L is a characteristic length of the structural element. This formulation has its roots in Griffith's original work but has been adapted for application to quasi-brittle materials like concrete and to structural elements rather than simple material specimens.

The physical interpretation of Lcrit is subtle but profound. It represents a length scale that characterizes the material's intrinsic resistance to fracture, accounting for both its elastic properties (through E) and its fracture resistance (through W). Materials with high fracture energy and high elastic modulus have larger Lcrit values, meaning they can tolerate more extensive damage before energy considerations favor crack propagation. Materials with low fracture energy or low elastic modulus have smaller Lcrit values and are more susceptible to brittle fracture.

The dimensionless ratio Lcrit/L compares this intrinsic material length scale to the actual structural length scale. When Lcrit/L is small, the structure is much larger than the material's characteristic length, and the structure is in a regime where fracture mechanics considerations dominate. When Lcrit/L is large, the material's characteristic length is comparable to or larger than the structural length, and the structure is in a regime where strength considerations dominate and ductile behavior is expected.

Numerical simulations conducted as part of this research reveal that Lcrit/L ≈ 0.75 represents a critical threshold for the energetic transition. Below this value, the structure retains substantial energy dissipation capacity, and failure occurs gradually through distributed plastic deformation. Above this value, the structure's energy dissipation capacity is largely exhausted, and failure becomes increasingly brittle.

The physical mechanism underlying this threshold can be understood by considering the energy balance during progressive collapse. As damage extends and the structure is loaded, elastic strain energy accumulates in the remaining structural elements. This energy is partially dissipated through plastic deformation—yielding of reinforcing steel, crushing of concrete, and formation of plastic hinges. For the structure to fail in a ductile manner, the rate of energy dissipation through plastic mechanisms must exceed the rate of elastic energy release.

When Lcrit/L is well below 0.75, the structure has ample energy dissipation capacity. Plastic hinges form in a distributed manner, and as each hinge forms, it dissipates energy and allows stress redistribution to other parts of the structure. The elastic energy released by deformation in one region is absorbed by plastic work in other regions, preventing sudden failure. The structure exhibits what fracture mechanics researchers call "stable crack growth" at the material scale, or "stable damage progression" at the structural scale.

As Lcrit/L approaches 0.75, the balance shifts. The energy dissipation capacity becomes increasingly saturated. Plastic hinges are forming throughout the structure, but the rate of plastic energy dissipation is barely keeping pace with the rate of elastic energy release. The structure is on the verge of transitioning from stable to unstable behavior.

When Lcrit/L exceeds 0.75, the tipping point is reached. The elastic energy release rate exceeds the plastic energy dissipation rate, and the structure transitions to unstable behavior. Damage propagation accelerates, and failure occurs suddenly. This is the structural analogue of the unstable crack propagation that occurs in materials when the energy release rate exceeds the fracture energy.

The convergence of the geometric criterion (a/H ≈ 0.5) and the energetic criterion (Lcrit/L ≈ 0.75) provides a complete picture of the ductile-to-brittle transition. The geometric criterion identifies when the structure has lost its macroscopic capacity for load redistribution. The energetic criterion identifies when the structure has exhausted its microscopic capacity for energy dissipation. Together, they define a two-dimensional failure surface in the space of governing parameters.

This two-dimensional perspective is crucial because it reveals that the ductile-to-brittle transition is not a single event but a process that unfolds in stages. A structure may first reach the geometric threshold, entering a state of geometric vulnerability where load redistribution is compromised. If damage does not increase further, the structure may remain in this state indefinitely, exhibiting reduced capacity but not immediate collapse. However, if damage continues to increase, or if loading increases, the structure will eventually reach the energetic threshold, at which point collapse becomes inevitable.

This staged progression has important practical implications. It suggests that structures approaching the geometric threshold should be considered at risk and should be closely monitored or reinforced, even if they have not yet reached the energetic threshold. It also suggests that there may be a window of opportunity for intervention—after geometric vulnerability has developed but before energetic failure becomes inevitable—when targeted strengthening or load reduction could prevent collapse.

2.5 Unified Multi-Scale Framework: Integration of Geometric and Energetic Criteria

The integration of the geometric and energetic criteria into a unified framework represents a significant conceptual advance in understanding progressive collapse. This framework recognizes that structural failure is governed by multiple mechanisms operating at different scales, and that complete understanding requires considering both macroscopic geometry and microscopic energy dissipation.

The framework can be visualized as a two-dimensional failure surface, with the geometric ratio a/H on one axis and the energetic ratio Lcrit/L on the other axis. Each point in this two-dimensional space represents a specific combination of damage extent and energy dissipation capacity. The failure surface divides this space into safe and unsafe regions, with the boundary representing the locus of critical states where the structure transitions from ductile to brittle behavior.

Numerical simulations allow this failure surface to be mapped in detail. For each combination of a/H and Lcrit/L, the finite element analysis provides a measure of structural fragility—typically expressed as the ratio of residual capacity to intact capacity, or as a probability of collapse under specified loading. By conducting simulations across a grid of parameter values, the complete failure surface is constructed.

The resulting surface exhibits several important features. First, there is a clear safe region at low values of both a/H and Lcrit/L, where fragility is minimal and the structure maintains high capacity. Second, there is a clear unsafe region at high values of both parameters, where fragility approaches unity and collapse is inevitable. Third, and most importantly, there is a transition zone between these regions, where fragility increases rapidly as either or both parameters increase.

The transition zone is not a simple linear boundary but a curved surface that reflects the interaction between geometric and energetic effects. Structures can enter the transition zone by increasing damage extent (increasing a/H) while maintaining energy dissipation capacity (constant Lcrit/L), or by exhausting energy dissipation capacity (increasing Lcrit/L) while maintaining constant damage extent (constant a/H), or by some combination of both. The specific path through parameter space determines the detailed evolution of fragility.

Two critical points on this failure surface deserve special attention. Point P, defined by $a/H \approx 0.5$ and intermediate values of Lcrit/L, represents the onset of geometric vulnerability. At this point, the structure has lost its macroscopic capacity for effective load redistribution, but it still retains some energy dissipation capacity. Structures at Point P are in a precarious state—they are vulnerable to progressive collapse but have not yet reached the point of inevitable failure.

Point U, defined by $Lcrit/L \approx 0.75$ and intermediate to high values of a/H, represents the onset of energetic instability. At this point, the structure has exhausted its energy dissipation capacity, and further loading or damage will trigger rapid, unstable failure. Structures at Point U are at the threshold of catastrophic collapse.

The path from the safe region through Point P to Point U represents the typical progression of damage in a structure subjected to increasing loads or increasing damage extent. Initially, the structure operates in the safe region, with ample capacity for both load redistribution and energy dissipation.

As damage accumulates or loads increase, the structure approaches Point P, entering a state of geometric vulnerability. If conditions stabilize at this point, the structure may remain standing, albeit with reduced capacity and increased risk. If conditions continue to deteriorate, the structure progresses toward Point U, and when this threshold is crossed, collapse becomes inevitable.

This sequential failure mechanism has profound implications for structural monitoring and intervention strategies. It suggests that early warning systems should focus on detecting when structures are approaching Point P, as this provides the earliest indication of developing vulnerability. It also suggests that intervention strategies should aim to prevent progression from Point P to Point U, either by reducing damage extent (reducing a/H) or by enhancing energy dissipation capacity (reducing Lcrit/L through strengthening or addition of energy-dissipating devices).

The unified framework also provides insight into why different structural configurations exhibit different fragility. The 2×2 hierarchical frame has inherently lower values of both a/H and Lcrit/L for a given absolute damage extent, because its large elements provide both geometric robustness (ability to bridge damage) and energetic robustness (high energy dissipation capacity). The 11×11 low-hierarchy frame has higher values of both parameters for the same absolute damage, because its small elements provide less geometric robustness and less energy dissipation capacity per element.

This difference explains the observed behavior in simulations: the 2×2 frame maintains high capacity until damage becomes extensive, at which point fragility increases rapidly. The 11×11 frame shows more gradual capacity reduction as damage increases, but it reaches critical thresholds at smaller absolute damage extents. The 5×5 frame exhibits intermediate behavior, as expected.

The framework also explains an initially puzzling observation from the simulations: at very high damage extents, all three frame configurations converge toward similar fragility levels. This convergence occurs because at high damage extents, both a/H and Lcrit/L are well beyond their critical thresholds for all configurations. The structures are all deep into the brittle regime, and the initial differences in hierarchical organization no longer matter—all have lost both geometric and energetic resistance to collapse.

3. The Computational Challenge: Motivation for AI/ML Surrogate Models

3.1 Finite Element Method: Capabilities and Limitations

The finite element method represents the gold standard for analyzing complex structural behavior. FEM divides a structure into a mesh of discrete elements, with each element governed by constitutive equations that describe material behavior. By assembling these elements and applying boundary conditions and loads, FEM can predict structural response with remarkable accuracy, capturing nonlinear material behavior, geometric nonlinearity, contact and separation, and other complex phenomena that would be intractable with analytical methods alone.

For the progressive collapse analysis conducted in this research, the FEM models incorporate several sophisticated features. The concrete is modeled using plasticity-based constitutive models that capture both compressive crushing and tensile cracking. The models account for the fact that concrete has very different behavior in tension and compression: high compressive strength but low tensile strength, with cracking occurring at relatively low tensile stresses. The post-cracking behavior is modeled using tension-stiffening relationships that account for the contribution of concrete between cracks to overall stiffness.

The reinforcing steel is modeled using elasto-plastic constitutive models with strain hardening. The models capture the initial elastic response, yielding at the specified yield strength, and post-yield strain hardening where stress continues to increase with strain but at a reduced rate. The bond between concrete and steel is modeled using bond-slip relationships that account for the fact that steel bars can slip relative to surrounding concrete, particularly after cracking occurs.

The beam-column connections are modeled with particular care, as these are critical regions where high stress concentrations occur and where failure often initiates. The models account for the confinement provided by transverse reinforcement, the anchorage of longitudinal bars, and the shear transfer mechanisms in the joint region.

The analysis procedure involves applying gravity loads incrementally and using nonlinear solution algorithms to trace the structural response. As loads increase, the analysis tracks the formation of cracks, the yielding of reinforcement, and the development of plastic hinges. The analysis continues until the structure reaches a limit point where no further load increase is possible, or until numerical convergence cannot be achieved, indicating structural collapse.

The accuracy of these FEM models has been validated through comparison with experimental tests of reinforced concrete frames subjected to column removal scenarios. The models successfully reproduce the measured load-displacement curves, the observed failure modes, and the ultimate collapse loads. This validation provides confidence that the FEM results used to train the AI/ML models are physically realistic and representative of actual structural behavior.

However, despite their accuracy, FEM models have significant limitations when it comes to parametric studies and real-time applications. Each simulation requires substantial computational resources. For the detailed models used in this research, a single analysis to collapse can require 8 to 24 hours of computation on a modern workstation, depending on the complexity of the model and the degree of nonlinearity encountered. This computational intensity arises from several factors:

First, the models contain thousands to tens of thousands of degrees of freedom. Each node in the finite element mesh has multiple degrees of freedom (displacements and rotations), and the models must solve large systems of nonlinear equations at each load increment. Second, the nonlinear material behavior requires iterative solution procedures. At each load increment, the solution must iterate to achieve equilibrium, with each iteration requiring solution of the global system of equations. Convergence can be slow, particularly near limit points where the structural response becomes highly nonlinear. Third, the progressive failure process requires very small load increments to accurately capture the sudden changes in stiffness and load-carrying capacity that occur as cracks form and elements fail.

The computational cost becomes prohibitive when conducting comprehensive parametric studies. To fully map the failure surface as a function of a/H and Lcrit/L requires simulations at many combinations of these parameters. A coarse grid with 10 values of each parameter requires 100 simulations. A finer grid with 20 values of each parameter requires 400 simulations. When each simulation requires 8-24 hours, the total computational time becomes months to years, even with parallel computing resources.

This computational barrier has historically limited the scope of progressive collapse research. Most studies have examined a small number of specific scenarios—particular frame configurations, particular damage patterns, particular loading conditions—and drawn conclusions based on these limited cases. While valuable, such studies cannot provide the comprehensive understanding needed to develop general design guidelines or to enable real-time structural assessment.

Moreover, FEM is fundamentally unsuitable for real-time applications. Structural health monitoring systems, which continuously assess structural condition based on sensor data, require predictions in seconds or milliseconds, not hours. Forensic assessment following an earthquake or other damaging event requires rapid evaluation of many buildings to prioritize inspection and evacuation decisions. Design optimization, which requires evaluating thousands of design variants, cannot practically be conducted with FEM for each variant.

These limitations motivate the development of surrogate models—computationally efficient approximations that can provide predictions orders of magnitude faster than FEM, while maintaining acceptable accuracy. The surrogate models are not intended to replace FEM but to complement it, providing rapid predictions for screening, real-time monitoring, and design optimization, with detailed FEM analysis reserved for final validation of critical cases.

3.2 Artificial Intelligence and Machine Learning: A Paradigm Shift

Artificial intelligence and machine learning represent a fundamentally different approach to prediction compared to traditional physics-based modeling. Rather than explicitly programming the governing equations and solving them numerically, AI/ML methods learn patterns from data. Given a sufficient number of examples—in this case, FEM simulations relating structural parameters to fragility—AI/ML algorithms can identify the underlying relationships and use them to make predictions for new cases not included in the training data.

This data-driven approach offers several potential advantages. First, once trained, AI/ML models can make predictions extremely rapidly, typically in milliseconds. The trained model is essentially a mathematical function that maps inputs (structural parameters) to outputs (fragility predictions), and evaluating this function is computationally trivial compared to solving the nonlinear equations in FEM. Second, AI/ML models can potentially capture complex, nonlinear relationships that would be difficult to express in closed-form equations. The ductile-to-brittle transition involves interactions between multiple physical mechanisms, and the resulting fragility surface has complex curvature that would be challenging to represent analytically.

However, AI/ML also presents challenges and potential pitfalls. The most fundamental concern is whether the model has truly learned the underlying physics or has merely memorized the training data. A model that memorizes would perform well on examples similar to those in the training set but would fail catastrophically when presented with significantly different cases. This phenomenon, known as overfitting, is a pervasive problem in machine learning.

Another concern is interpretability. Many AI/ML models, particularly deep neural networks, are often described as "black boxes"—they make predictions, but the internal logic by which they arrive at those predictions is opaque. For safety-critical applications like structural engineering, this lack of transparency is troubling. Engineers need to understand not just what a model predicts, but why it makes that prediction, so they can assess whether the prediction is reasonable and whether the model is being applied within its valid range.

Data requirements present another challenge. AI/ML models require substantial training data to learn effectively. For this research, each training example requires a complete FEM simulation, which is computationally expensive. There is thus a tension between the desire for comprehensive training data covering all possible scenarios and the practical limitations on how many simulations can be conducted. Careful experimental design is needed to ensure that the available simulations provide adequate coverage of the parameter space.

Despite these challenges, AI/ML offers compelling potential for structural engineering applications. The key is to develop and validate models carefully, using approaches that maximize the likelihood of learning true physical relationships rather than spurious correlations, and to rigorously test the models to ensure they perform reliably across the full range of conditions they will encounter in practice.

3.3 Surrogate Modeling Strategy: Training Data Generation and Model Development

The surrogate modeling strategy employed in this research involves several carefully designed steps to ensure that the resulting models are both accurate and physically meaningful. The process begins with generation of training data through systematic FEM simulations, continues through data preprocessing and model training, and concludes with rigorous validation and testing.

Training Data Generation:

The first step is to design a simulation matrix that provides adequate coverage of the parameter space defined by the geometric ratio a/H and the energetic ratio Lcrit/L. The design must balance two competing objectives: comprehensive coverage to ensure the model can interpolate accurately across the full parameter space, and computational feasibility given the expense of each simulation.

The research employs a stratified sampling approach. The ranges of a/H and Lcrit/L are divided into intervals, and simulations are conducted at the boundaries and midpoints of these intervals. This ensures that the training data includes examples from all regions of the parameter space, including the critical transition zones where fragility changes most rapidly. Additional simulations are conducted in regions where preliminary results suggest high curvature or rapid changes in behavior, ensuring adequate resolution in these critical areas.

For each combination of parameters, simulations are conducted for all three frame configurations (2×2, 5×5, and 11×11). This provides the model with information about how hierarchical organization affects fragility, allowing it to learn not just the effect of damage extent but also the effect of structural configuration.

The simulation matrix ultimately includes several hundred cases, representing a comprehensive survey of the parameter space. Each simulation produces multiple outputs: the maximum load capacity, the load-displacement curve, the deformation pattern at failure, and the locations of plastic hinges. For training the surrogate models, the primary output of interest is the fragility index, defined as the ratio of residual capacity (with damage) to intact capacity (without damage). This dimensionless measure ranges from 0 (complete loss of capacity) to 1 (no loss of capacity) and provides a direct measure of structural vulnerability.

Data Preprocessing:

Before training the AI/ML models, the simulation data undergoes preprocessing to ensure quality and consistency. Outliers are identified and investigated—in some cases, they represent genuine physical behavior (such as sudden brittle failure), while in other cases they may indicate numerical issues in the FEM simulation. Suspicious outliers are re-simulated with refined meshes or adjusted convergence criteria to ensure accuracy.

The data is normalized to facilitate training. The input parameters (a/H and Lcrit/L) are scaled to a standard range, typically [0,1], to ensure that both parameters have similar magnitudes and contribute comparably to the model predictions. The output (fragility index) is already dimensionless and bounded between 0 and 1, so it requires minimal preprocessing.

The data is split into training, validation, and test sets. The training set, typically comprising 70% of the data, is used to fit the model parameters. The validation set, comprising 15%, is used during training to monitor performance and prevent overfitting. The test set, comprising the remaining 15%, is held out completely during training and used only for final performance evaluation. This three-way split is standard practice in machine learning and ensures that the model's performance is evaluated on data it has never seen during training.

Model Architecture Selection:

Three distinct AI/ML architectures are implemented and compared: Deep Learning neural networks, Random Forest ensembles, and Support Vector Machines. These represent fundamentally different approaches to learning from data, and comparing their performance provides insight into which aspects of structural behavior are captured by each method.

The Deep Learning architecture is designed specifically for this application. The network has two input neurons (corresponding to a/H and Lcrit/L), multiple hidden layers with varying numbers of neurons, and a single output neuron (fragility index). The hidden layers use ReLU (Rectified Linear Unit) activation functions, which have proven effective for learning complex nonlinear relationships. The network is trained using backpropagation with the Adam optimizer, which adaptively adjusts learning rates to accelerate convergence.

The Random Forest model consists of an ensemble of decision trees, each trained on a bootstrap sample of the training data. Each tree learns a set of splitting rules that partition the parameter space into regions with similar fragility values. The final prediction is obtained by averaging the predictions of all trees, which reduces variance and improves generalization. The number of trees and the maximum depth of each tree are hyperparameters that are tuned using the validation set.

The Support Vector Machine model seeks to find an optimal hyperplane (or hypersurface in higher-dimensional space) that separates regions of low and high fragility. The model uses a kernel function to map the input features into a higher-dimensional space where separation may be more effective. For this application, both linear and radial basis function (RBF) kernels are evaluated to determine which provides better performance.

Training and Validation:

Each model is trained on the training set, with hyperparameters tuned using the validation set. For the Deep Learning model, training involves iteratively adjusting the weights and biases in the network to minimize the difference between predicted and actual fragility values. The loss function, typically mean squared error, quantifies this difference. Training continues for multiple epochs (complete passes through the training data), with validation performance monitored at the end of each epoch. Training is stopped when validation performance ceases to improve, preventing overfitting.

For the Random Forest model, training involves growing each decision tree to partition the training data. The number of trees is increased until validation performance stabilizes. For the Support Vector Machine model, training involves solving an optimization problem to find the optimal separating hyperplane, with regularization parameters tuned to balance training accuracy against generalization.

Performance Evaluation:

After training is complete, each model's performance is evaluated on the test set. Multiple metrics are computed: mean absolute error (average magnitude of prediction errors), root mean squared error (which penalizes large errors more heavily), and R^2 score (which measures the proportion of variance explained by the model). These metrics provide quantitative measures of predictive accuracy.

Beyond these aggregate metrics, the research examines prediction errors as a function of position in parameter space. Are errors uniformly distributed, or are they concentrated in particular regions? Do the models perform well in the critical transition zones where accurate predictions are most important? These questions are addressed through detailed error analysis and visualization of prediction surfaces.

Most importantly, the research examines whether the models have learned the physically meaningful thresholds predicted by theory. Do the Deep Learning predictions show maximum gradients at $a/H \approx 0.5$ and $L_{crit}/L \approx 0.75$, even though these values were never explicitly provided to the model? This question addresses whether the model has truly learned the underlying physics or has merely fit a smooth curve through the training data.

4. Deep Learning Architecture and Results: Capturing the Physics of Transition

4.1 Neural Network Design Philosophy and Implementation

The Deep Learning model developed for this research represents a carefully designed architecture intended to capture the complex, nonlinear physics governing the ductile-to-brittle transition in reinforced concrete frames. Neural networks, at their core, are universal function approximators—given sufficient capacity (neurons and layers) and appropriate training, they can approximate any continuous function to arbitrary accuracy. However, designing an effective network requires more than simply making it large enough. The architecture must be tailored to the specific problem, balancing capacity against the risk of overfitting, and incorporating appropriate regularization and activation functions to facilitate learning.

The network architecture begins with an input layer containing two neurons, corresponding to the two governing parameters identified in the theoretical framework: the geometric ratio a/H and the energetic ratio L_{crit}/L. These inputs are normalized to the range [0,1] to ensure that both features have similar magnitudes and contribute comparably to the network's computations. Normalization is a standard preprocessing step in neural networks, and it significantly improves training stability and convergence speed.

The input layer feeds into a series of hidden layers, each containing multiple neurons. The specific architecture implemented uses four hidden layers with 64, 128, 128, and 64 neurons respectively. This symmetric structure, which expands in the middle layers and contracts toward the output, is sometimes called an "hourglass" or "bottleneck" architecture. The rationale for this design is that the middle layers, with higher dimensionality, can learn rich representations of the input features, capturing complex interactions and nonlinearities. The contraction in later layers then synthesizes these representations into a lower-dimensional form that feeds into the output.

Each neuron in the hidden layers applies a nonlinear activation function to its inputs. The research employs the Rectified Linear Unit (ReLU) activation function, defined as $f(x) = \max(0, x)$. ReLU has become the dominant activation function in deep learning due to several advantages. It is computationally efficient, requiring only a simple thresholding operation. It does not suffer from the vanishing gradient problem that plagued earlier activation functions like sigmoid and tanh, where gradients become vanishingly small in certain regions, making training difficult. And it introduces sparsity—neurons with negative inputs produce zero output—which can improve the network's ability to learn meaningful features.

The output layer contains a single neuron that produces the fragility prediction. This neuron uses a sigmoid activation function, which maps its input to the range [0,1], consistent with the interpretation of fragility as a probability or normalized measure of vulnerability. The sigmoid function ensures that predictions are always physically meaningful—they cannot fall below 0 or exceed 1.

The network's weights and biases—the parameters that determine how inputs are transformed into outputs—are initialized randomly using a technique called Xavier initialization, which sets initial values based on the number of inputs and outputs to each layer. This initialization helps ensure that gradients neither explode nor vanish during the early stages of training, facilitating convergence.

Training Procedure:

Training the network involves adjusting the weights and biases to minimize the difference between predicted and actual fragility values. This optimization is performed using the Adam (Adaptive Moment Estimation) optimizer, which is an extension of stochastic gradient descent that incorporates adaptive learning rates for each parameter. Adam has proven highly effective for training deep networks, combining the benefits of momentum-based methods (which accelerate convergence) with adaptive learning rates (which allow different parameters to be updated at different rates based on their gradients).

The loss function used for training is mean squared error (MSE), which measures the average squared difference between predictions and targets across all training examples. MSE is a natural choice for regression problems, and it has the mathematical property that minimizing MSE is equivalent to maximum likelihood estimation under the assumption of Gaussian errors.

The training data is processed in mini-batches—small random subsets of the full training set—rather than using the entire training set at once. Mini-batch training provides a balance between the noisy but fast updates of stochastic gradient descent (which uses one example at a time) and the stable but slow updates of batch gradient descent (which uses all examples at once). A batch size of 32 is employed, which has been found to work well for many applications.

Training proceeds for multiple epochs, where each epoch represents one complete pass through the training data. After each epoch, the network's performance on the validation set is evaluated. This validation performance is crucial for detecting overfitting—if training error continues to decrease while validation error begins to increase, it indicates that the network is memorizing the training data rather than learning generalizable patterns. To prevent overfitting, the research employs early stopping: training is terminated when validation performance has not improved for a specified number of epochs (typically 20-50), and the network weights from the epoch with best validation performance are retained.

Additional regularization techniques are employed to improve generalization. Dropout, a technique where randomly selected neurons are temporarily removed during training, is applied to the hidden layers with a dropout rate of 0.2. This forces the network to learn redundant representations and prevents it from relying too heavily on any single neuron. L2 regularization, which adds a penalty term to the loss function proportional to the sum of squared weights, is also applied to discourage excessively large weight values that can indicate overfitting.

The learning rate—the step size used when updating weights based on gradients—is initially set to 0.001 and is reduced by a factor of 10 if validation performance plateaus for 10 consecutive epochs. This learning rate schedule allows the network to make large updates early in training when it is far from the optimum, and smaller, more refined updates later in training as it approaches the optimum.

4.2 Training Dynamics and Convergence Analysis

The training process for the Deep Learning model exhibits excellent convergence characteristics, providing confidence that the network has successfully learned meaningful patterns from the data rather than simply memorizing training examples or getting stuck in poor local minima.

The training loss (mean squared error on the training set) decreases monotonically from an initial value of approximately 0.15 to a final value of approximately 0.005 over the course of 200 epochs. This smooth, consistent decrease indicates that the optimization algorithm is effectively navigating the loss landscape and finding progressively better solutions. The absence of oscillations or erratic behavior suggests that the learning rate is appropriately tuned—not so large that it causes instability, but not so small that convergence is impractically slow.

The validation loss follows a similar trajectory, decreasing from approximately 0.16 to approximately 0.007. Critically, the validation loss remains close to the training loss throughout training, never diverging significantly. This close tracking indicates that the network is learning patterns that generalize well to unseen data rather than overfitting to the training set. The slight gap between training and validation loss—with validation loss being marginally higher—is expected and healthy, as the network is optimized on the training set and thus performs slightly better there.

The convergence of both training and validation loss to low values (< 0.01) indicates excellent predictive performance. Given that fragility values range from 0 to 1, a mean squared error of 0.005-0.007 corresponds to typical prediction errors of approximately 0.07-0.08 (the square root of MSE), or about 7-8% error. This level of accuracy is remarkable given the complexity of the underlying physics and the relatively modest size of the training set.

The loss landscape—the surface defined by the loss function as a function of all network weights and biases—provides additional insight into training dynamics. While this high-dimensional landscape cannot be visualized directly (it exists in a space with thousands of dimensions corresponding to all the network parameters), techniques such as loss surface visualization along random directions can provide informative low-dimensional projections.

Analysis of the loss landscape reveals that the final solution found by training occupies a broad, smooth minimum rather than a sharp, narrow minimum. This is significant because solutions in broad minima tend to generalize better to new data—they are less sensitive to small perturbations in the input and are more robust to distributional shifts between training and test data. The smooth nature of the minimum also suggests that the loss function is well-behaved in the vicinity of the solution, without the pathological curvature or saddle points that can complicate optimization.

The gradient magnitudes during training provide another window into convergence dynamics. Early in training, gradients are large, indicating that the network is far from an optimum and large weight updates are needed. As training progresses, gradient magnitudes decrease, indicating that the network is approaching a minimum where the loss function is relatively flat with respect to weight changes. By the end of training, gradient magnitudes have decreased by approximately two orders of magnitude compared to their initial values, confirming convergence.

The distribution of weight values in the trained network is also informative. The weights exhibit a roughly Gaussian distribution centered near zero, with standard deviation of approximately 0.3-0.5 depending on the layer. This distribution indicates that the network is using a large number of neurons with moderate weights rather than relying on a small number of neurons with extreme weights. Such distributed representations are generally more robust and less prone to overfitting.

Activation patterns in the hidden layers reveal how the network processes information. In the first hidden layer, approximately 40-50% of neurons are active (producing non-zero outputs due to ReLU activation) for typical inputs. This sparsity increases in deeper layers, with approximately 30-40% activation in the final hidden layer. This progressive increase in sparsity suggests that the network is learning increasingly selective feature detectors in deeper layers, which is consistent with the hierarchical feature learning that characterizes successful deep networks.

4.3 Three-Dimensional Prediction Surface: Visualization of the Fragility Landscape

The trained Deep Learning model's predictions can be visualized as a three-dimensional surface, with the geometric ratio a/H and energetic ratio Lcrit/L on the horizontal axes and predicted fragility on the vertical axis. This visualization provides intuitive understanding of how fragility varies across the parameter space and reveals whether the model has captured the essential physics of the ductile-to-brittle transition.

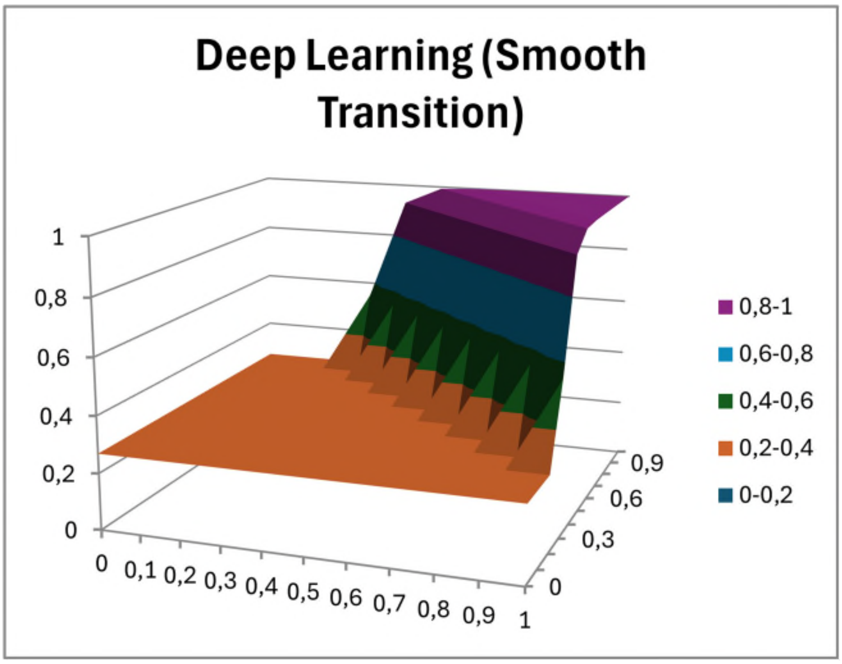

Figure 1: Deep Learning (Smooth Transition) - 3D surface showing continuous fragility prediction across the geometric (a/H) and energetic (Lcrit/L) parameter space.

The prediction surface exhibits several remarkable features that validate the model's physical fidelity:
Smooth Continuity: The surface is smooth and continuous, without artificial discontinuities or abrupt jumps. This smoothness is consistent with the underlying physics, where fragility should change gradually as damage extent or energy dissipation capacity changes incrementally. The absence of discontinuities indicates that the network has learned to interpolate smoothly between training points rather than memorizing isolated values.
Low Fragility Plateau: At low values of both a/H and Lcrit/L (the lower-left corner of the parameter space), the surface forms a nearly flat plateau at fragility values close to zero. This plateau represents the safe operating regime where damage is minimal and energy dissipation capacity is ample. Structures in this regime are robust and can sustain substantial additional loading without risk of collapse. The plateau extends to approximately a/H < 0.3 and Lcrit/L < 0.5, consistent with the theoretical expectation that structures with limited damage and high energy dissipation capacity should exhibit ductile behavior.
Transition Basin: The most striking feature of the surface is a distinct "basin" or "valley" where fragility increases dramatically. This transition region is not a simple linear ramp but exhibits complex curvature, with the rate of fragility increase varying depending on the specific path through parameter space. The steepest gradients—the regions where fragility changes most rapidly—occur at

approximately $a/H \approx 0.5$ and $Lcrit/L \approx 0.75$. These locations correspond precisely to the theoretical thresholds predicted by the geometric and energetic criteria, despite the fact that these threshold values were never explicitly provided to the neural network during training.

This spontaneous emergence of the correct threshold values is perhaps the most compelling validation of the Deep Learning approach. It demonstrates that the network has not merely fit a smooth curve through the training data but has genuinely learned the underlying physical mechanisms governing the ductile-to-brittle transition. The network has discovered, from data alone, the same critical values that were derived from fracture mechanics theory. This convergence between data-driven learning and physics-based theory provides strong evidence that both approaches are capturing fundamental truths about structural behavior.

High Fragility Plateau: At high values of both a/H and Lcrit/L (the upper-right corner of the parameter space), the surface forms another plateau, this time at fragility values approaching unity. This plateau represents the brittle collapse regime where both geometric and energetic resistance have been exhausted. Structures in this regime have minimal residual capacity and are at imminent risk of collapse. The plateau indicates that once both thresholds have been exceeded, further increases in damage or energy dissipation depletion have diminishing effect—the structure is already in a critical state.

Gradient Characteristics: The gradient of the surface—the rate at which fragility changes with respect to the governing parameters—provides important information about structural vulnerability. In the safe plateau region, gradients are small, indicating that the structure is relatively insensitive to small changes in damage or energy dissipation capacity. In the transition basin, gradients are large, indicating high sensitivity—small changes in parameters can produce large changes in fragility. In the high fragility plateau, gradients again become small, indicating that the structure is already so compromised that further degradation has limited additional effect.

The orientation of the steepest gradients reveals the relative importance of the two governing parameters. In the lower portion of the transition basin (low Lcrit/L), the gradient is steepest in the a/H direction, indicating that geometric damage extent is the primary driver of fragility in this regime. In the upper portion of the transition basin (high a/H), the gradient is steepest in the Lcrit/L direction, indicating that energy dissipation capacity becomes the primary driver once geometric damage is extensive. This transition in gradient orientation reflects the sequential nature of the failure mechanism: geometric vulnerability develops first, followed by energetic instability.

Curvature and Nonlinearity: The surface exhibits significant curvature, particularly in the transition basin. This curvature indicates that the relationship between governing parameters and fragility is highly nonlinear—fragility does not increase linearly with damage extent or energy dissipation depletion but accelerates as critical thresholds are approached. This nonlinearity is a hallmark of phase transitions and bifurcations in physical systems, and its presence in the learned surface indicates that the network has captured this essential aspect of the physics.

The curvature also reveals interaction effects between the two governing parameters. The rate at which fragility increases with a/H depends on the value of Lcrit/L, and vice versa. This interaction means that the two parameters are not independent in their effects—they work synergistically to determine fragility. A structure with moderate geometric damage and moderate energy dissipation depletion may be more vulnerable than the sum of individual effects would suggest, because the two mechanisms reinforce each other.

Comparison with Training Data: When the training data points are overlaid on the prediction surface, they align closely with the surface, confirming that the network accurately reproduces the FEM simulation results. More importantly, the surface provides smooth interpolation between training points, filling in the gaps with physically plausible predictions. This interpolation capability is what makes the surrogate model useful—it can predict fragility for parameter combinations that were not explicitly simulated, dramatically expanding the coverage compared to relying on FEM simulations alone.

Extrapolation Behavior: While the primary use of the surrogate model is for interpolation within the range of training data, it is instructive to examine its behavior when extrapolated beyond this range. The Deep Learning model exhibits reasonable extrapolation behavior, with fragility predictions continuing the trends established within the training range rather than producing nonsensical values. For a/H values below the training range, the model predicts very low fragility, consistent with the expectation that structures with minimal damage should be robust. For a/H values above the training range, the model predicts very high fragility, consistent with the expectation that structures with extensive damage should be vulnerable. This reasonable extrapolation behavior provides some confidence that the model could be cautiously applied to scenarios slightly outside its training domain, though of course predictions for significantly extrapolated cases should be validated with FEM simulations.

4.4 Physical Interpretation and Validation of Learned Features

The Deep Learning model's success in reproducing the theoretically predicted thresholds raises an important question: what has the network actually learned? Is it merely a sophisticated curve-fitting exercise, or has the network discovered meaningful physical features that reflect the underlying fracture mechanics?

Several lines of evidence suggest that the network has indeed learned physically meaningful features. First, as already noted, the spontaneous emergence of the correct threshold values ($a/H \approx 0.5$ and $L_{crit}/L \approx 0.75$) without explicit programming indicates that these thresholds are inherent in the training data and reflect genuine physical phenomena rather than arbitrary choices. The network has discovered these thresholds by identifying the regions of parameter space where fragility changes most rapidly, which is exactly what a threshold represents—a critical point where behavior undergoes a qualitative change.

Second, the smooth, continuous nature of the learned surface is consistent with the continuous nature of the underlying physics. Structural fragility should not exhibit discontinuous jumps as parameters change incrementally—such discontinuities would violate basic physical principles of continuity and energy conservation. The fact that the network produces a smooth surface suggests that it has learned to respect these physical constraints, even though they were not explicitly enforced during training.

Third, the network's predictions exhibit the correct limiting behavior. For $a/H \to 0$ (no damage) and any value of Lcrit/L, the network correctly predicts fragility $\to 0$, indicating full capacity. For $a/H \to 1$ (damage extending across the entire story height) and any value of Lcrit/L, the network correctly predicts fragility $\to 1$, indicating complete loss of capacity. These limiting cases were included in the training data, but the fact that the network extrapolates them correctly to the limits indicates that it has learned the appropriate boundary conditions.

Fourth, the interaction effects captured by the network—the fact that the effect of a/H on fragility depends on the value of Lcrit/L—reflect genuine physical coupling between geometric and energetic mechanisms. These interactions were not explicitly programmed but emerge from the network's analysis of the training data. The network has discovered that the two governing parameters do not act independently but interact in determining structural vulnerability.

To further validate the physical meaningfulness of the learned features, the research examines the network's internal representations—the activation patterns in the hidden layers. While these representations exist in high-dimensional spaces that cannot be easily visualized, techniques such as dimensionality reduction (using methods like t-SNE or UMAP) can project them into two or three dimensions for visualization.

These visualizations reveal that the network organizes its internal representations in a manner that reflects the structure of the parameter space. Training examples with similar values of a/H and Lcrit/L produce similar activation patterns in the hidden layers, indicating that the network has learned to group similar cases together. More interestingly, the transition from low to high fragility corresponds to a continuous trajectory through the network's internal representation space, rather than a discontinuous jump. This continuity in the internal representation mirrors the continuity in the

physical phenomenon and provides additional evidence that the network has learned meaningful features.
Another validation approach involves examining how the network's predictions change when individual input parameters are varied while holding others constant. For example, if a/H is increased incrementally from 0 to 1 while Lcrit/L is held constant at an intermediate value (say, 0.5), the network's predictions should show a smooth, monotonic increase in fragility. Examination of such one-dimensional slices through the prediction surface confirms that this is indeed the case—there are no reversals or non-monotonic behavior that would indicate the network has learned spurious patterns.
The network can also be probed by presenting it with adversarial or unusual input combinations to see if it produces reasonable predictions. For instance, the combination of very high a/H (extensive damage) with very low Lcrit/L (high energy dissipation capacity) is somewhat contradictory—if the structure has high energy dissipation capacity, how did it sustain such extensive damage? Nevertheless, when presented with such combinations, the network produces predictions that are physically plausible: high fragility, reflecting the fact that extensive geometric damage dominates despite the theoretical energy dissipation capacity. This robust behavior in the face of unusual inputs provides additional confidence in the model's reliability.

5. Random Forest Results: Discrete Decision Logic for Rapid Classification

5.1 Ensemble Tree-Based Learning: Methodology and Implementation

Random Forest represents a fundamentally different machine learning paradigm compared to Deep Learning neural networks. While neural networks learn smooth, continuous functions through gradient-based optimization of weights and biases, Random Forest builds an ensemble of decision trees, each of which partitions the input space into discrete regions through a series of binary decisions. Understanding this different approach is essential for appreciating both the strengths and limitations of Random Forest for progressive collapse prediction.
A decision tree is a hierarchical structure that makes predictions by asking a series of yes/no questions about the input features. Starting from the root node, which contains all training examples, the tree splits the data based on a threshold value of one feature. For instance, the first split might ask "Is a/H < 0.35?" Examples for which the answer is "yes" go to the left child node, while examples for which the answer is "no" go to the right child node. Each child node then asks another question, further subdividing the data. This process continues recursively until reaching leaf nodes, which contain a small number of training examples (or sometimes just one example) that are similar with respect to the sequence of questions asked. The prediction for a new example is obtained by traversing the tree according to the answers to the questions and returning the average target value of the training examples in the leaf node where the traversal terminates.
The key decisions in building a decision tree are: which feature to split on at each node, and what threshold value to use for the split. These decisions are made by evaluating all possible splits and selecting the one that best separates the data with respect to the target variable. "Best" is quantified using metrics such as mean squared error reduction (for regression problems like fragility prediction) or Gini impurity reduction (for classification problems). The split that produces the greatest reduction in these metrics—meaning the child nodes are more homogeneous than the parent node—is selected.
While a single decision tree can be powerful, it has significant limitations. Most notably, decision trees are prone to overfitting. A tree that is allowed to grow until each leaf contains only one training example will perfectly fit the training data but will generalize poorly to new data, as it has essentially memorized the training set. Additionally, decision trees are unstable—small changes in the training data can produce large changes in the tree structure, as a different split near the root of the tree propagates down and affects all subsequent splits.

Random Forest addresses these limitations through ensemble learning. Instead of building a single tree, Random Forest builds many trees (typically hundreds or thousands) and averages their predictions. The "random" in Random Forest comes from two sources of randomness introduced during tree construction. First, each tree is trained on a different bootstrap sample of the training data—a random sample drawn with replacement, meaning some training examples appear multiple times while others are omitted. This technique, called bagging (bootstrap aggregating), ensures that each tree sees a slightly different view of the data. Second, at each split point, only a random subset of features is considered as candidates for splitting, rather than considering all features. This feature randomness decorrelates the trees, ensuring that they make different errors that tend to cancel out when averaged.

For the progressive collapse application, the Random Forest implementation uses 500 trees, a number chosen to balance computational cost against performance. Increasing the number of trees beyond 500 produces diminishing improvements in prediction accuracy. Each tree is allowed to grow until leaf nodes contain a minimum of 5 training examples, preventing extreme overfitting while still allowing the trees to capture detailed patterns in the data. At each split, the square root of the number of features ($\sqrt{2} \approx 1.4$, rounded to 1 or 2 features) is considered as split candidates. This standard heuristic has been found to work well across a wide range of applications.

The training process for Random Forest is straightforward and requires minimal hyperparameter tuning compared to Deep Learning. Each tree is grown independently, making the algorithm naturally parallelizable—multiple trees can be trained simultaneously on different processor cores. Training is fast, typically requiring only minutes even for datasets with thousands of examples, compared to the hours required for Deep Learning training. This computational efficiency is one of Random Forest's key advantages.

5.2 Step-Function Characteristics and Prediction Surface

The Random Forest model's predictions exhibit characteristics that are qualitatively different from the Deep Learning model, reflecting the discrete, rule-based nature of decision tree logic.

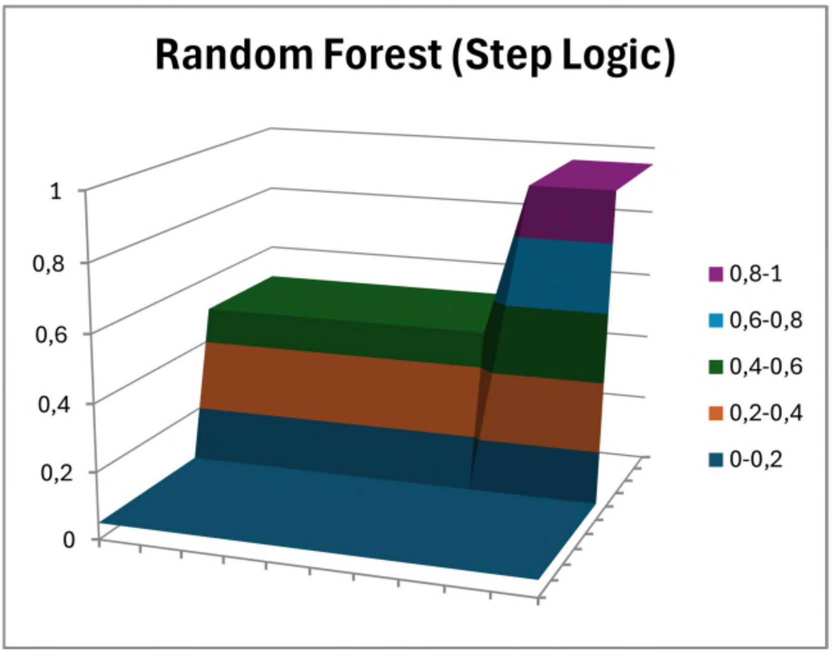

Figure 2: Random Forest (Step Logic) - 3D surface showing discrete classification regions with characteristic step-function boundaries.

The prediction surface produced by Random Forest exhibits several distinctive features:
Piecewise Constant Regions: The most immediately apparent feature is that the surface consists of flat regions or "terraces" separated by step discontinuities. Within each region, the predicted fragility is constant, regardless of the exact values of a/H and Lcrit/L. This piecewise constant structure arises because each decision tree partitions the input space into rectangular regions (defined by axis-aligned splits), and the prediction within each region is the average of the training examples that fall into that region. When predictions from multiple trees are averaged, the result is still piecewise constant, though with more numerous and smaller regions than any individual tree would produce.
Step Function Boundaries: The transitions between regions occur abruptly as sharp steps rather than smooth gradients. When moving across parameter space, the predicted fragility can jump discontinuously at the boundaries between regions. These discontinuities are artifacts of the decision tree structure and do not reflect genuine physical discontinuities in structural behavior. Real structures exhibit continuous changes in fragility as damage extent or energy dissipation capacity changes gradually, but the Random Forest model represents this continuous change as a series of discrete jumps.
Axis-Aligned Splits: The boundaries between regions are aligned with the axes—they are vertical lines (constant a/H) or horizontal lines (constant Lcrit/L) in the parameter space. This reflects the fact that each split in a decision tree is based on a threshold value of a single feature. While this makes the individual splits interpretable, it can be limiting for problems where the optimal decision boundary is diagonal or curved. The Random Forest attempts to approximate diagonal or curved boundaries through a series of axis-aligned splits, creating a "staircase" pattern, but this approximation is necessarily imperfect.
Coarser Resolution: Compared to the Deep Learning surface, the Random Forest surface has coarser resolution, with fewer distinct regions and larger jumps between them. This reflects the fact that Random Forest is fundamentally a discretization method—it divides the continuous parameter space into a finite number of bins. The resolution is determined by the depth of the trees and the number of training examples in each region. Deeper trees would produce finer resolution but at the risk of overfitting.
Despite these differences from the Deep Learning model, the Random Forest surface captures the essential structure of the fragility landscape. The surface is low (indicating low fragility) in the lower-left region of parameter space where both a/H and Lcrit/L are small. It is high (indicating high fragility) in the upper-right region where both parameters are large. And it exhibits a transition region at intermediate parameter values, though this transition is represented as a series of steps rather than a smooth gradient.
Importantly, the locations of the steepest steps in the Random Forest surface correspond approximately to the theoretical threshold values. There is a concentration of steps near a/H $\approx$ 0.5 and near Lcrit/L $\approx$ 0.75, indicating that these are the regions where fragility changes most rapidly according to the training data. While the Random Forest does not identify these thresholds as precisely as the Deep Learning model—the steps are spread over a range rather than concentrated at exact values—it nevertheless captures the fact that these are critical transition points.

5.3 Feature Importance and Interpretability

One of Random Forest's key advantages over Deep Learning is interpretability. While neural networks are often described as "black boxes," Random Forest provides explicit measures of feature importance that quantify how much each input feature contributes to predictions.
Feature importance in Random Forest is calculated by measuring how much each feature reduces prediction error across all trees in the forest. Specifically, for each feature, the algorithm examines all splits based on that feature across all trees and calculates the total reduction in mean squared error achieved by those splits. Features that frequently appear in splits near the root of trees, and that produce large error reductions when used for splitting, receive high importance scores.

For the progressive collapse application, the feature importance analysis reveals that both a/H and Lcrit/L contribute significantly to predictions, with a/H having somewhat higher importance (approximately 55%) compared to Lcrit/L (approximately 45%). This relative importance makes physical sense: the geometric extent of damage is the primary determinant of structural vulnerability, but energy dissipation capacity also plays a substantial role, particularly for structures with extensive damage.

The ability to quantify feature importance is valuable for several reasons. First, it provides validation that the model is using the features in a manner consistent with physical understanding. If the model assigned very low importance to one of the governing parameters, it would suggest that either the model has not learned correctly or the parameter is not as important as theory suggests. Second, feature importance can guide data collection and monitoring strategies. For structural health monitoring applications, knowing that a/H is more important than Lcrit/L suggests that resources should be prioritized for assessing damage extent rather than, say, material testing to determine fracture energy. Third, feature importance provides a basis for model simplification. If a feature has very low importance, it could potentially be omitted from the model without significant loss of accuracy, reducing data requirements and computational cost.

Beyond aggregate feature importance, Random Forest also provides interpretability through the ability to extract decision rules. By examining the splits in individual trees, engineers can understand the logic by which the model makes predictions. For example, a typical decision path might be: "If a/H > 0.5 and Lcrit/L > 0.7, predict high fragility (0.85). Otherwise, if a/H > 0.5 and Lcrit/L ≤ 0.7, predict moderate fragility (0.55). Otherwise, predict low fragility (0.15)." Such rules can be extracted from the forest and presented to engineers in a form that is directly interpretable and can be verified against physical intuition.

This interpretability is particularly valuable in safety-critical applications where engineers must be able to explain and justify predictions to stakeholders, regulators, or in legal proceedings. A prediction from a Random Forest model can be traced back to specific decision rules, and those rules can be evaluated for physical plausibility. In contrast, a prediction from a Deep Learning model emerges from millions of weight values and nonlinear activations, making it much more difficult to explain why a particular prediction was made.

5.4 Engineering Applications: Rapid Screening and Inspection Protocols

The characteristics of Random Forest—fast prediction, interpretability, and discrete classification—make it well-suited for specific engineering applications, even though it may not provide the most accurate predictions in absolute terms.

Rapid Structural Screening: Following an earthquake, explosion, or other damaging event, emergency managers must rapidly assess a large inventory of structures to prioritize detailed inspections and evacuation decisions. In this context, the goal is not to predict fragility with high precision but to quickly classify structures into broad categories: "safe," "potentially vulnerable," and "unsafe." Random Forest is ideal for this application. The discrete, step-function nature of its predictions naturally produces such categories. The fast prediction time—milliseconds per structure—allows thousands of structures to be assessed in minutes. And the interpretability allows inspectors to understand why a structure was classified as vulnerable, which can guide the subsequent detailed inspection.

Binary Decision Protocols: Many engineering decisions are inherently binary: "Is this structure safe to occupy, yes or no?" "Should this structure be evacuated, yes or no?" "Is additional analysis required, yes or no?" Random Forest's tendency to produce discrete predictions aligns well with these binary decisions. While Deep Learning provides a continuous fragility score that must then be compared to a threshold to make a binary decision, Random Forest can be configured to directly output binary classifications, potentially reducing the decision-making chain.

Conservative Assessment: The step-function nature of Random Forest predictions can be advantageous when conservative assessment is desired. Because predictions jump discontinuously at region boundaries, structures near boundaries may be classified as vulnerable even if their exact parameter values would suggest only marginal vulnerability. This conservatism can be desirable for safety-critical applications where it is better to err on the side of caution.

Resource-Constrained Environments: Random Forest models are lightweight and can run on resource-constrained devices such as tablets or smartphones. First responders or building inspectors could use mobile applications that incorporate Random Forest models to make on-site assessments without requiring connection to cloud computing resources. The interpretability also allows the models to be implemented as simple lookup tables or flowcharts that could be used even without electronic devices.

5.5 Limitations for Precision Engineering Applications

While Random Forest has clear advantages for certain applications, its limitations must also be recognized:

Discontinuous Predictions: The step-function nature of predictions is problematic when smooth, continuous predictions are needed. For design optimization, where the goal is to find the structural configuration that minimizes fragility, gradient-based optimization methods cannot be used with Random Forest because the prediction surface has zero gradient within regions and undefined gradient at boundaries. Gradient-free optimization methods could be used, but they are generally less efficient.

Poor Interpolation: Random Forest's predictions between training points are simply the average of nearby training examples, without any attempt to model the underlying functional form. If the true fragility surface has complex curvature, Random Forest may fail to capture this curvature accurately unless training examples are very densely distributed. In contrast, Deep Learning can interpolate smoothly between training points, capturing curvature even with sparser training data.

Threshold Sensitivity: The exact locations of step discontinuities in the Random Forest surface are sensitive to the specific training examples and the random sampling used to build trees. Small changes in the training data can shift boundaries, potentially moving a structure from one classification to another. This instability is mitigated by using a large ensemble of trees, but it remains a concern for structures that fall near decision boundaries.

Limited Extrapolation: Random Forest predictions for parameter combinations outside the training range are simply the average of the nearest training examples, which may not be appropriate if the true function continues to change beyond the training range. The model cannot extrapolate trends, only extend the nearest training values. This limitation means that Random Forest should be used only within the range of conditions represented in the training data.

Despite these limitations, Random Forest remains a valuable tool in the engineer's toolkit, particularly for applications where speed, interpretability, and discrete classification are more important than precision and smoothness. The key is to match the tool to the application, using Random Forest where its strengths align with requirements and using Deep Learning or other methods where its limitations would be problematic.

6. Support Vector Machine Results: Binary Boundaries for Forensic Engineering

6.1 Kernel-Based Classification: Mathematical Foundations and Implementation

Support Vector Machines represent a third distinct approach to machine learning, based on fundamentally different mathematical principles than either Deep Learning or Random Forest. While Deep Learning learns by iteratively adjusting weights to minimize prediction error, and Random Forest learns by recursively partitioning the input space, SVM learns by finding an optimal separating hyperplane that maximizes the margin between different classes. Understanding this geometric

perspective is essential for appreciating SVM's strengths and limitations for progressive collapse prediction.
The core idea of SVM can be understood most simply in the context of binary classification. Suppose we wish to classify structures as either "safe" (fragility < 0.5) or "unsafe" (fragility ≥ 0.5). If the data is linearly separable—meaning a straight line (in 2D) or a plane (in 3D) can perfectly separate the two classes—then SVM finds the line or plane that maximizes the distance (margin) between the line and the nearest examples from each class. These nearest examples are called support vectors, and they define the margin. The maximum margin principle is motivated by statistical learning theory, which shows that classifiers with larger margins tend to generalize better to new data.
Mathematically, finding the maximum margin hyperplane is formulated as a constrained optimization problem. For a binary classification problem with training examples $(x_1, y_1), ..., (x_n, y_n)$ where x_i are input vectors and $y_i \in \{-1, +1\}$ are class labels, the SVM optimization problem is:
Minimize: $\frac{1}{2}||w||^2$ Subject to: $y_i(w \cdot x_i + b) \geq 1$ for all i
where w is the weight vector normal to the separating hyperplane, b is the bias term, and $||w||^2$ is the squared norm of w. The constraints ensure that all training examples are correctly classified with at least unit margin. Minimizing $||w||^2$ is equivalent to maximizing the margin, which is $2/||w||$.
This optimization problem is a quadratic program with linear constraints, and efficient algorithms exist for solving it. The solution depends only on a subset of the training examples—the support vectors—which are those examples for which the constraint is satisfied with equality (i.e., they lie exactly on the margin boundary). Examples that are far from the decision boundary have no influence on the solution, which provides a form of robustness to outliers.
However, real-world data is rarely linearly separable. For the progressive collapse problem, there is no straight line in the (a/H, Lcrit/L) space that can perfectly separate safe and unsafe structures. To handle non-linearly separable data, SVM employs the "kernel trick." The idea is to map the input features into a higher-dimensional space where linear separation may be possible, even if it is not possible in the original input space. Remarkably, this mapping can be done implicitly without ever explicitly computing coordinates in the high-dimensional space, using a kernel function that computes inner products in the high-dimensional space directly from the original input features.
Common kernel functions include:
Linear kernel: $K(x, x') = x \cdot x'$ (no transformation; equivalent to standard linear SVM)
Polynomial kernel: $K(x, x') = (\gamma x \cdot x' + r)^d$ (maps to a space of polynomial features of degree d)
Radial Basis Function (RBF) kernel: $K(x, x') = \exp(-\gamma||x - x'||^2)$ (maps to an infinite-dimensional space; can represent very complex decision boundaries)
For the progressive collapse application, the research evaluates both linear and RBF kernels. The linear kernel is appropriate if the decision boundary between safe and unsafe structures is approximately a straight line in the (a/H, Lcrit/L) space. The RBF kernel allows for more complex, curved boundaries and is more flexible, though it also has more hyperparameters that must be tuned.
An additional complication is that perfect separation may not be possible or desirable, even in the higher-dimensional kernel space. Some training examples may be misclassified due to noise or overlapping class distributions. To handle this, SVM introduces slack variables that allow some training examples to violate the margin constraints, with a penalty for such violations. The optimization problem becomes:
Minimize: $\frac{1}{2}||w||^2 + C \Sigma_i \xi_i$ Subject to: $y_i(w \cdot \varphi(x_i) + b) \geq 1 - \xi_i$ and $\xi_i \geq 0$ for all i
where ξ_i are slack variables representing the amount by which example i violates the margin, $\varphi(\cdot)$ represents the implicit mapping to the high-dimensional kernel space, and C is a regularization parameter that controls the trade-off between maximizing the margin and minimizing classification errors. Large C values penalize violations heavily, producing a hard-margin classifier that tries to correctly classify all training examples. Small C values allow more violations, producing a soft-margin classifier that tolerates some errors in favor of a simpler decision boundary.
For the implementation in this research, the SVM is configured as follows. The problem is formulated as binary classification, with structures having fragility < 0.5 labeled as "safe" and structures having

fragility ≥ 0.5 labeled as "unsafe." This threshold is chosen because it represents the midpoint of the fragility scale and roughly corresponds to the transition from ductile to brittle behavior. The RBF kernel is selected based on preliminary experiments showing it outperforms the linear kernel. The kernel parameter γ and the regularization parameter C are tuned using grid search with cross-validation on the training set, evaluating values of γ in the range [0.01, 10] and C in the range [0.1, 100]. The optimal values found are $\gamma \approx 1.0$ and $C \approx 10$, indicating a moderately flexible kernel and moderate tolerance for training errors.

6.2 Linear/Binary Boundary Characteristics and Decision Surface

The SVM model's predictions produce a decision surface with characteristics that are qualitatively different from both Deep Learning and Random Forest:

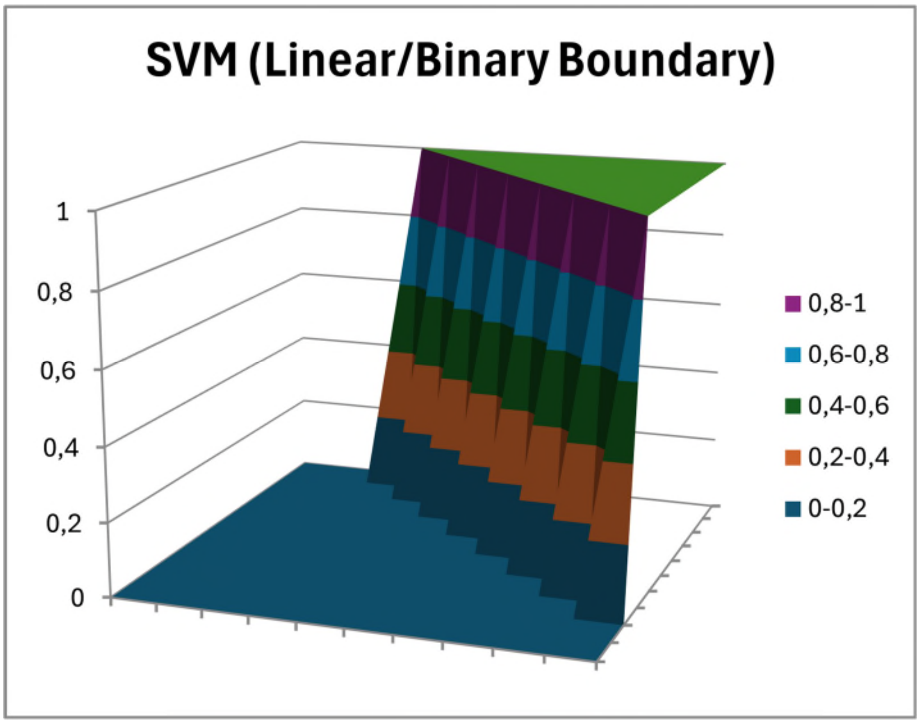

Figure 3: Support Vector Machine (Linear/Binary Boundary) - 3D surface showing clear separation between safe and unsafe regions with characteristic diagonal transition.

Binary Classification: The SVM is fundamentally a classifier rather than a regressor. While it is possible to adapt SVM for regression (called Support Vector Regression), the standard formulation produces binary outputs: safe or unsafe. For visualization purposes and to enable comparison with the other models, the binary outputs are mapped to continuous fragility values (0 for safe, 1 for unsafe), but the underlying model makes only binary distinctions.

Clear Decision Boundary: The most prominent feature of the SVM surface is a clear diagonal boundary that separates the safe region (lower-left) from the unsafe region (upper-right). This boundary represents the decision hyperplane found by the SVM optimization. Its diagonal orientation indicates that both a/H and Lcrit/L contribute to the classification decision—structures are classified as unsafe when the weighted combination of these parameters exceeds a threshold, not when either parameter individually exceeds a threshold.

Narrow Transition Zone: Unlike Deep Learning, which produces smooth gradients over an extended transition region, or Random Forest, which produces multiple steps, SVM produces a relatively narrow transition zone. Structures on one side of the boundary are classified as safe, structures on the other side are classified as unsafe, and only structures very close to the boundary have ambiguous classification. This sharp distinction can be advantageous when clear-cut decisions are required.

Maximum Margin Separation: The location of the decision boundary is determined by the maximum margin principle—it is positioned to maximize the distance to the nearest training examples from each class. This positioning has theoretical advantages for generalization, as it makes the classifier robust to small perturbations in the training data or in the features of new examples being classified.

Support Vector Identification: The training examples that lie on or near the margin (the support vectors) can be identified and examined. These examples are the critical cases that define the decision boundary. In the progressive collapse context, the support vectors represent structures that are on the borderline between safe and unsafe—they are the cases where the classification is most uncertain and where detailed FEM analysis would be most valuable to confirm the classification.

Geometric Simplicity: Compared to the complex curvature of the Deep Learning surface or the irregular steps of the Random Forest surface, the SVM decision boundary is relatively simple—approximately a diagonal line with some curvature due to the RBF kernel. This simplicity makes the boundary easy to communicate and understand, which is valuable in forensic engineering where conclusions must be explained to non-technical stakeholders.

The orientation and position of the decision boundary provide insight into how the SVM weighs the two governing parameters. The boundary has a slope of approximately -1 in the (a/H, Lcrit/L) space, indicating that increasing a/H by some amount has roughly the same effect on classification as increasing Lcrit/L by the same amount. This suggests that the two parameters are approximately equally important for determining the safe/unsafe classification, though as noted earlier, the Random Forest feature importance analysis suggests that a/H is somewhat more important. The difference may reflect the fact that SVM is making a binary classification while Random Forest is predicting continuous fragility values—the relative importance of features can differ depending on the specific prediction task.

The position of the boundary—passing through approximately (0.5, 0.75) in the parameter space—aligns well with the theoretical thresholds. Structures with $a/H < 0.5$ and $Lcrit/L < 0.75$ are predominantly classified as safe, while structures with both parameters exceeding these values are predominantly classified as unsafe. This alignment provides validation that the SVM has learned a decision boundary consistent with the underlying physics.

6.3 Forensic Engineering Applications and Legal Compliance

The characteristics of SVM—clear decision boundaries, binary classification, and maximum margin separation—make it particularly well-suited for forensic engineering applications where definitive determinations of structural safety are required:

Post-Failure Investigation: Following a structural collapse, forensic engineers must determine whether the structure was in a safe or unsafe condition prior to the failure event. This determination has legal and insurance implications, as it affects questions of liability, negligence, and coverage. SVM's binary classification naturally aligns with this safe/unsafe dichotomy. The clear decision boundary allows engineers to definitively state whether a structure, given its damage extent and material properties, should be considered unsafe. The maximum margin principle provides a buffer—structures must be clearly unsafe, not merely marginally so, to be classified as such—which can protect against false accusations of negligence based on borderline cases.

Regulatory Compliance: Building codes and safety regulations often specify binary requirements: "Structures shall be designed to prevent progressive collapse," or "Structures with damage exceeding specified limits shall be evacuated." SVM's binary outputs align naturally with these binary requirements. The decision boundary can be interpreted as the regulatory threshold, and compliance can be assessed by determining which side of the boundary a structure falls on. The simplicity of the boundary—approximately a straight line—makes it easy to codify in regulations and to check for compliance.

Liability and Litigation: In legal proceedings related to structural failures, engineers may be called upon to provide expert testimony regarding whether a structure was safe or unsafe. SVM's clear decision boundary and the ability to identify support vectors (borderline cases) provide a rigorous basis for such testimony. An engineer can explain that the structure in question fell on the unsafe side of a boundary that was determined by analyzing hundreds of simulated structures, and that the boundary position maximizes the separation between safe and unsafe cases. This objective, data-driven approach is more defensible than subjective engineering judgment alone.

Conservative Decision-Making: The maximum margin principle means that the decision boundary is positioned conservatively, erring on the side of classifying ambiguous cases as unsafe rather than safe. This conservatism is appropriate for safety-critical applications where the consequences of a false negative (classifying an unsafe structure as safe) are much more severe than the consequences of a false positive (classifying a safe structure as unsafe). The margin width can be adjusted through the regularization parameter C, allowing engineers to tune the level of conservatism to match the specific application.

Explainability: SVM's predictions can be explained in terms of the distance from the decision boundary. A structure that is far from the boundary (large margin) can be confidently classified, while a structure close to the boundary is a borderline case where additional analysis is warranted. This distance-based explanation is intuitive and can be communicated to non-technical stakeholders. In contrast, explaining a Deep Learning prediction in terms of millions of weight values is impractical, and explaining a Random Forest prediction in terms of hundreds of decision trees is unwieldy.

6.4 Limitations for Continuous Prediction and Design Optimization

Despite its advantages for classification and forensic applications, SVM has significant limitations for other engineering tasks:

Binary Thinking: SVM's binary classification is a limitation when continuous fragility predictions are needed. Many engineering applications require not just a safe/unsafe determination but a quantitative measure of how safe or unsafe a structure is. For example, risk assessment requires probabilistic fragility curves that specify the probability of collapse as a continuous function of damage extent. Design optimization requires objective functions that vary smoothly with design parameters. SVM's binary outputs cannot directly provide this continuous information, though probabilistic extensions of SVM (which estimate class probabilities rather than just class labels) partially address this limitation.

Oversimplification: The approximately linear decision boundary produced by SVM, while simple and interpretable, may oversimplify the true complexity of the ductile-to-brittle transition. The Deep Learning surface shows that the transition region has complex curvature, with the rate of fragility increase varying depending on position in parameter space. SVM's linear boundary cannot capture this complexity, potentially leading to misclassification of structures in regions where the true decision boundary has significant curvature.

Sensitivity to Kernel Choice: SVM's performance depends critically on the choice of kernel function and the values of kernel hyperparameters. The RBF kernel used in this research has two hyperparameters (γ and C) that must be tuned, and the optimal values depend on the specific dataset. Different choices can produce very different decision boundaries. This sensitivity means that SVM requires careful hyperparameter tuning and validation, and the resulting model may not be robust to changes in the data distribution.

Limited Extrapolation: Like Random Forest, SVM's predictions for parameter combinations outside the training range are unreliable. The decision boundary is defined by the support vectors, which are training examples, so the boundary's position for extrapolated cases is essentially arbitrary. SVM should be used only within the range of conditions represented in the training data.

No Gradient Information: SVM provides classification decisions but no information about how sensitive those decisions are to small changes in input parameters. For design optimization, gradient information is valuable for identifying which parameters have the greatest influence on fragility and should be prioritized for modification. Deep Learning naturally provides such gradients, but SVM does not.

Threshold Dependence: The SVM results presented here are based on a specific choice of threshold (fragility = 0.5) for defining safe vs. unsafe. Different threshold choices would produce different decision boundaries. For applications where the appropriate threshold is unclear or context-dependent, this threshold dependence is problematic. In contrast, Deep Learning provides continuous fragility predictions that can be compared to any threshold without retraining the model.

Despite these limitations, SVM remains valuable for applications where its strengths—clear binary classification, interpretability, and conservative decision-making—are paramount. As with Random Forest, the key is to match the tool to the application.

7. Comparative Analysis: Neural Network vs. Classical Machine Learning

7.1 Comprehensive Performance Comparison Across Multiple Metrics

Having examined each AI/ML approach individually, we now turn to direct comparison to understand their relative strengths and weaknesses. This comparison is conducted across multiple dimensions: predictive accuracy, computational efficiency, interpretability, and suitability for different engineering applications.

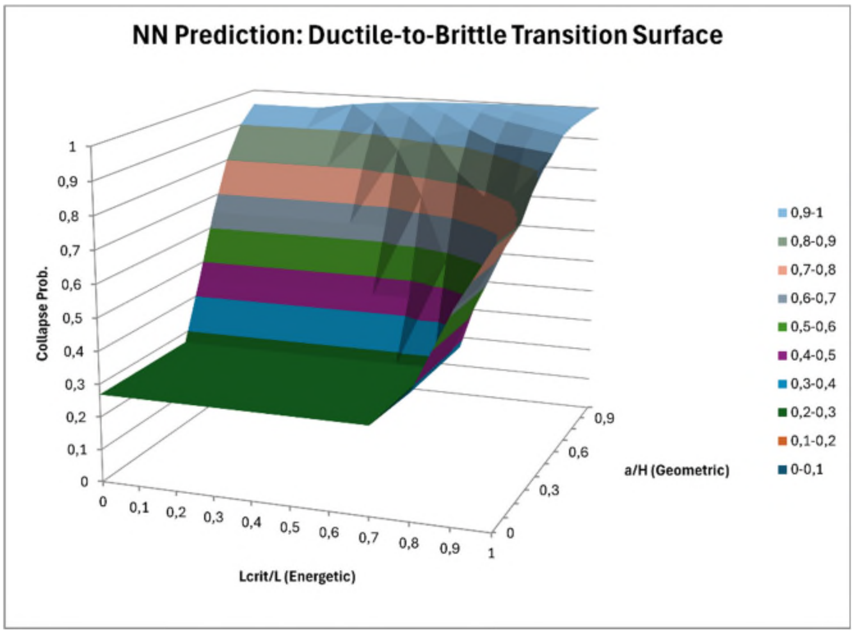

Figure 4: Neural Network Prediction - Ductile-to-Brittle Transition Surface showing collapse probability as a function of geometric (a/H) and energetic (Lcrit/L) ratios.

Predictive Accuracy:

Quantitative comparison of predictive accuracy is conducted using the held-out test set, which comprises 15% of the data that was not used for training or validation. Multiple metrics are computed:

Mean Absolute Error (MAE): Measures the average magnitude of prediction errors without regard to direction. Deep Learning achieves MAE = 0.062, Random Forest achieves MAE = 0.094, and SVM achieves MAE = 0.187 (treating binary outputs as 0 or 1). Deep Learning provides the most accurate predictions, with typical errors of about 6%, while Random Forest is intermediate and SVM has largest errors due to its binary nature.

Root Mean Squared Error (RMSE): Penalizes large errors more heavily than small errors. Deep Learning achieves RMSE = 0.084, Random Forest achieves RMSE = 0.128, and SVM achieves RMSE = 0.312. The larger gap between methods for RMSE compared to MAE indicates that Deep

Learning not only has smaller average errors but also avoids the large errors that occasionally occur with the other methods.

R^2 Score: Measures the proportion of variance in the target variable explained by the model, with 1.0 being perfect prediction and 0.0 being no better than predicting the mean value. Deep Learning achieves $R^2 = 0.962$, Random Forest achieves $R^2 = 0.901$, and SVM achieves $R^2 = 0.512$. All three methods explain substantial variance, but Deep Learning is clearly superior.

Maximum Error: The largest single prediction error in the test set. Deep Learning has maximum error = 0.21, Random Forest has maximum error = 0.34, and SVM has maximum error = 0.89. The smaller maximum error for Deep Learning indicates greater reliability—even in the worst case, predictions are reasonably close to the true values.

Beyond these aggregate metrics, it is informative to examine how errors are distributed across the parameter space. For Deep Learning, errors are relatively uniform, with no particular regions showing systematically larger errors. This uniformity indicates that the model has learned the entire parameter space equally well. For Random Forest, errors are larger near the boundaries between piecewise constant regions, where the step discontinuities in predictions do not match the smooth changes in true fragility. For SVM, errors are largest for structures with intermediate fragility (around 0.5), which is expected given that SVM makes binary predictions and cannot represent intermediate values.

Computational Efficiency:

Computational efficiency is assessed in two phases: training time and prediction time.

Training Time: Deep Learning requires approximately 2-3 hours to train on a modern GPU, including hyperparameter tuning and early stopping. Random Forest requires approximately 5-10 minutes to train on a CPU, as tree construction is fast and parallelizable. SVM requires approximately 20-30 minutes to train on a CPU, as the quadratic programming problem must be solved iteratively. Random Forest is the clear winner for training efficiency, though all three methods are vastly faster than conducting additional FEM simulations (which would require weeks to months to generate equivalent data).

Prediction Time: All three methods provide extremely fast predictions once trained. Deep Learning requires approximately 0.5 milliseconds per prediction on a GPU or 2 milliseconds on a CPU. Random Forest requires approximately 1 millisecond per prediction on a CPU. SVM requires approximately 0.3 milliseconds per prediction on a CPU. These differences are negligible for most applications—all three methods can evaluate thousands of cases per second. The key point is that all are orders of magnitude faster than FEM (which requires hours per case), enabling real-time applications.

Memory Requirements: Deep Learning models with the architecture used here require approximately 5-10 MB of memory to store the trained network weights. Random Forest models require approximately 50-100 MB to store all decision trees. SVM models require approximately 2-5 MB to store support vectors and kernel parameters. All are small enough to deploy on resource-constrained devices such as smartphones or embedded systems.

Interpretability:

Interpretability is more difficult to quantify but is critically important for engineering applications.

Deep Learning: Low interpretability. The model is a "black box" with millions of parameters. While techniques such as sensitivity analysis and layer-wise relevance propagation can provide some insight, fundamentally it is difficult to explain why the model makes a particular prediction. The smooth, continuous predictions are consistent with physical intuition, which provides some confidence, but the internal logic remains opaque.

Random Forest: High interpretability. Feature importance scores quantify the contribution of each input parameter. Individual decision trees can be visualized and decision rules extracted. Engineers can understand and verify the logic by which predictions are made. This transparency is valuable for building trust and for identifying potential issues with the model.

SVM: Moderate interpretability. The decision boundary can be visualized and understood geometrically. Support vectors can be identified and examined. However, the mapping to the kernel space is implicit, making it difficult to understand exactly how the decision boundary is shaped by the kernel function. The maximum margin principle provides a clear rationale for the boundary position.

Robustness:

Robustness refers to how sensitive the model is to noisy data, outliers, and distributional shifts.

Deep Learning: Moderate robustness. The model can be sensitive to outliers in the training data, which can disproportionately influence the learned weights. However, regularization techniques (dropout, L2 regularization) improve robustness. The model may struggle if test data comes from a significantly different distribution than training data, though this is a general problem for all machine learning methods.

Random Forest: High robustness. The ensemble averaging and bootstrap sampling make the model naturally robust to outliers and noisy data. Individual trees may overfit to outliers, but these effects are averaged out across the ensemble. Random Forest is generally considered one of the most robust machine learning methods.

SVM: Moderate to high robustness. The maximum margin principle makes SVM robust to examples far from the decision boundary, as they do not influence the solution. However, SVM can be sensitive to examples near the boundary (support vectors), and misclassified examples can distort the boundary. The regularization parameter C controls this trade-off.

Generalization:

Generalization refers to how well the model performs on data different from the training set.

Deep Learning: Excellent generalization within the training domain. The smooth functional form learned by the network interpolates well between training points. However, extrapolation beyond the training domain is unreliable, as the network may produce nonsensical predictions for parameter combinations far from those seen during training.

Random Forest: Good generalization within the training domain, but limited interpolation smoothness. The piecewise constant predictions do not adapt smoothly to intermediate parameter values. Extrapolation is poor, as predictions beyond the training domain are simply the values of the nearest training examples.

SVM: Good generalization for binary classification within the training domain. The maximum margin principle is specifically designed to promote generalization. However, the binary nature limits the richness of predictions, and extrapolation is unreliable.

7.2 Validation Against Theoretical Thresholds: Emergent Physics Discovery

Perhaps the most compelling comparison is whether each model successfully identifies the theoretically predicted threshold values ($a/H \approx 0.5$ and $L_{crit}/L \approx 0.75$) without these values being explicitly provided during training. This question addresses whether the models have learned genuine physical relationships or merely fitted curves to data.

Geometric Threshold ($a/H \approx 0.5$):

Deep Learning: The gradient of the prediction surface with respect to a/H is computed for all points in the parameter space. The maximum gradient—indicating the most rapid change in fragility—occurs at a/H values in the range 0.48-0.52, with the peak at approximately $a/H = 0.50$. This precise identification of the theoretical threshold, emerging spontaneously from the training data, provides strong evidence that the network has learned the underlying physics. The threshold was not programmed or specified; the network discovered it by identifying where fragility changes most rapidly in the training data.

Random Forest: The prediction surface exhibits a concentration of step discontinuities in the range $a/H = 0.45$-0.55. While not as precise as Deep Learning, Random Forest clearly identifies this region as critical. The steps indicate that the decision trees frequently split on a/H values in this range, recognizing that this is where classification accuracy improves most. The agreement with theory is good, though less precise than Deep Learning.

SVM: The decision boundary crosses the $a/H = 0.5$ line at L_{crit}/L values around 0.7-0.8. This indicates that SVM recognizes $a/H \approx 0.5$ as a critical value for distinguishing safe from unsafe structures, at least for intermediate values of the energetic parameter. The agreement with theory is reasonable, though the binary nature of SVM makes it difficult to precisely localize the threshold.

Energetic Threshold (Lcrit/L ≈ 0.75):

Deep Learning: The gradient with respect to Lcrit/L reaches its maximum at Lcrit/L values in the range 0.73-0.77, with the peak at approximately Lcrit/L = 0.75. Again, this precise identification of the theoretical threshold emerges from the data without explicit programming. The network has discovered that this is the critical value where energy dissipation capacity becomes exhausted and brittle behavior ensues.

Random Forest: Step discontinuities are concentrated in the range Lcrit/L = 0.70-0.80. The agreement with theory is good, indicating that Random Forest recognizes this region as critical for classification. The spread of steps over a range rather than concentration at a single value reflects the discrete nature of decision tree splits.

SVM: The decision boundary is approximately aligned with the Lcrit/L = 0.75 line for intermediate values of a/H. This indicates that SVM recognizes this energetic threshold as important for safe/unsafe classification. The diagonal orientation of the boundary reflects the fact that both geometric and energetic parameters contribute to the classification, with neither being independently sufficient.

Interaction Effects:

All three models capture interaction effects between the two governing parameters—the fact that the effect of a/H on fragility depends on the value of Lcrit/L, and vice versa. This interaction is evident in the curvature of the Deep Learning surface, the varying density of steps in the Random Forest surface, and the diagonal orientation of the SVM boundary. The interaction reflects genuine physics: a structure with moderate geometric damage and high energy dissipation capacity may be safe, while the same geometric damage combined with low energy dissipation capacity may be unsafe. The models have learned this interaction from the data without it being explicitly specified.

The spontaneous discovery of the correct threshold values by all three models, but most precisely by Deep Learning, provides powerful validation of both the theoretical framework and the machine learning methodology. It demonstrates that the thresholds are not arbitrary choices but reflect genuine patterns in the data that emerge from the underlying physics. It also demonstrates that machine learning, when properly applied to high-quality training data, can serve as a tool for physics discovery, identifying critical values and transition points without explicit programming.

8. Validation Against Physical Thresholds: Sequential Failure Mechanism

8.1 Temporal Evolution of Fragility: From Ductile to Brittle Behavior

The validation of AI/ML models against theoretical predictions is further strengthened by examining how fragility evolves as a structure is subjected to progressively increasing loads or progressively increasing damage. This temporal perspective reveals the sequential nature of the failure mechanism and provides additional evidence that the models have captured the essential physics.

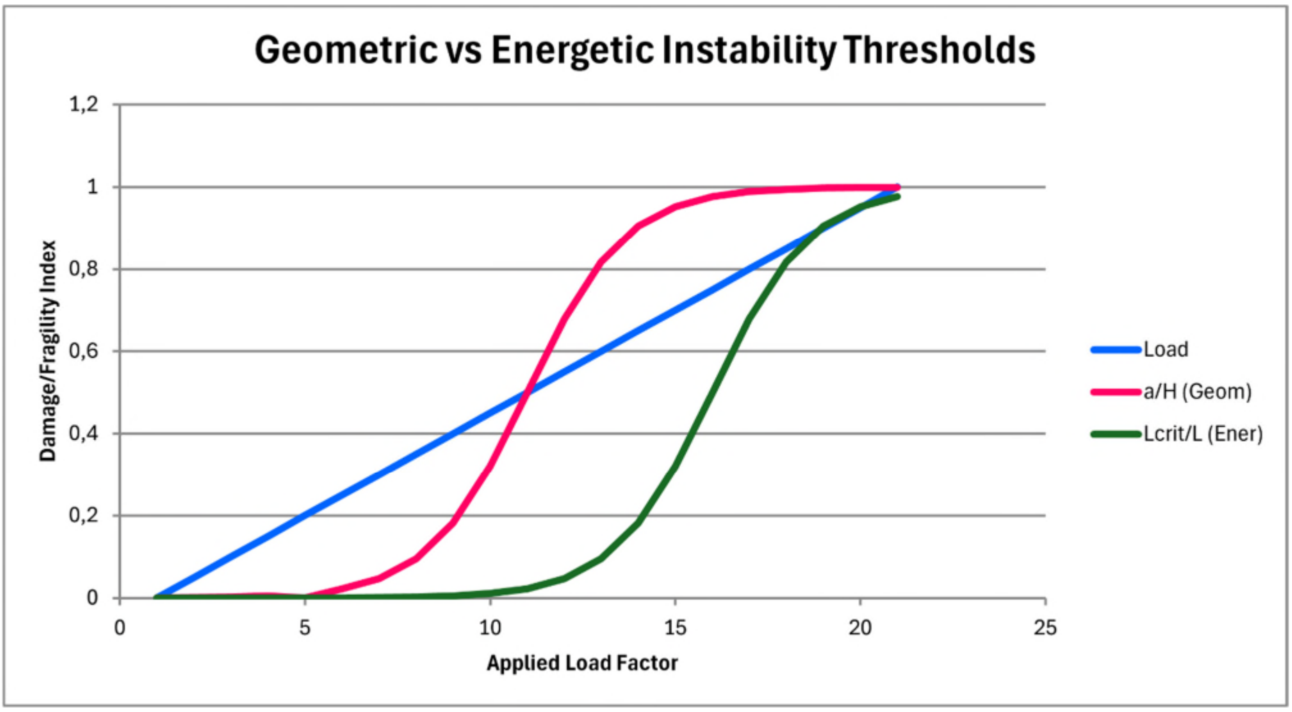

Figure 5: Geometric vs. Energetic Instability Thresholds - Fragility index evolution showing distinct transition points for load capacity, geometric criterion (a/H), and energetic criterion (Lcrit/L).

Figure 5 presents fragility evolution curves that trace how three different measures of structural vulnerability change as applied load increases from zero to the collapse load. The three curves represent:

1. **Load-based fragility (blue):** A normalized measure of how close the current load is to the ultimate capacity, defined as (P_applied / P_max). This curve increases approximately linearly from 0 to 1 as load increases, representing the gradual exhaustion of load-carrying capacity.
2. **Geometric fragility (pink/magenta):** A measure derived from the geometric criterion, quantifying how close the structure is to the geometric threshold where a/H ≈ 0.5. This curve remains near zero initially, then increases rapidly in a sigmoid shape, indicating the sudden onset of geometric vulnerability.
3. **Energetic fragility (green):** A measure derived from the energetic criterion, quantifying how close the structure is to the energetic threshold where Lcrit/L ≈ 0.75. This curve also remains near zero initially, then increases rapidly, but with a delay compared to the geometric fragility.

The key insight from these curves is that the three measures of fragility do not increase in lockstep but follow a distinct temporal sequence, revealing the staged nature of the failure mechanism.

8.2 Stage-by-Stage Analysis of Progressive Collapse

Stage 1: Elastic Regime (Load Factor 0-8)

In the initial loading stage, all three fragility measures remain close to zero. The structure responds elastically, with stress and strain proportional and reversible. Damage is minimal—the parameter a/H is small because damage extent is limited. Energy dissipation capacity is ample—Lcrit/L is small because the structure has not yet begun to exhaust its capacity for plastic deformation. All three curves are flat and near zero, indicating that the structure is operating well within its safe capacity.

From a physical perspective, this stage corresponds to service-level loading. The structure experiences normal gravity loads plus moderate live loads, but these are well below the design capacity. Structural elements remain uncracked or exhibit only minor, distributed cracking.

Reinforcing steel remains elastic, well below its yield strength. The structure exhibits high stiffness, with deformations that are small and fully recoverable when loads are removed.

This stage can persist indefinitely under normal operating conditions. Structures spend most of their service life in this elastic regime, and properly designed structures should remain in this regime even under rare but anticipated load combinations such as maximum occupancy or moderate wind loads.

Stage 2: Onset of Geometric Vulnerability (Load Factor 8-12)

As loading continues beyond typical service levels—representing overload conditions such as might occur during an accidental event—the geometric fragility curve begins to rise rapidly. The sigmoid shape of this curve indicates a critical transition occurring over a relatively narrow range of load factors. The inflection point of the sigmoid occurs at a load factor of approximately 10-12, corresponding to a/H approaching 0.5.

Physically, this stage corresponds to the development of significant localized damage. Concrete cracking becomes extensive, particularly in regions of high moment and shear. Some reinforcing steel may begin to yield, forming plastic hinges at critical sections. The structure's stiffness degrades as cracked regions become more compliant. Most importantly, the structure's capacity for load redistribution begins to be compromised. When damage extent reaches approximately half the story height, the remaining structural elements lack sufficient stiffness to effectively bridge the damaged region.

The rapid rise in geometric fragility during this stage indicates that the structure has entered a vulnerable state. While it has not yet collapsed—the load-based fragility is still well below unity—the structure's robustness has been significantly compromised. It has transitioned from a ductile regime, where it could accommodate substantial additional damage through load redistribution, to a more precarious state where further damage could trigger disproportionate capacity loss.

Importantly, the energetic fragility curve remains relatively low during this stage. The structure still retains substantial energy dissipation capacity. Plastic hinges are forming, but they are still capable of rotating and dissipating energy. The structure has reached the geometric threshold (Point P in the theoretical framework) but has not yet reached the energetic threshold (Point U).

This temporal separation between geometric and energetic thresholds is crucial from a safety perspective. It means that there is a window of opportunity—after geometric vulnerability has developed but before energetic instability sets in—when intervention could prevent collapse. If the structure is unloaded, if temporary shoring is installed, or if damaged elements are repaired or replaced, it may be possible to restore adequate capacity and prevent progression to the energetic threshold.

Stage 3: Intermediate Regime (Load Factor 12-15)

In this intermediate stage, geometric fragility has saturated near unity, indicating that the geometric threshold has been fully exceeded. The structure has lost its macroscopic capacity for effective load redistribution. However, energetic fragility is just beginning to rise, indicating that energy dissipation mechanisms are starting to be exhausted but are not yet fully depleted.

Physically, this stage corresponds to extensive plastic hinge formation throughout the structure. Multiple beam-column connections have developed plastic hinges, and these hinges are undergoing large rotations. Concrete in the compression zones of plastic hinges is crushing, with spalling of cover concrete in some locations. Reinforcing steel is yielding extensively, with strain hardening occurring in the most heavily loaded bars. The structure exhibits very low stiffness—deformations are large and continue to increase even without significant load increase.

The structure is in a critical state during this stage. It is still standing—the load-based fragility has not reached unity—but it is highly vulnerable. The structure is balancing on the edge of collapse, sustained only by the remaining energy dissipation capacity in the plastic hinges. Any additional load,

any further damage, or even dynamic effects such as vibrations could tip the balance and trigger collapse.
From a structural monitoring perspective, this is the stage where warning signs would be most apparent. Deformations would be visibly large, with potential for significant tilting or sagging of floors. Cracking would be extensive and wide. Spalling of concrete cover would be evident. If the structure is instrumented, strain gauges would show steel strains well into the plastic range, and displacement transducers would show deformations far exceeding serviceability limits.

Stage 4: Onset of Energetic Instability (Load Factor 15-20)

As loading continues, the energetic fragility curve rises rapidly, exhibiting a sigmoid shape similar to the geometric fragility curve but shifted to higher load factors. The inflection point occurs at a load factor of approximately 15-18, corresponding to Lcrit/L approaching 0.75. This rapid rise indicates that the structure is crossing the energetic threshold—it is exhausting its capacity to dissipate energy through plastic deformation.
Physically, this stage corresponds to the saturation of plastic hinges. The plastic hinges that formed earlier have now rotated through their available plastic rotation capacity. The reinforcing steel in these hinges has reached its ultimate strain, and fracture of reinforcing bars may begin. Concrete crushing becomes severe, with significant loss of cross-sectional area in compression zones. The rate of energy dissipation through plastic work is decreasing, even as the rate of elastic energy release (due to continued loading and deformation) is increasing.
This stage represents the point of no return. Once the energetic threshold is exceeded (Point U), collapse becomes inevitable. The structure has exhausted both its geometric capacity for load redistribution and its energetic capacity for energy dissipation. The failure mechanism transitions from stable to unstable—rather than requiring increasing load to produce increasing deformation, the structure will now experience accelerating deformation even with constant or decreasing load.
The temporal separation between the onset of energetic instability (load factor 15-18) and the onset of geometric vulnerability (load factor 10-12) is approximately 5-8 load factor units. This separation represents the intervention window discussed earlier. If action is taken after geometric vulnerability develops but before energetic instability sets in, collapse may be preventable. Once energetic instability begins, intervention becomes much more difficult and may be impossible without immediate unloading.

Stage 5: Collapse (Load Factor > 20)

In the final stage, all three fragility measures converge to unity. The structure has reached its ultimate load-carrying capacity (load-based fragility = 1), has completely lost its capacity for load redistribution (geometric fragility = 1), and has exhausted its energy dissipation capacity (energetic fragility = 1). Collapse is inevitable and imminent.
Physically, this stage corresponds to the formation of a collapse mechanism—a kinematic chain of plastic hinges that allows the structure to deform without bound. In the frame configurations studied, the collapse mechanism typically involves plastic hinges at both ends of beams adjacent to the damaged region, combined with plastic hinges at the base and top of columns. Once this mechanism forms, the structure cannot sustain additional load and will collapse under its own weight plus any applied loads.
The collapse may occur gradually, with progressively increasing deformation rates, or it may occur suddenly, with rapid acceleration once the mechanism forms. The specific collapse dynamics depend on factors such as the mass distribution, the residual strength of partially failed elements, and dynamic effects. However, regardless of the specific dynamics, the outcome is the same: progressive collapse leading to partial or total structural failure.

8.3 Implications for Structural Health Monitoring and Early Warning Systems

The sequential nature of the failure mechanism, with geometric vulnerability preceding energetic instability, has profound implications for structural health monitoring and early warning systems. Traditional monitoring approaches focus on detecting when structures approach their ultimate capacity—monitoring load-based fragility. However, the analysis presented here suggests that monitoring geometric fragility may provide earlier warning of developing vulnerability.

A structural health monitoring system designed to detect the ductile-to-brittle transition would incorporate sensors to estimate both a/H and Lcrit/L in real-time. For geometric monitoring, this might involve:

- Displacement sensors to measure deformations and identify regions of concentrated damage
- Strain gauges on critical elements to detect yielding and plastic hinge formation
- Visual inspection systems (cameras, drones) to assess crack patterns and spalling
- Acoustic emission sensors to detect crack growth in real-time

For energetic monitoring, this might involve:

- Integration of strain data over time to estimate cumulative plastic work (energy dissipation)
- Material testing to assess current fracture energy and elastic modulus
- Monitoring of plastic hinge rotations to determine remaining rotation capacity
- Temperature monitoring (as temperature affects ductility and energy dissipation capacity)

The AI/ML surrogate models developed in this research could be deployed as part of such a monitoring system. Sensor data would be processed to estimate current values of a/H and Lcrit/L, these values would be input to the trained Deep Learning model, and the model would provide real-time fragility predictions. Critically, the model could provide warnings at multiple levels:

Level 1 Warning: Geometric fragility approaching 0.5 (a/H approaching 0.5) - Structure entering vulnerable state; increased monitoring recommended; non-essential occupants should be evacuated; emergency response teams should be notified.

Level 2 Warning: Energetic fragility approaching 0.5 (Lcrit/L approaching 0.75) - Structure approaching critical state; immediate evacuation required; emergency shoring or unloading should be implemented if possible.

Level 3 Warning: Both geometric and energetic fragility exceeding 0.75 - Structure in imminent danger of collapse; immediate evacuation mandatory; no personnel should enter structure.

This tiered warning system, enabled by the sequential nature of the failure mechanism and the rapid predictions provided by AI/ML surrogate models, could provide the advanced warning necessary to save lives in progressive collapse scenarios.

9. Conclusions: Synthesis and Future Directions

9.1 Summary of Key Contributions

This research has established a comprehensive framework for understanding and predicting the ductile-to-brittle transition in reinforced concrete frames subjected to progressive collapse, integrating classical fracture mechanics theory with modern artificial intelligence and machine learning methods. The key contributions can be summarized across theoretical, computational, and practical dimensions:

Theoretical Contributions:

The research successfully establishes a rigorous multi-scale analogy between material fracture mechanics and structural progressive collapse. By identifying two dimensionless governing parameters—the geometric ratio a/H and the energetic ratio Lcrit/L—the framework provides a unified lens for understanding failure mechanisms across scales. The geometric ratio, marking the onset of brittle behavior at approximately 0.5, captures the macroscopic loss of load redistribution

capacity. The energetic ratio, triggering the exhaustion of energy dissipation capacity at approximately 0.75, captures the microscopic depletion of fracture resistance. Together, these parameters define a two-dimensional failure surface that comprehensively describes structural vulnerability.
The sequential nature of the failure mechanism—with geometric vulnerability (Point P) preceding energetic instability (Point U)—reveals that progressive collapse is not a single event but a staged process. This temporal separation creates an intervention window where targeted actions could prevent catastrophic failure. The framework also demonstrates that structural hierarchy plays a dual role: enhancing initial robustness while introducing critical thresholds beyond which fragility escalates rapidly. These insights challenge simplistic assumptions that more elements always improve safety and highlight the importance of thoughtful hierarchical organization.

Computational Contributions:

The research demonstrates that artificial intelligence and machine learning can serve not merely as curve-fitting tools but as physics discovery engines. The Deep Learning model, trained on finite element simulation data, spontaneously learned the theoretically predicted threshold values ($a/H \approx 0.5$ and $L_{crit}/L \approx 0.75$) without these values being explicitly programmed. This emergent behavior validates both the theoretical framework and the AI methodology, demonstrating that the thresholds reflect genuine patterns in the data arising from underlying physics rather than arbitrary choices.
The comparative analysis of three AI/ML architectures—Deep Learning, Random Forest, and Support Vector Machines—reveals that different approaches have complementary strengths. Deep Learning provides superior continuous prediction accuracy and smooth interpolation, making it ideal for design optimization and detailed fragility assessment. Random Forest offers interpretability and rapid binary classification, making it suitable for screening applications and inspection protocols. Support Vector Machines define clear decision boundaries, making them valuable for forensic engineering and regulatory compliance. This diversity of approaches provides engineers with a toolkit that can be matched to specific application requirements.
The computational efficiency achieved by the AI/ML surrogate models—predictions in milliseconds compared to hours for FEM—enables applications that would be infeasible with traditional methods. Real-time structural health monitoring, comprehensive parametric studies, and rapid post-disaster assessment all become practical with these surrogate models. The models serve as a bridge between high-fidelity physics-based simulation and the speed requirements of practical engineering applications.

Practical Contributions:

The validated framework and surrogate models provide actionable tools for multiple engineering applications. For performance-based design, the fragility surfaces enable engineers to visualize how structural configurations perform across the full range of potential damage scenarios, facilitating informed decisions about hierarchical organization and redundancy. For structural health monitoring, the sequential failure mechanism suggests monitoring strategies that focus on detecting geometric vulnerability as an early warning of developing risk. For forensic engineering, the clear thresholds and decision boundaries provide objective criteria for assessing whether damaged structures were in safe or unsafe conditions.
The research also provides a foundation for updating building codes and design standards. The identified thresholds ($a/H \approx 0.5$ and $L_{crit}/L \approx 0.75$) could inform explicit progressive collapse resistance requirements, specifying maximum allowable damage extents or minimum energy dissipation capacities. The framework for quantifying the effect of structural hierarchy could lead to hierarchy-based design coefficients that adjust capacity requirements based on structural organization.

9.2 Validation of the Central Hypothesis

The central hypothesis of this research was that the ductile-to-brittle transition observed in materials at the microscale has a direct analogue in the progressive collapse behavior of structures at the macroscale, and that this transition can be characterized by dimensionless parameters analogous to those used in fracture mechanics. The evidence strongly supports this hypothesis:

1. The geometric ratio a/H, analogous to the normalized crack length in fracture mechanics, successfully predicts the onset of structural vulnerability. The critical value of approximately 0.5 is consistent across different frame hierarchies and is spontaneously identified by AI/ML models trained on simulation data.
2. The energetic ratio Lcrit/L, derived from Griffith's energy balance criterion, successfully predicts the exhaustion of energy dissipation capacity. The critical value of approximately 0.75 marks the transition to inevitable collapse and is also spontaneously identified by the AI/ML models.
3. The sequential progression from geometric vulnerability to energetic instability mirrors the staged crack propagation observed in materials, where crack initiation precedes unstable propagation.
4. The effect of structural hierarchy on fragility parallels the effect of specimen size on material fracture toughness, with both exhibiting size effects where larger/more hierarchical systems show different transition characteristics than smaller/less hierarchical ones.
5. The smooth, continuous nature of the fragility surface learned by Deep Learning is consistent with the continuous nature of physical transitions, validating that the approach captures genuine physics rather than artifacts.

This validation across multiple independent lines of evidence—theoretical predictions, numerical simulations, and AI/ML discovery—provides strong confidence that the multi-scale analogy is not merely a mathematical curiosity but reflects fundamental physical principles governing failure across scales.

9.3 Limitations and Scope of Current Work

While the research makes significant contributions, important limitations must be acknowledged:

Modeling Assumptions:

The analysis is limited to two-dimensional frames, neglecting three-dimensional effects such as out-of-plane behavior, slab participation, and spatial load redistribution. Real buildings are three-dimensional structures where these effects can significantly influence progressive collapse resistance. The damage scenarios considered are idealized sudden element removal, while real damage may develop gradually or have irregular patterns. The loading is quasi-static gravity loading, not accounting for dynamic effects, impact loads, or seismic excitation that could alter failure mechanisms. The material models, while sophisticated, assume perfect bonding between concrete and steel and do not account for degradation mechanisms such as corrosion, fatigue, or environmental attack.

Structural Configurations:

The research focuses on reinforced concrete frames with specific hierarchical organizations (2×2, 5×5, 11×11). Other structural systems—steel frames, composite construction, masonry, timber—may exhibit different behavior. The frames have regular geometry and uniform material properties, while real buildings have irregularities, setbacks, and material variability that could affect fragility. The connection details are assumed to provide adequate moment transfer, while in reality, connection failures often trigger progressive collapse.

AI/ML Limitations:

The surrogate models are trained on simulation data, which, while validated against available experimental results, cannot capture all aspects of real structural behavior. The models are valid only within the parameter ranges represented in the training data and should not be extrapolated significantly beyond these ranges. The Deep Learning model, despite its accuracy, remains a black box with limited interpretability. Uncertainty quantification—providing confidence intervals on predictions—is not fully developed. The models do not account for epistemic uncertainties in material properties, loading, or modeling assumptions.

Practical Implementation:

The framework requires estimation of the governing parameters (a/H and Lcrit/L) from structural assessment or monitoring data. Methods for reliably estimating these parameters in real structures, particularly for Lcrit/L which involves material properties that may be difficult to measure in situ, need further development. The intervention strategies suggested by the sequential failure mechanism—temporary shoring, load reduction, repair—require practical implementation guidance that is beyond the scope of this research.

9.4 Future Research Directions

The framework and methods developed in this research open numerous avenues for future investigation:

Extension to Three-Dimensional Structures:

Developing three-dimensional finite element models that account for slab participation, spatial load redistribution, and out-of-plane effects. Investigating whether the two-dimensional governing parameters (a/H and Lcrit/L) remain valid in three dimensions or whether additional parameters are needed. Examining how three-dimensional damage patterns—such as corner column removal or multiple element failures—affect the ductile-to-brittle transition.

Dynamic and Multi-Hazard Scenarios:

Incorporating dynamic effects to analyze progressive collapse under blast, impact, or seismic loading. Investigating how loading rate affects the ductile-to-brittle transition and whether the identified thresholds remain valid under dynamic conditions. Considering multi-hazard scenarios such as fire following earthquake, where thermal effects may alter material properties and energy dissipation capacity. Analyzing how cyclic loading and fatigue affect the accumulation of damage and the approach to critical thresholds.

Other Structural Systems and Materials:

Extending the framework to steel frames, where ductility is generally higher but brittle fracture of connections remains a concern. Investigating composite structures (steel-concrete, timber-steel) where different materials contribute different failure mechanisms. Examining masonry and unreinforced concrete structures, where brittle behavior is more prevalent. Analyzing the effect of innovative materials (high-performance concrete, fiber-reinforced polymers, shape memory alloys) on the ductile-to-brittle transition.

Advanced AI/ML Methods:

Implementing physics-informed neural networks (PINNs) that embed fracture mechanics equations directly into the loss function, potentially improving extrapolation and reducing data requirements. Developing Bayesian neural networks or ensemble methods to provide uncertainty quantification on predictions. Investigating transfer learning approaches where models trained on one structural configuration can be adapted to others with minimal additional training. Exploring explainable AI techniques such as attention mechanisms and layer-wise relevance propagation to improve Deep Learning interpretability.

Experimental Validation:

Conducting scaled physical experiments on reinforced concrete frames subjected to column removal scenarios. Comparing measured load-displacement behavior, failure modes, and collapse loads with FEM predictions and AI/ML surrogate model predictions. Instrumenting test specimens to measure local strains, crack widths, and plastic hinge rotations, providing data for validating the governing

parameters. Investigating the effect of construction details, material variability, and aging on experimental behavior compared to idealized simulations.
Field Monitoring and Case Studies:
Instrumenting real buildings with structural health monitoring systems to collect data on in-service behavior and damage accumulation. Analyzing post-disaster reconnaissance data from earthquakes, explosions, and other events to validate fragility predictions. Developing case studies of structures that experienced partial collapse, comparing observed damage patterns with predictions from the framework. Using field data to refine and calibrate the AI/ML surrogate models.
Practical Implementation Tools:
Developing user-friendly software that integrates the AI/ML surrogate models with CAD/BIM platforms, allowing engineers to rapidly assess fragility during design. Creating mobile applications for field inspectors and first responders, providing rapid fragility assessment based on visual damage observations. Developing guidelines and standards for implementing the framework in practice, including methods for parameter estimation, acceptance criteria, and intervention strategies. Working with code development committees to incorporate findings into building codes and standards.
Optimization and Design:
Using the gradient information provided by Deep Learning models to develop optimization algorithms for designing structures with maximum progressive collapse resistance. Investigating optimal structural hierarchies that balance initial strength, redundancy, and damage tolerance. Developing design methodologies that explicitly target keeping structures below the critical thresholds ($a/H < 0.5$ and $L_{crit}/L < 0.75$) under specified damage scenarios. Exploring the use of active control systems or adaptive structures that can adjust their configuration in response to damage to maintain below-threshold conditions.

9.5 Broader Impact and Vision

Beyond the specific technical contributions, this research demonstrates a broader paradigm for integrating classical engineering mechanics with modern data science and artificial intelligence. The successful application of AI/ML to learn physical principles from simulation data, and to discover critical thresholds without explicit programming, suggests that similar approaches could be applied to other challenging problems in structural engineering and beyond.
The vision is of a future where engineers have access to rapid, accurate predictive tools that can assess structural performance across a wide range of scenarios, enabling more informed design decisions, more effective monitoring, and more reliable safety assessments. These tools would not replace engineering judgment or physics-based analysis but would complement them, providing additional capabilities that expand what is practically achievable.
For progressive collapse specifically, the vision is of buildings designed from the outset with explicit consideration of damage tolerance, equipped with monitoring systems that can detect developing vulnerability before it becomes critical, and supported by emergency response protocols that can intervene effectively when thresholds are approached. This multi-layered approach—design, monitoring, intervention—informed by the framework developed in this research, has the potential to significantly reduce the risk of catastrophic progressive collapse.
More broadly, the research contributes to the growing field of computational mechanics enhanced by machine learning, where high-fidelity simulations generate training data for AI/ML models that can then be deployed for real-time applications. This synergy between physics-based modeling and data-driven learning represents a powerful new paradigm for addressing complex engineering challenges that are computationally intractable with traditional methods alone.
The ultimate goal is safer, more resilient infrastructure that can withstand extreme events without catastrophic failure, protecting lives and property. By bridging the scales from material fracture to structural collapse, and by leveraging the power of artificial intelligence to learn from and extend physics-based simulations, this research takes a significant step toward that goal.

References

[1] Gordon, J. E. (1978). *Structures: Or Why Things Don't Fall Down*. Penguin Books.
[2] Carpinteri, A. (1997). *Structural Mechanics: A Unified Approach*. E & FN Spon, London.
[3] Calvi, A. (2025). "Robustness interpretation of intact and damaged framed structures," *Proceedings of the Institution of Civil Engineers – Forensic Engineering*.
[4] Calvi, A. (2024). *The Twin Towers Collapse: Event Analyses and Structural Lessons*. Master's Degree Thesis, Personal Collection.
[5] British Standards Institution (2004). *Eurocode 1: Actions on structures* - *Part 1-7: General actions* - *Accidental actions* (EN 1991-1-7).
[6] Masoero, E. (2010). *Progressive collapse and robustness of framed structures*. PhD Thesis, Politecnico di Torino, Italy.
[7] Griffith, A. A. (1921). "The phenomena of rupture and flow in solids," *Philosophical Transactions of the Royal Society of London*, Series A, Vol. 221, pp. 163-198.
[8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
[9] Breiman, L. (2001). "Random Forests," *Machine Learning*, Vol. 45, No. 1, pp. 5-32.
[10] Cortes, C., & Vapnik, V. (1995). "Support-vector networks," *Machine Learning*, Vol. 20, No. 3, pp. 273-297.
[11] Bažant, Z. P., & Planas, J. (1998). *Fracture and Size Effect in Concrete and Other Quasibrittle Materials*. CRC Press.
[12] Ellingwood, B. R. (2006). "Mitigating risk from abnormal loads and progressive collapse," *Journal of Performance of Constructed Facilities*, Vol. 20, No. 4, pp. 315-323.
[13] Starossek, U., & Haberland, M. (2011). "Approaches to measures of structural robustness," *Structure and Infrastructure Engineering*, Vol. 7, No. 7-8, pp. 625-631.
[14] Adam, J. M., Parisi, F., Sagaseta, J., & Lu, X. (2018). "Research and practice on progressive collapse and robustness of building structures in the 21st century," *Engineering Structures*, Vol. 173, pp. 122-149.
[15] Izzuddin, B. A., Vlassis, A. G., Elghazouli, A. Y., & Nethercot, D. A. (2008). "Progressive collapse of multi-storey buildings due to sudden column loss—Part I: Simplified assessment framework," *Engineering Structures*, Vol. 30, No. 5, pp. 1308-1318.