

OpenPyStruct: Open-Source Toolkit for Machine Learning-Driven Structural Optimization

Danny Smyl^a, Bozhou Zhuang^a, Sam Rigby^b, Edvard Bruun^a, Brandon Jones^a, Patrick Kastner^c, Iris Tien^a, Adrien Gallet^d

^a*School of Civil and Environmental Engineering, Georgia Institute of Technology, 790 Atlantic Drive, Atlanta, 30332-0355, GA, USA*

^b*School of Mechanical, Aerospace and Civil Engineering, University of Sheffield, S1 3JD, UK*

^c*School of Architecture, Georgia Institute of Technology, 280 Ferst Drive NW, Atlanta, GA, 30332-0355, USA*

^d*Unipart Construction Technologies, Advanced Manufacturing Park, Brunel Way, Sheffield, S60 5WG, UK*

Abstract

OpenPyStruct is a first-version open-source toolkit that provides finite element model based optimization frameworks for generating training data and machine learning models for global structural optimization of indeterminate continuous structures. The key machine learning feature of **OpenPyStruct** is its ability to optimize subject to single or multiple arbitrary simultaneous loading and/or support conditions. The framework utilizes multi-core CPU and GPU-enhanced implementations integrating **OpenSeesPy** forward solvers in structural optimization, leveraging **PyTorch** for accelerated computations. Accompanying machine learning scripts enable users to train high-fidelity predictive models using transformer architectures with diffusion modules, physics-informed neural networks (PINNs), convolutional operations, and contemporary machine learning techniques to analyze and optimize structural designs. By incorporating state-of-the-art optimization tools, robust datasets, and flexible machine learning resources, **OpenPyStruct** aims to establish a scalable – fully transparent – engine for structural optimization by engaging the structural engineering community in this open-source project.

Keywords: Structural Design, Finite Element Method, Machine Learning, Optimization, Python, Structural Engineering

Contents

1	Introduction	3
2	The OpenPyStruct Framework	4
2.1	Overview	4
2.2	OpenPyStruct Scripts and Data Overview	4
3	OpenPyStruct for Classical Model Based Optimization	7
3.1	Mathematical Model Based Optimization Approach	7
3.2	Optimization Implementation and Loss Computation	7
3.3	Optimization Workflow	8
3.4	GPU-Accelerated Optimizer	9
3.5	Single-Core Optimizer	9
3.6	Multi-Core Optimizer	9
3.7	Comparison of Program Variants	10
3.8	Visualization Features	10
3.9	Flexibility in Design Parameters	10
3.10	Flexibility in Loss Functions	10
3.11	Data Generation Workflow Summary	11
4	OpenPyStruct for Machine Learning Optimization	13
4.1	Loss Function Approach	13
4.2	Data Preprocessing and Label Aggregation Approach	13
4.3	Learning Dynamics Approach	15
4.4	Evaluation and Inference Approach	15
4.5	Feedforward Neural Network Optimizer	16
4.6	Physics-Informed Neural Network Optimizer	16
4.7	Transformer-Diffusion Optimizer	17
5	Examples and Discussion	18
5.1	Case Study 1: PINN and FNN Simultaneous Optimization for 6 Load Cases with fixed roller locations	18
5.2	Case Study 2: Transformer-Diffusion Simultaneous Optimization for Arbitrary Load Cases, Spans and Support Conditions	20
5.3	Computational Considerations	23
5.4	Discussion	23
5.4.1	Potential for Extended Capabilities: Simultaneous Load and Support Optimization	23
5.4.2	Potential for Extended Capabilities: Integration of Novel Standardization	23
5.4.3	Potential for Extended Capabilities: Integration of Novel Losses	24
5.4.4	Potential for Extended Capabilities: Simplification Using Standard Shapes and Inclusion of Environmental Effects	24
6	Conclusion and Future Work	25
6.1	Summary of Key Contributions	25
6.2	Future Directions for OpenPyStruct and Machine Learning Structural Optimization	25
7	Data Availability Statement	26
8	Statements and Declarations	26
9	CRediT Author Statement	26
10	Appendix A: Summary of Variables and Functions	26
11	Appendix B: Classical Optimization Pseudocode	27
12	Appendix C: ML Optimization Pseudocode (Transformer-Diffusion Example)	29

1. Introduction

Structural design can be fundamentally viewed as an inverse problem [1, 2, 3], where the goal is to determine the optimal structural configurations and material properties that satisfy desired performance subject to specified loading and boundary conditions. Unlike forward problems, which involve computing structural responses under known parameters, inverse problems require reasoning backwards from performance targets to identify solutions [4]. For structural design, this process is inherently nonlinear, multi-constrained, and computationally demanding due to the complexity of physical laws, such as equilibrium, compatibility, and material behavior.

Moreover, inverse structural design is ill-posed, meaning that small variations in the desired performance criteria or input conditions can lead to significant changes in the solution, or the solution may not exist or have multiple valid solutions. Ill-posedness arises from the underdetermined and indeterminate nature of many structural design problems, where the available information about the desired performance is insufficient to fully specify the structural configuration [5]. This lack of robustness and determinacy further complicates the optimization process, as the solutions must be regularized or constrained to ensure physical feasibility and stability. Addressing this ill-posedness thus often requires advanced computational methods that can incorporate prior knowledge, enforce physical constraints, and handle uncertainty in the input data, all while maintaining computational efficiency [1].

Traditional approaches to structural optimization, including functionalities such as gradient-based methods and finite element simulations, have provided physically rigorous solutions to these problems [6, 7]. However, their applicability may be limited by high computational costs [8] when considering multiple simultaneous loading and boundary conditions, exploring high-dimensional design spaces or addressing scenarios with significant variability in input conditions. While this is the source of ongoing research [9], many such methods may struggle to incorporate and benefit from the increasing availability of large-scale datasets that capture diverse structural behaviors across a wide range of configurations and conditions.

Machine learning (ML) provides a powerful alternative solution for addressing these limitations by using data-driven techniques to model complex relationships between structural inputs and performance outcomes [10, 11]. ML models, such as neural networks and ensemble methods, excel at capturing nonlinear patterns in high-dimensional datasets and generalizing across diverse conditions. Examples include publications on ML applications of topology optimization [12, 13, 14, 15, 16]). With access to large-scale datasets, ML enables rapid predictions, and thus reduces the computational burden of exploring design alternatives. In structural design, ML has been proven to be capable of making mappings between load configurations, boundary conditions, and optimal design parameters at a fast and accurate manner [1].

Incorporating physical laws into ML further enhances its applicability in structural engineering. Physics-informed neural networks (PINNs) embed governing equations from structural design into the learning process [17, 18]. Therefore, relevant mechanical constraints such as equilibrium and compatibility can be satisfied. The integration of physics with ML mitigates the "black-box" nature of traditional data-driven models and improves the reliability and interpretability in engineering applications.

In our approach, we propose **OpenPyStruct**, an open-source platform that integrates single- and multi-core physics-based optimization with ML workflows. By utilizing **OpenSeesPy** [19] with **PyTorch** [20] for accelerated computation on central processing units (CPUs) and graphics processing units (GPUs), **OpenPyStruct** efficiently generates high-fidelity datasets with diverse structural responses across a wide range of loading conditions, support configurations, and structural geometries. Users can employ ML scripts to train predictive models for single and multiple load and support structural optimization with or without physics constraints. This combination of rigorous physics-based simulations, large-scale data generation, and ML tools offers engineers and researchers a scalable, flexible, and high-performance platform for solving complex structural optimization problems. The contributions of this paper are summarized as follows:

- Develop an open-source toolkit that integrates physics-based structural simulation with ML for structural design and optimization;
- Establish a scalable data-generation pipeline that efficiently produces high-fidelity datasets with diverse loading and support conditions;
- Release datasets for use in ML-driven optimization;
- Demonstrate the efficacy of the toolkit.

2. The OpenPyStruct Framework

2.1. Overview

OpenPyStruct is designed to provide a modular framework for continuous structural optimization. **OpenPyStruct** contains the following core functionalities: (1) single load case simulation-based optimization, (2) data generation with a single-core/single-GPU and a multi-core operations and (3) ML-driven design employing feedforward neural networks (FNNs), PINNs, and transformer models with diffusion processes. In addition, **OpenPyStruct** includes advanced tools for data visualization, statistical analysis, and enhancing the machine learning process to make it a versatile platform for tackling different structural design problems. All optimization tasks, from design variable refinement to neural network training, are handled using computational graphs in **PyTorch**.

As shown in Fig. 1a, the single load simulation-based optimization module uses **OpenSeesPy** as the forward model in optimizing design variables such as the second moments of area. While optimizing the second moment of area is the default, it should be mentioned that this module can be modified to optimize multiple design targets. This component is tailored for scenarios where the objective is to optimize a single structural configuration under specified loading and boundary conditions. By directly solving the governing mechanics, the module ensures that results adhere to physical laws for both standalone optimization tasks and dataset generation.

Establishing a scalable structural data-generation pipeline is a core objective of the **OpenPyStruct** toolkit. To achieve this, three complementary modules are provided, a single-core generator, a GPU generator and a multi-core generator, as illustrated in Fig. 1b. The single-core and GPU generator execute data generation sequentially. However, sequential generation may be slow for generating large datasets. To address this limitation, the multi-core generator employs parallelized data generation.

The ML scripts in **OpenPyStruct** provide advanced tools for structural optimization, as demonstrated in Fig. 1c. In the ML suite, FNNs offer a straightforward and computationally efficient approach for predictive modeling. Meanwhile, PINNs incorporate governing physical information into learning – ensuring that predictions are physically consistent. Transformer models with diffusion modules are also provided, having capabilities to handle highly nonlinear relationships and noisy inputs. All scripts include preprocessing utilities for data standardization, trend visualization, and unscaling operations to ensure that results are interpretable and consistent with physical units, as shown in Fig. 1d.

OpenPyStruct also integrates advanced ML optimization protocols to enhance the performance and flexibility of its models. Learning rate scheduling enables adaptive training by dynamically adjusting optimization rates to balance convergence speed and stability. Regularization techniques (e.g., weight decay, noise scheduling, dropout, etc.) aim to mitigate overfitting and improve the generalization of models on unseen data. Additionally, the architecture of the frameworks allow for flexible customization of neural network configurations, allowing users to tailor models to specific problem domains. These features, combined with robust visualization and statistical tools, position **OpenPyStruct** as a flexible framework that bridges physics-based optimization with modern data-driven methodologies for structural design and analysis.

2.2. OpenPyStruct Scripts and Data Overview

Several scripts are provided in **OpenPyStruct** as summarized in Table 1. Beginning first with the core ML optimization modules, the **OpenPyStruct_FNN_MultiCase.py** script employs FNNs with residual connections for scalability and robust performance across diverse structural configurations. Its architecture is optimized for efficient handling of multi-load cases while incorporating advanced regularization to enhance generalization capabilities. Aiming to improve performance beyond the FNN, the **OpenPyStruct_PINN_MultiCase.py** script integrates PINNs to simultaneously optimize design parameters, beam deflections and rotations. By combining data-driven learning with physics-based loss functions, this module efficiently captures complex structural behaviors and boundary conditions. The most advanced core ML module, the **OpenPyStruct_TransformerDiffusionModule_MultiCase.py** script, introduces cutting-edge transformer-based encoders and diffusion models. With multi-head attention mechanisms and diffusion-based latent space exploration, this script excels in predicting deflection, rotation, and stress distributions for complex structural systems. Its advanced feature extraction enables improved precision during optimization under intricate load scenarios.

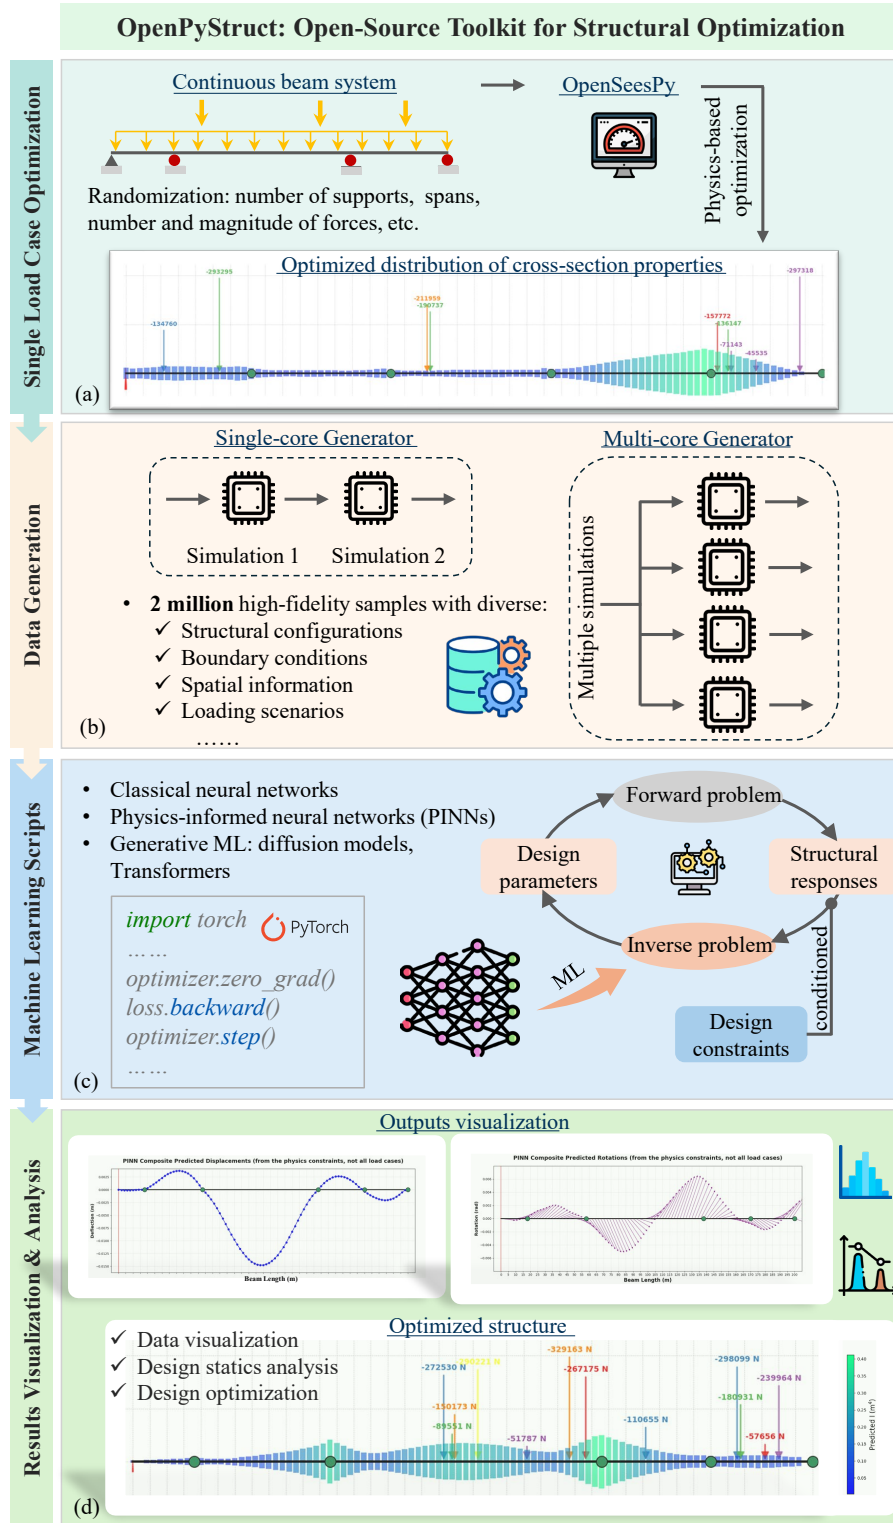


Figure 1: OpenPyStruct workflow summarizing key functionalities, dependencies and capabilities. (a) Single load case optimization, (b) Data generation, (c) Machine learning scripts, and (d) Results visualization & analysis.

In addition to the core ML modules, cutting-edge experimental models (labeled 'Beta') are also provided in the toolkit for user exploration. These models include Graph Neural Network (GNN), Bayesian Transformer Diffusion, trainable output weighting (meta parameters), Fourier Neural Operation (FNO) and evolutionary approaches. In addition, a discrete (i.e., constant design variables per member) *frame* optimizer is provided. These models are fully functional and offer a nucleation point for future toolkit community development.

Supporting these optimization modules are computational backends including modules `OpenPyStruct_BeamOpt_training_GPU.py` and `OpenPyStruct_BeamOpt_training_MultiCore.py`, which leverage GPU acceleration and parallel processing, respectively. These backends enable rapid and large-scale data generation, which is crucial for creating extensive datasets required for ML training and evaluation.

Two primary datasets, **StructDataLite** and **StructDataHeavy**, were generated using the aforementioned backends following the description in Section 3.11. The data are described as follows:

- **StructDataLite**: This dataset contains 100,000 samples and serves as a compact resource for initial training, debugging, and exploratory analysis. Support conditions are fixed in this dataset.
- **StructDataHeavy**: A comprehensive dataset comprising 2,000,000 samples for high-fidelity model training. This dataset provides extensive variability across structural configurations and load cases.

Table 1: Summary of `OpenPyStruct` Core Modules and experimental ('Beta') modules.

Core Module/Script Shorthand	Functionality	Key Features
<code>PINN_MultiCase</code>	PINN-based optimization for structural responses	Combines data-driven and physics-informed methods
<code>FNN_MultiCase</code>	Scalable FNN for multi-load case problems	Residual connections and regularization techniques
<code>TransformerDiffusionModule_MultiCase</code>	Transformer and diffusion-based optimization	Multi-head attention, latent space exploration
<code>BeamOpt_training_GPU</code>	GPU-accelerated data generation	Rapid, large-scale dataset creation
<code>BeamOpt_training_MultiCore</code>	Parallelized data generation on multiple cores	Efficient handling of large computational loads
<code>BeamOpt_training_SingleCore</code>	Single-core data generation	Optimized for environments without parallel processing
<code>BeamOpt</code>	Single-load optimization tool	Focused on individual structural load cases
Beta Module/Script Shorthand		
<code>Bayesian_TFModule_MultiCase_Beta</code>	Bayesian Transformer and diffusion-based optimization	Bayesian representation of MLP, multi-head attention, diffusion module
<code>Bayesian_TFModule_MultiCase_Meta_Beta</code>	Bayesian Transformer and diffusion-based optimization with output weighting	Bayesian representation of MLP, multi-head attention, diffusion module, trainable output weights, output statistical analyzer
<code>GNN_MultiCase_Beta</code>	Graph NN representation of the structural system	Graph connectivity, adjacency, chain attention
<code>FNO_MultiCase_Beta</code>	FNO-based optimization for multi-load cases	Fourier Neural Operators, spectral convolutions, customizable layer stacking
<code>BeamOpt_Evolutionary_Beta</code>	Evolutionary algorithm based single-load optimization tool	Population-based search, tournament selection, uniform crossover, mutation, and elitism
<code>FrameOpt_Discrete_Beta</code>	Single-load discrete <i>frame</i> optimization tool	Focused on individual load cases for arbitrary bays and stories

3. OpenPyStruct for Classical Model Based Optimization

3.1. Mathematical Model Based Optimization Approach

In **OpenPyStruct**, the classical structural optimization approach is formulated to refine design parameters, such as I , to achieve an optimal balance between performance, material distribution, and structural feasibility. The framework employs a composite loss function \mathcal{L} that integrates multiple components representing different structural behaviors. The optimization process is driven by the minimization of this loss function to make sure physical constraints such as equilibrium and compatibility are satisfied. This formulation is flexible to allow users to adjust weights for bending and shear penalties, and to extend or replace the design parameters I with others, such as cross-sectional areas A or user-specified material properties.

The structural optimization problem in **OpenPyStruct** is defined as:

$$\min_I \mathcal{L}(I) = \mathcal{L}_{\text{primary}}(I) + \alpha_{\text{moment}} \mathcal{L}_{\text{bending}}(M, E, I) + \alpha_{\text{shear}} \mathcal{L}_{\text{shear}}(G, k, V, I) \quad (1)$$

where the primary loss term, $\mathcal{L}_{\text{primary}}(I)$, is defined as the sum of the second moments of area I_e for each beam element e :

$$\mathcal{L}_{\text{primary}}(I) = \sum_e I_e. \quad (2)$$

This term minimizes the design parameters I_e , aiming to indirectly optimize material quantity for each element. The bending loss term, $\mathcal{L}_{\text{bending}}(M, E, I)$, penalizes elastic bending energy via:

$$\mathcal{L}_{\text{bending}}(M, E, I) = \sum_e \frac{M_e^2}{2EI_e}. \quad (3)$$

Similarly, the shear energy loss term, $\mathcal{L}_{\text{shear}}(G, k, V, I)$, penalizes shear energy via:

$$\mathcal{L}_{\text{shear}}(G, k, V, I) = \sum_e \frac{V_e^2}{GA_e}. \quad (4)$$

Here, $A_e \approx kI_e^{\frac{1}{2}}$ is the approximate cross-sectional area of the e -th element, which is approximated through a user-defined transformation coefficient k^1 , and is tied to the second moment of area I_e for optimization. Moreover, E is the Young's modulus, and G is the shear modulus. Meanwhile, the parameters α_{moment} and α_{shear} are user-defined coefficients that control the relative importance of the bending and shear penalties in the total loss function. These coefficients can also be treated as learnable parameters during the optimization process to allow for more flexible and adaptive control over the loss components.

3.2. Optimization Implementation and Loss Computation

The classical optimization process in **OpenPyStruct** is implemented iteratively using **PyTorch**'s automatic differentiation to compute gradients of the loss function $\mathcal{L}(I)$. The integration of **PyTorch** ensures that gradients with respect to the design parameters I are computed efficiently. Below is a detailed description of the algorithm for setting up the structural model and computing the loss components.

Structural Model Setup: The structural model is configured using **OpenSeesPy**, where nodes, supports, elements, and loads are defined based on the input design parameters and configurations. The code snippet for the setup is as follows:

```
# Define the OpenSeesPy model
def setup_model(I_tensor, node_positions, roller_nodes, ...):
    ops.model('basic', '-ndm', 2, '-ndf', 3)

    # Define nodes based on their spatial positions
    for i, x in enumerate(node_positions):
        ops.node(i + 1, x, 0.0) # Node IDs start at 1 in OpenSeesPy
```

¹Alternatively, the user may define A and I explicitly or omit the shear energy loss as they see fit.

```

# Apply supports
ops.fix(1, 1, 1, 0) # Fixed (pin) support at the first node
ops.fix(len(node_positions), 0, 1, 1) # Roller support at the last node
for node in roller_nodes:
    ops.fix(node, 0, 1, 1) # Roller supports at specified locations

# Define elements with second moments of area from I_tensor
ops.geomTransf('Linear', 1)
for i in range(len(node_positions) - 1):
    ops.element('elasticBeamColumn', i + 1, i + 1, i + 2, A, E, I_tensor[i].item(), 1)

# Apply point loads
ops.timeSeries('Linear', 1)
ops.pattern('Plain', 1, 1)
for node, force in zip(force_nodes, force_values):
    ops.load(node, 0.0, force, 0.0)

# Apply uniform distributed load (UDL)
for elem_id in range(1, len(node_positions)):
    ops.eleLoad('-ele', elem_id, '-type', '-beamUniform', uniform_udl, uniform_udl)

```

Loss Function Computation. The loss function comprises multiple components, including the primary design parameter, bending energy, and shear energy, each weighted to balance their contributions. The bending energy penalizes large bending moments, while the shear energy accounts for the impact of shear forces. The primary loss minimizes the total second moments of area. The following snippet computes these components:

```

# Define the loss components
def compute_combined_loss(I_tensor, E, G, alpha_moment, alpha_shear):
    bending_moments = []
    shear_forces = []

    # Retrieve element responses
    for elem_id in range(1, len(I_tensor) + 1):
        response = ops.eleResponse(elem_id, 'forces')
        shear_forces.append(response[1]) # Shear force
        bending_moments.append(response[2]) # Moment at the start of the element

    # Convert to PyTorch tensors
    bending_moments = torch.tensor(bending_moments, dtype=torch.float32, requires_grad=True)
    shear_forces = torch.tensor(shear_forces, dtype=torch.float32, requires_grad=True)

    # Compute bending energy
    bending_energy = torch.sum((bending_moments**2) / (2 * E * I_tensor + 1e-6))

    # Compute shear energy
    A = k * I_tensor**0.5 # Approximate area proportional to I
    shear_energy = torch.sum(shear_forces**2 / (G * A))

    # Compute primary loss
    primary_loss = torch.sum(I_tensor)

    # Combine all losses
    total_loss = primary_loss + alpha_moment * bending_energy + alpha_shear * shear_energy

    return total_loss

```

This workflow ensures that gradients of the combined loss function with respect to the second moments of area I are efficiently computed.

3.3. Optimization Workflow

The general workflow for the optimization workflow is as follows (detailed description provided in Appendix B):

1. **Initialization:** Design parameters I_e are initialized with a constant guess I_0 :

$$I_e^{(0)} = I_0, \quad \forall e.$$

2. **Simulation Setup:** Structural properties, node positions, and boundary conditions are defined in OpenSeesPy.
3. **Gradient-Based Update:** Gradients of $\mathcal{L}(I)$ are computed using PyTorch’s computational graph, and design parameters are updated.
4. **Stopping Criteria:** Optimization continues until the following is satisfied at iterate s with patience p :

$$|\mathcal{L}(I^{(s)}) - \mathcal{L}(I^{(s-1)})| < \text{tolerance}.$$

5. **Final Solution:** The optimized design parameters I_e^{opt} are obtained and can be further analyzed/visualized.

3.4. GPU-Accelerated Optimizer

The GPU-Accelerated Optimizer utilizes GPUs to perform data generation and optimization tasks efficiently. It uses PyTorch for tensor operations on GPUs. The program is designed for environments where a local GPU is available and maximizes parallelism in tensor computations with the following key functionalities:

- **GPU Utilization:** The program detects and utilizes available GPUs for moment of inertia tensor operations to improve performance for large datasets.
- **Progress Monitoring:** Includes a visual progress bar (`tqdm`) to track the generation of training samples in real-time.

3.5. Single-Core Optimizer

The Single-Core CPU Optimizer is designed for systems where parallel processing capabilities or GPUs are not available. It performs optimization and training data generation sequentially. This version is ideal for debugging or small-scale experimentation with the following key features:

- **Sequential Execution:** Processes training samples one at a time to reduce complexity and avoiding issues with resource contention.
- **Flexible Input Parameters:** Allows customization of key parameters, such as the number of nodes, roller supports, and applied forces.
- **Optimization Loop:** Employs a similar optimization framework as the GPU version with the Adam optimizer and a learning rate scheduler for loss minimization.
- **Error Handling:** Provides robust handling of failures during the analysis phase to ensure smooth execution even in edge cases.

3.6. Multi-Core Optimizer

The Multi-Core Optimizer utilizes the `joblib` library to parallelize training data generation across multiple CPU cores. It is optimized for high-performance computing environments with multiple processors to enable the generation of large datasets in a reduced time with following attributes:

- **Parallel Processing:** Distributes the workload across multiple cores using the `loky` backend of `joblib` to achieve high scalability.
- **Batch Processing:** Processes samples in user-defined batches (e.g., 500 samples per batch) and provides intermediate progress updates.
- **Configurable Worker Count:** Allows the user to specify the number of CPU cores to allocate.
- **Shared Data Management:** Ensures that training data structures are correctly updated across all worker processes without conflicts.

3.7. Comparison of Program Variants

Each program variant is tailored to specific computational environments. The GPU-accelerated optimizer is fast yet only scalable for multiple-GPU computing platforms utilizing TensorCores. The single-core optimizer is robust and simple, and thus is suitable for small-scale or debugging tasks. The multi-core optimizer strikes a balance between speed and accessibility. It is well-suited for multi-processor systems without GPUs. For example, 128 cores were used in generating the training data on an HPC herein.

3.8. Visualization Features

Visualization plays a convenient role for understanding the structural behavior and the outcomes of the optimization process in `OpenPyStruct`. The framework provides detailed graphical representations of the structural configurations, internal forces, and design parameter distributions. These visualizations allow users to analyze the effects of design choices, evaluate optimization performance, and gain insights into the structural response under varying loads and boundary conditions. `OpenPyStruct` supports customizable plotting for beam geometry, support conditions, applied loads, second moments of area, shear force diagrams, and bending moment diagrams.

A key feature of the visualization module is the ability to scale beam thickness based on the optimized second moments of area. This ability provides an intuitive understanding of how structural properties are distributed along the span. Support conditions, such as pinned and roller supports, are highlighted using distinct markers, and applied loads are depicted with arrows indicating their magnitude and direction. In addition to geometric visualization, `OpenPyStruct` generates shear force and bending moment diagrams for each structural configuration. Shear forces are visualized using stepped plots, while bending moments are plotted as continuous curves.

3.9. Flexibility in Design Parameters

The modularity of `OpenPyStruct` aims to provide flexibility in defining and optimizing structural design parameters. While the second moments of area I are the primary design parameters in the current implementation, the framework is designed to accommodate alternative or supplementary parameters, such as the cross-sectional area A , elastic modulus E , or even load distributions. This adaptability allows users to tailor the optimization problem to the specific needs of their application, ranging from single-variable optimization to complex multi-variable formulations.

For example, replacing I with A enables the optimization of cross-sectional areas directly, which is particularly useful in scenarios where material usage is a critical factor. Similarly, incorporating E as a variable allows for material property optimization, such as selecting the optimal modulus of elasticity for composite or mixed-material structures. In cases where multiple design parameters need to be optimized simultaneously, `OpenPyStruct` supports multi-objective formulations, enabling trade-offs between competing objectives like cost, weight, and performance.

This flexibility is achieved through the modular structure of the framework, where the loss function, simulation setup, and optimization workflow are parameterized to allow seamless substitution or addition of design variables. Below is an example of modifying the loss function to include cross-sectional area A as a design parameter. Moreover, `OpenPyStruct`'s flexible architecture makes it easy to incorporate additional constraints or physical relationships. For example, if A is proportional to I , the proportionality constant can be explicitly enforced during the optimization process to maintain physical consistency. This adaptability is particularly beneficial in complex structural optimization problems, where multiple design variables interact or where additional constraints must be imposed. By providing a robust and modular framework, `OpenPyStruct` allows researchers and engineers to explore a wide variety of structural optimization tasks without being constrained by a rigid implementation.

3.10. Flexibility in Loss Functions

`OpenPyStruct` supports the integration of a diverse range of loss functions, which can be readily incorporated into the optimization and ML workflows. Below are some potential loss functions that can be integrated within the framework. Each offers flexibility to address different aspects of structural optimization problems. These loss functions can be applied to various stages of the optimization process, from design refinement to advanced ML tasks, enhancing the model's adaptability and performance.

Table 2: Potential loss functions applicable to `OpenPyStruct`: higher-order derivatives are also available via AutoGrad.

Loss Function	Description	Formula
L_2	Penalizes large deviations from target design	$\mathcal{L}_{L_2} = \sum_e (x_e - x_{\text{target},e})^2$
L_1	Minimizes absolute differences	$\mathcal{L}_{L_1} = \sum_e x_e - x_{\text{target},e} $
Total Variation	Reduces abrupt changes in design parameters	$\mathcal{L}_{\text{TV}} = \sum_e x_e - x_{e+1} $
Elasticity	Penalizes deviations from optimal elasticity	$\mathcal{L}_{\text{elasticity}} = \sum_e \left(\frac{1}{E_e} - \frac{1}{E_{\text{target},e}} \right)^2$
Shear Strength	Minimizes shear deformations using shear strength	$\mathcal{L}_{\text{shear}} = \sum_e \left(\frac{V_e^2}{\zeta_e} \right)$
Stress Constraint	Penalty for exceeding stress limits	$\mathcal{L}_{\text{stress}} = \sum_e \max(0, \bar{\sigma}_e - \bar{\sigma}_{\text{max}})^2$
Geometric	Penalizes deviations from nonlinear behavior	$\mathcal{L}_{\text{geometric}} = \sum_e \left(\frac{g_e}{g_{\text{linear}}} - 1 \right)^2$
Fatigue	Penalizes high-cycle stresses and strains	$\mathcal{L}_{\text{fatigue}} = \sum_e \left(\frac{\bar{\sigma}_e}{\bar{\sigma}_{\text{fatigue}}} \right)^2$
Buckling	Prevents buckling by penalizing high risk	$\mathcal{L}_{\text{buckling}} = \sum_e \left(1 - \frac{\bar{P}_{\text{critical},e}}{\bar{P}_{\text{applied},e}} \right)^2$
Energy Dissipation	Optimizes energy dissipation and stability	$\mathcal{L}_{\text{energy}} = \sum_e \frac{\Xi_{\text{dissipated},e}}{\Xi_{\text{input},e}}$
Displacement	Penalizes large displacements, ensuring stability	$\mathcal{L}_{\text{displacement}} = \sum_e (\delta_e - \delta_{\text{target},e})^2$
Rotation	Minimizes rotational displacement within limits	$\mathcal{L}_{\text{rotation}} = \sum_e (\theta_e - \theta_{\text{target},e})^2$
Curvature	Ensures curvature stays within desired bounds	$\mathcal{L}_{\text{curvature}} = \sum_e (\kappa_e - \kappa_{\text{target},e})^2$
Deflection	Minimizes deflection at specific points	$\mathcal{L}_{\text{deflection}} = \sum_e (\delta_e - \delta_{\text{max},e})^2$

Each of these loss functions can be integrated into `OpenPyStruct`'s optimization workflow by simply modifying the loss function component in the design parameter update rule. This flexibility allows users to prioritize specific structural performance criteria, such as reducing stress, enhancing material usage efficiency, or minimizing fatigue damage, all while ensuring that the solution adheres to physical laws and constraints.

Moreover, `PyTorch`'s computation graph enables efficient calculation of gradients for these diverse loss functions, ensuring that complex multi-objective optimization problems can be tackled in a straightforward manner. By adjusting the weights or combining different loss functions, users can customize their optimization objectives to match the requirements of specific structural design tasks, facilitating the creation of robust, efficient, and cost-effective structures.

3.11. Data Generation Workflow Summary

Lastly, we provide pseudo code in Algorithm 1 summarizing the general workflow for the data generators and the JSON data structures the generators save. In the present configuration, the design variables distributions being saved are I , and the additional variables being saved are used as ML model inputs: [support locations, point force locations, point force magnitudes, node locations], physics constraints or supplementary information used in visualization.

Algorithm 1: Beam Data Generation and Storage

Input : Beam parameters, load settings, optimization settings
Output : JSON file containing beam analysis data

- 1 Initialize beam parameters, loads, optimization settings, and an empty data structure
- 2 **Function SetupModel():**
 - 3 | Define beam nodes and elements
 - 4 | Apply supports and loads
 - 5 | Run static analysis
- 6 **Function GenerateSample():**
 - 7 | Randomize beam properties (if enabled)
 - 8 | Assign roller supports and force locations
 - 9 | Initialize moment of inertia values
 - 10 **for** *each optimization step* **do**
 - 11 | **SetupModel()**
 - 12 | Compute shear forces and bending moments
 - 13 | Update inertia values using optimization
 - 14 | Check stopping criteria
 - 15 | Store computed sample data
- 16 **Function Main():**
 - 17 **for** *each batch of samples* **do**
 - 18 | Generate multiple samples in parallel by calling **GenerateSample()**
 - 19 | Filter and store valid results
 - 20 | Print progress
 - 21 | Save all data as a JSON file
- 22 **Main()**

JSON Data Structure:

- "roller_x_locations": List of roller support positions
- "force_x_locations": List of applied force positions
- "force_values": List of applied force magnitudes
- "I_values": Moment of inertia values for each element
- "shear_forces": Computed shear forces
- "bending_moments": Computed bending moments
- "node_positions": List of node positions along the beam
- "roller_nodes": Node indices with roller supports
- "force_nodes": Node indices with applied forces
- "num_nodes": Total number of nodes in the beam
- "L": Length of the beam
- "rotations": Nodal rotations from analysis
- "deflections": Vertical deflections at the nodes.

4. OpenPyStruct for Machine Learning Optimization

In this section, we discuss the three core ML models integrated into the `OpenPyStruct` framework: the FNN, PINN, and Transformer-Diffusion model. These models are used to predict optimal structural design parameters based on various input features, such as nodal coordinates, support conditions, and load scenarios. All models are set up for handling multiple or single load cases and support conditions. While the FNN and Transformer-Diffusion model rely purely on data-driven learning, the PINN incorporates physics information. Key functionalities of the ML package, including data handling, loss functions, unique features, architecture, and an overview of the workflow, are illustrated in the flow chart below.

Step	Description
Start: Input Data	The initial data provided by the user or dataset, containing structural configurations and load scenarios.
Data Preprocessing	Cleaning, normalizing, and preparing the input data for further analysis. Includes padding sequences and feature scaling.
Feature Extraction	Extracting and aggregating meaningful features from the input data to ensure relevant information is passed to the model.
Model Processing	The main computational engine processes the extracted features to understand patterns and relationships.
Prediction Generation	Producing predictions based on the processed data. Predictions include key output metrics for structural evaluation.
Output Refinement	Post-processing the predictions to ensure they align with physical constraints or domain-specific requirements.
Evaluation	Comparing predictions against validation data to assess accuracy and model performance (e.g., using metrics like R^2).
Finish: Deployment	Preparing the model for real-world applications, allowing users to input new data for live inference and predictions.

Table 3: General workflow for the ML optimization approaches used in `OpenPyStruct`.

4.1. Loss Function Approach

The loss function is a critical component of any ML framework as it largely controls the learning process and the degree to which a network responds to outliers. While there are many loss choices, we select a joint loss for all models employing a combination of L_1 and L_2 norms weighted dynamically by a trainable parameter, α . This allows the model to adaptively emphasize absolute or squared differences during optimization. Moreover, the loss function includes constraints that penalize predictions exceeding predefined physical boundaries, thereby ensuring the outputs remain realistic and interpretable. In our approach, we write our loss as:

$$\mathcal{L} = \alpha \cdot L_1(y_i, \hat{y}_i) + (1 - \alpha) \cdot L_2(y_i, \hat{y}_i) + \beta \cdot P(\hat{y}_i) + \lambda \cdot L_2(W), \quad (5)$$

where y_i is the true value for the i -th sample, \hat{y}_i is the predicted value, λ is a regularization parameter, W represents the network’s weights to be regularized in an attempt to prevent overfitting and β controls the penalty magnitude. The penalty terms P act as a regularizer, applying additional loss for predictions outside an allowable range for the design parameters (i.e. a box constraint). This box penalty aims to improve convergence and enforce adherence to domain-specific constraints, for example that physical parameters such as I cannot be negative.

4.2. Data Preprocessing and Label Aggregation Approach

The following steps are adopted using data from a generator or using the provided `StructData-Heavy/StructDataLite` datasets:

Data Preprocessing. The raw input data consists of diverse structural configurations, including roller locations, force positions, force values, and node positions. These inputs vary in size and dimension, thus preprocessing to standardize and align the data is required. The steps include:

1. *Padding Sequences:* Input arrays are padded to a consistent length based on the maximum observed size within the dataset. This ensures that all samples conform to the expected dimensionality.
2. *Feature Scaling:* Each input feature is normalized using a standard scaler to zero mean and unit variance. This prevents numerical instabilities and accelerates convergence during training.
3. *Dimensional Padding:* The feature dimensions are padded to be a multiple of the number of transformer attention heads to ensure compatibility with the multi-head attention mechanism.

Label Aggregation. The design variables for each structural configuration are aggregated across multiple design cases, as shown schematically in Fig. 2. The aggregation process computes a unified label that captures the variability across cases while maintaining physical interpretability. The approach involves:

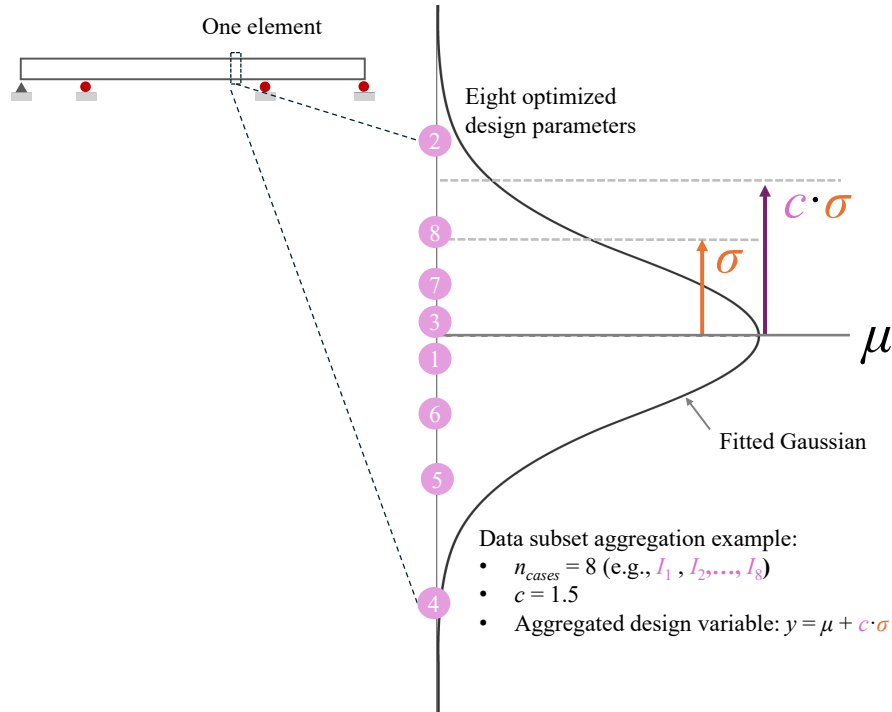


Figure 2: Schematic aggregation of design parameters for a single data subset across multiple load cases on single element.

1. *Statistical Aggregation:* For each element, the mean μ and standard deviation σ of design variables are computed over each subset of output data containing arrays of dimension [number of design cases, number of elements] and used for label unification. This is a key facet in handling optimization over multiple simultaneous structural configurations. The unified label we adopt is written:

$$y = \mu + c \cdot \sigma \quad (6)$$

where c is a scalar coefficient controlling design conservativeness. This value is user-selectable, where $c > 0$ emphasizes a higher threshold of design parameters, for example, c standard deviations above the mean (representing a more conservative choice for I in our case). Conversely, $c < 0$ could be interpreted as a minimal design.

2. *Handling Outliers:* Outliers in the design parameter space are thus mitigated through the unification approach, thus promoting stability in training.

This preprocessing pipeline ensures that the input data and labels are consistent, normalized, and aligned. In addition, functions including mode and median statistics are available in the script which can be used depending on the user’s desired statistical approach.

4.3. Learning Dynamics Approach

The neural networks learn by iteratively minimizing the loss function \mathcal{L} over the training dataset. Gradients of \mathcal{L} with respect to model parameters are computed via backpropagation:

$$\nabla \mathcal{L} = \frac{\partial \mathcal{L}}{\partial W}. \quad (7)$$

These gradients are handled using the PyTorch AutoGrad functionality and are used to update parameters using the Adam optimizer via:

$$W_{s+1} = W_s - l_s \nabla \mathcal{L} \quad (8)$$

where l_s is the learning rate at step s which is gradually adjusted with a scheduler. In our approach, we utilize the GradScaler functionality and the Adam optimizer in PyTorch with Autocast for computational speedup via mixed-precision utilization.

4.4. Evaluation and Inference Approach

The evaluation and inference are designed to assess model performance and enable predictions. These processes are implemented for the user to ensure robustness, interpretability, and efficiency of their ML optimization model.

Evaluation Procedure. The evaluation phase involves:

- **Data Loading:** Validation data is processed into batches using the same preprocessing pipeline as the training data. This ensures consistency in scaling and dimensionality.
- **Batch-wise Predictions:** The model generates predictions for each batch of validation inputs.
- **Metric Computation:** The predictions are compared against ground truth values using the R^2 score after unstandardizing. This metric provides a measure of how well the model’s predictions align with the true values.

Inference Pipeline (user input phase). Inference is structured to handle user-provided input data and generate predictions with the trained model:

- **Input Preparation:** User inputs, such as force positions, values, and node positions, are scaled and padded to align with the model’s expected input dimensions. Features are concatenated to form a unified input tensor.
- **Model Forward Pass:** The prepared input tensor is passed through the model to obtain predictions.
- **Post-Processing:** Predicted values are unscaled and optionally clipped to predefined physical ranges.
- **Visualization:** Results are visualized through custom plots, including representations of forces applied to beams and corresponding second moments of area. Semi-transparent blocks and arrows illustrate predicted and input forces for enhanced interpretability.

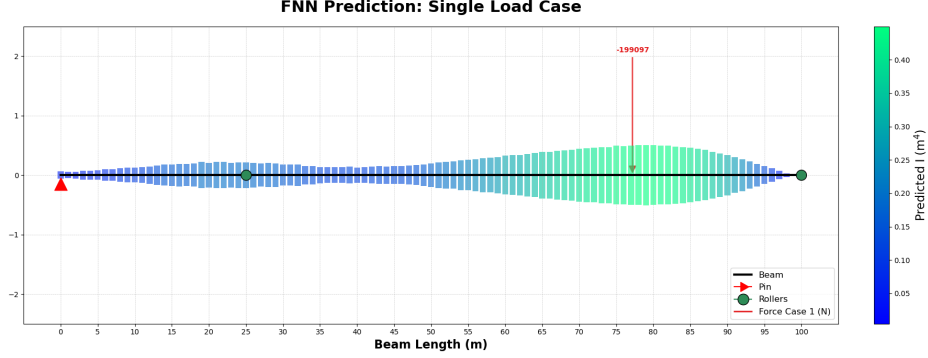


Figure 3: Proof of concept: FNN optimized structure subject to one point load and simple support conditions.

4.5. Feedforward Neural Network Optimizer

The FNN optimizer serves as the foundational model within the `OpenPyStruct` framework. As the most basic neural architecture, the FNN processes data in a unidirectional flow, from input to output, through a sequence of hidden layers. Each layer applies a weighted linear transformation followed by a nonlinear activation, enabling the network to model relationships between structural input features and optimized design parameters.

This model is designed to handle structural optimization tasks with relatively low computational overhead, such as the optimized structure with one load case (more loads cases are possible with the FNN, as described in the Case Studies) shown in Fig. 3. The aim is that the FNN provides the groundwork for introducing more sophisticated models within the framework. In the context of this package, the FNN Optimizer represents the most accessible and computationally lightweight approach. Its straightforward implementation and training pipeline make it an ideal starting point for users new to machine learning in structural optimization. Furthermore, the FNN establishes a benchmark for comparison with more advanced architectures, highlighting the progressive capabilities of `OpenPyStruct` to address increasingly complex structural design challenges.

4.6. Physics-Informed Neural Network Optimizer

The PINN module in `OpenPyStruct` incorporates structural mechanics into the optimization problem to directly constrain and regularize the model during the learning process. Our aim is to enable PINNs to bridge the gap between data-driven machine learning models and physics-based methods by enforcing governing physics, which include equilibrium and compatibility conditions as part of the optimization framework.

Loss Function and Constraints. The loss function in the PINN module is designed to minimize discrepancies in structural parameters while ensuring adherence to physical constraints. The total loss function $\mathcal{L}_{\text{total}}$ combines contributions from the primary data-driven loss \mathcal{L} and physics-informed penalties $\mathcal{L}_{\text{physics}}$. This formulation is expressed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L} + \lambda_{\text{physics}} \mathcal{L}_{\text{physics}} \quad (9)$$

where λ_{physics} is a scaling parameter that balances the influence of the physics-informed penalty. Meanwhile, the data-driven loss is the same as described in Equation 5 while the physics loss $\mathcal{L}_{\text{physics}}$ incorporates penalties for violations of the governing equations. For example, the deflection and rotation residuals are computed as:

$$\mathcal{L}_{\text{deflection}} = \frac{1}{N} \sum_{i=1}^N \frac{|\delta_i - \delta_i^{\text{true}}|}{|\delta_i^{\text{true}}| + \varepsilon}, \quad \mathcal{L}_{\text{rotation}} = \frac{1}{N} \sum_{i=1}^N \frac{|\theta_i - \theta_i^{\text{true}}|}{|\theta_i^{\text{true}}| + \varepsilon}. \quad (10)$$

Here, δ_i and θ_i represent the predicted deflections and rotations, while δ_i^{true} and θ_i^{true} are the corresponding ground truth values. The small constant ε prevents division by zero. It should be noted that residual physics penalties were selected from empirical testing, generally providing better performance than standard L_1 and L_2 loss choices. For bespoke PINN losses for structural design, we refer the reader to [18].

4.7. Transformer-Diffusion Optimizer

This algorithm represents the state-of-the-art in ML based structural optimization – the approach utilizes a transformer-based architecture enhanced with diffusion mechanisms to predict structural responses across various load cases. The framework integrates attention mechanisms and a diffusion module to provide adaptive learning capabilities. Key features include:

- Integration of positional encodings and transformer layers for handling sequence-based input data efficiently;
- Diffusion modules that simulate noise and denoising processes to enhance model generalization and robustness;
- Scalability and adaptability to different structural configurations and load scenarios.

Functional Overview

Architecture. The core architecture comprises three main components: a diffusion module, a transformer encoder, and a fully connected prediction network. The overall structure is as follows:

1. **Diffusion Module:** This module applies Gaussian noise to input features and learns to remove it using a standard perception (a NN). The noise application follows a predefined schedule to enhance model robustness by training it to recover clean signals from noisy data.
2. **Positional Encoding and Transformer Encoder:** Input sequences are encoded with positional information to retain spatial correlations. The transformer encoder, consisting of multi-head attention and feedforward layers, processes these sequences to capture interdependencies between different load cases.
3. **Prediction Network:** The final network layers aggregate the learned features and output predictions for structural properties, such as second moments of area, across multiple elements.

Transformer-Based Processing. The transformer encoder employs multi-head self-attention to model interactions between input features across load cases. Positional encodings are added to retain spatial relationships, and the sequence is passed through stacked encoder layers to extract hierarchical representations. This architecture enables efficient handling of sequence-based inputs, making it ideal for multi-load case analysis.

Diffusion Mechanism. The diffusion mechanism is characterized by a controlled noise injection and removal process designed to enhance model generalization. At each time step τ , Gaussian noise ϵ is added to the input \mathbf{x} to produce the noisy \mathbf{x}_τ . The process can be described as:

$$\mathbf{x}_\tau = \sqrt{\alpha_\tau} \mathbf{x} + \sqrt{1 - \alpha_\tau} \epsilon, \quad (11)$$

where α_τ is a predefined noise scheduling parameter controlling the magnitude of noise. The multi-layer perceptron (MLP) is trained to predict the added noise ϵ such that the clean signal can be recovered:

$$\hat{\epsilon} = \text{MLP}(\mathbf{x}_\tau). \quad (12)$$

The denoised signal $\hat{\mathbf{x}}$ is then computed as:

$$\hat{\mathbf{x}} = \frac{\mathbf{x}_\tau - \sqrt{1 - \alpha_\tau} \hat{\epsilon}}{\sqrt{\alpha_\tau}}. \quad (13)$$

The iterative denoising mechanism ensures robustness against perturbations, stabilizing the learning process and reducing overfitting. The noise scheduling α_τ follows a geometric progression depending on the magnitude of the input \mathbf{x} .

Transformer-Based Processing. The Transformer-based processing employs multi-head self-attention to model dependencies within input sequences. Given an input sequence $\mathbf{X} \in \mathbb{R}^{n \times d}$, where n is the sequence length and d is the feature dimension, the self-attention mechanism computes attention scores \mathbf{A} as:

$$\mathbf{A} = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right), \quad (14)$$

where \mathbf{Q} , \mathbf{K} , and \mathbf{V} are query, key, and value matrices derived from \mathbf{X} , and d_k is the dimensionality of the queries and keys. The output of the attention mechanism is:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{A}\mathbf{V}. \quad (15)$$

Positional encodings \mathbf{P} are added to retain spatial relationships:

$$\mathbf{X}_{\text{encoded}} = \mathbf{X} + \mathbf{P}, \quad (16)$$

where \mathbf{P} is a sinusoidal positional encoding matrix. The encoded sequence is passed through L stacked encoder layers, each comprising:

$$\mathbf{X}^{(l+1)} = \text{LayerNorm} \left(\mathbf{X}^{(l)} + \text{MultiHeadAttention} \left(\mathbf{X}^{(l)} \right) \right), \quad (17)$$

$$\mathbf{X}^{(l+1)} = \text{LayerNorm} \left(\mathbf{X}^{(l+1)} + \text{FeedForward} \left(\mathbf{X}^{(l+1)} \right) \right), \quad (18)$$

where l denotes the layer index, and FeedForward refers to a two-layer fully connected network with appropriate activation. This architecture attempts to efficiently capture hierarchical features across sequences.

5. Examples and Discussion

5.1. Case Study 1: PINN and FNN Simultaneous Optimization for 6 Load Cases with fixed roller locations

In this subsection, a basic case study is conducted on a beam subjected to 6 simultaneous point load cases. The total length of the beam is a 200 m and roller positions are fixed. The data StructDataLite and the FNN optimizer are utilized, and results are compared with the PINN optimizer. We begin by accessing `OpenPyStruct_FNN_MultiCase.py`. First, we observe the configuration block:

```
#####
# 1) CONFIGURATION & HYPERPARAMETERS
#####

# Model and training configuration #
n_cases = 6           # Number of sub-cases per sample
nelem = 100          # Final output dimension per sample: (B, n_elem)
box_constraint_coeff = 5e-1 # Coefficient for box constraint penalty
hidden_units = 128   # Number of hidden units in MLP
dropout_rate = 0.5   # Dropout rate for regularization
num_blocks = 3       # Number of blocks (unused in current model)
num_epochs = 500     # Maximum number of training epochs
batch_size = 128     # Batch size for training
patience = 10       # Early stopping patience
learning_rate = 2e-4 # Learning rate for optimizer
weight_decay = 1e-2  # Weight decay (L2 regularization) for optimizer
train_split = 0.8    # Fraction of data used for training
sigma_0 = 0.03       # Initial Gaussian noise for input
gamma_noise = 0.97   # Decay rate for noise during training
gamma = 0.99         # Learning rate scheduler decay rate
initial_alpha = 0.5  # Initial alpha value for loss weighting
c = 1.0              # Parameter to adjust label aggregation (higher c = more conservative ...)
```

where all key parameters are defined. The total number of load cases is defined as 6 in the first line. Thereafter, the number of elements in the discretization (and length of the I array), constraints, hyperparameters for the

NN, and optimization parameters are input. At present, the defaults will provide a good starting point for optimizing this particular problem. However, user parameters such as the label aggregation parameter c may be throttled to modulate the degree of conservativeness one would like to consider in their final optimized structure. Thereafter, the code scales automatically with respect to data sizes and the optimizer executes seamlessly with respect to the predefined parameters.

During training, the training and validation loss curves are provided at each epoch as shown in Fig. 4. It is important to note that hyperparameters including the dropout rate, weight decay, initial added noise level σ_0 and the size of the network all influence the susceptibility for the network to underfitting or overfitting. As a rule of thumb it is good practice to adjust these such that the model reaches a balance between the training and validation set.




Figure 4: Training and validation losses plotted at each epoch for the case study 1.

Training will then complete once the maximum number of epochs has been attained or the process is terminated early via the provided stopping criteria. The best model is automatically saved at each epoch so that the most suitable model is used in prediction. Here, the default patience is 10, meaning that the training will stop if the validation loss has not improved in 10 epochs. This can be reduced if a particular design case has a consistently low slope after the initial learning is complete (i.e. after the initial knee point). Following, the R^2 values are computed from all validation data, comparing true validation data against all predictions from validation data. In general, as the number of input data (e.g. load cases) increases, so does the difficulty in solving the optimization problem, so users may expect successively decreased R^2 in direct proportion to the number of cases considered. For simple problems, for example cases < 4 , a good goal (based on our empirical trailing) is to exceed an R^2 value greater than 0.75 with the FNN ². For this case study, we report an $R^2 = 0.71$ which is deemed satisfactory given the large output space and aggregation of 6 I distributions.

Lastly, the user is able to make predictions of optimized design parameters at the inference and plotting phase. While the default setting predefines the roller locations and randomizes force magnitudes and locations consistent with the training data, these can be set by the user. For example, if all roller locations, force locations, force magnitudes and beam lengths are randomized, the inputs can be arbitrarily set. In the default configuration using the provided StructLite data, the user may input desired parameters from here:

```
# Example User Inputs
L_beam = 200 # Beam length (m)
Fmin_user = -355857 # Min point load to be randomized (N)
```

²The transformer-diffusion model is expected to outperform both of these models with R^2 in excess of 0.9 using the StructLite dataset

```

Fmax_user = Fmin_user / 10 # Max point load to be randomized (N)
user_rollers = [2*9, 2*29, 2*69, 2*85, 2*100]

```

where we emphasize again that these values should fit within the space of the training data (current entries are valid for StructLite and StructHeavy). After this point, the user may execute the rest of the code to visualize the output, as seen in Fig. 5 for the FNN. In this figure, the distribution of optimized design parameters I is shown with associated color mapping and sizing, along with the location and magnitude of the test forces represented by plotted arrows and the support locations. It is noted that the I is larger near the areas of expected maximum moment (above the interior supports and midspans) and smaller at the end spans, which is an anticipated result.

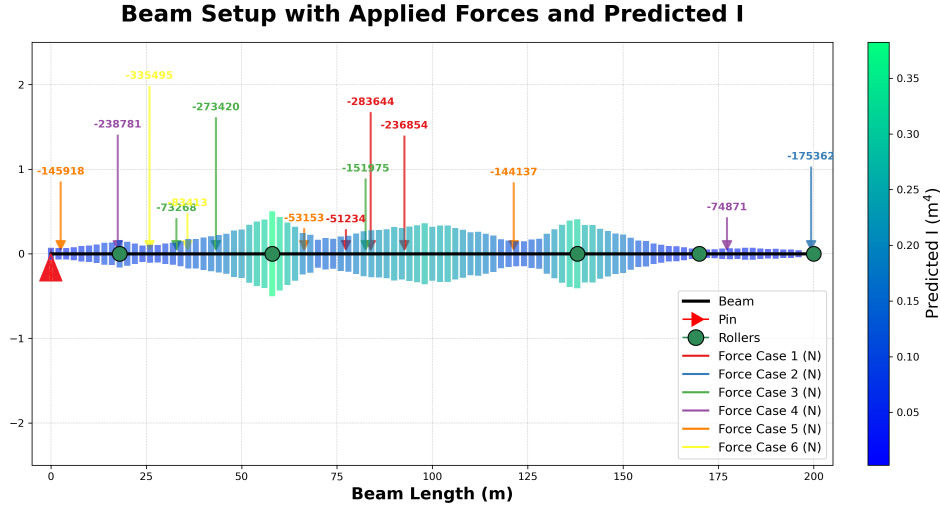


Figure 5: Optimized design parameters I predicted by the FNN optimizer including the locations of test forces, force magnitudes and roller locations.

To conclude this section, we will compare the FNN results with the PINN optimizer which integrates additional mechanics information. In the PINN evaluation, we report an $R^2 = 0.77$, which is a noteworthy improvement over the FNN and is owed to the NN learning the physics of the underlying design problem and resulting regularization. Next, we plot the optimized parameter space for the same randomized testing scenario as shown in Fig. 6. Here, we additionally plot the predicted displacements and rotations in the bottom two subplots respectively. We observe a general expected consistency with the problem mechanics, for example 0 displacements at the rollers and consistent rotations at the rollers. We emphasize that these predicted displacements and rotations are predicted from the aggregation of all output training data and are *not* a summation or combination of all load cases in an additive sense. However, we do observe some displacement artifacts near the rollers and pin at the left hand side. These are owed to their relatively lower contribution to the overall loss function, with respect to the larger displacements at midspan. In future iterations of this code, weighted loss could be incorporated to mitigate these artifacts.

5.2. Case Study 2: Transformer-Diffusion Simultaneous Optimization for Arbitrary Load Cases, Spans and Support Conditions

In this case study we aim to predict optimized I using maximally-randomized data – i.e. by randomizing roller number, locations, span widths and force locations/magnitudes in the training data with $c = 1.0$. For this, we set a maximum allowable structural length of 100 m, clamp the left side at $x = 0$ m and randomize the number of rollers and number of applied point loads between 1 and 8 and sample 500k data points. Given the ill-posed nature and high dimensionality of this inverse problem, we select the transformer-diffusion model as our optimizer and increase the number of neurons in the feedforward layers to 512, attention heads to 24 and transformer blocks to 4.

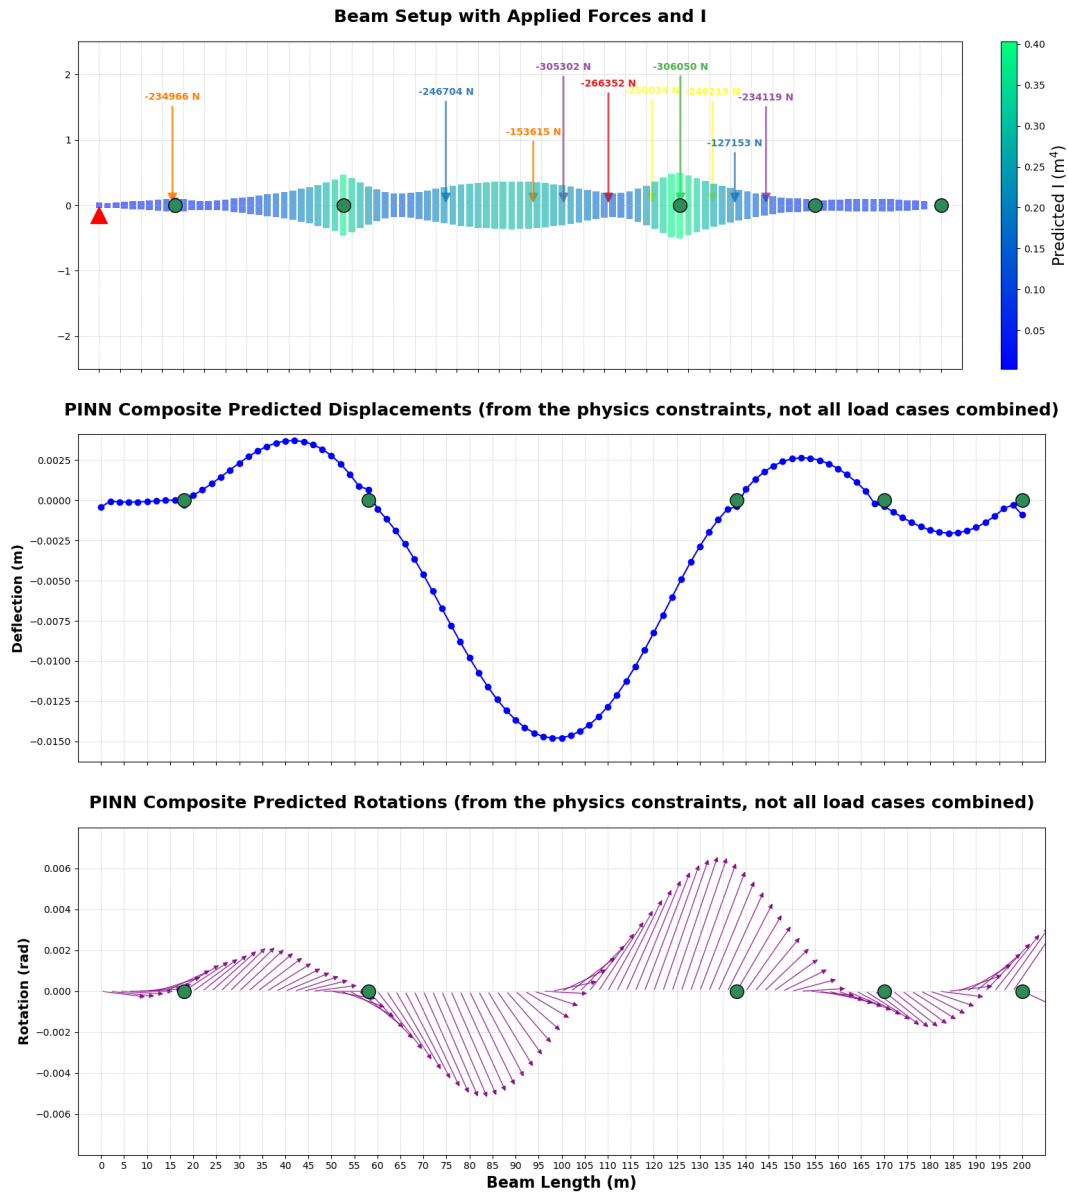


Figure 6: Top: Optimized section properties predicted by the PINN optimizer including the locations of test forces, force magnitudes and roller locations; Middle: Predicted displacement field; Bottom: Predicted rotations where arrow directions indicate rotation angle and magnitude. Note the beam lines are removed for better visualization of displacements and rotations.

After training, we test our model using 4 different trial support cases and 6 randomized force/magnitudes cases, as reported in Fig. 7. As was discussed for the previous case study, we consistently see higher I magnitudes at areas where the moment is expected to be elevated while the reverse is also observed (albeit, predictions in this case study are smoother due to the highly ill-posed learning problem presented here despite

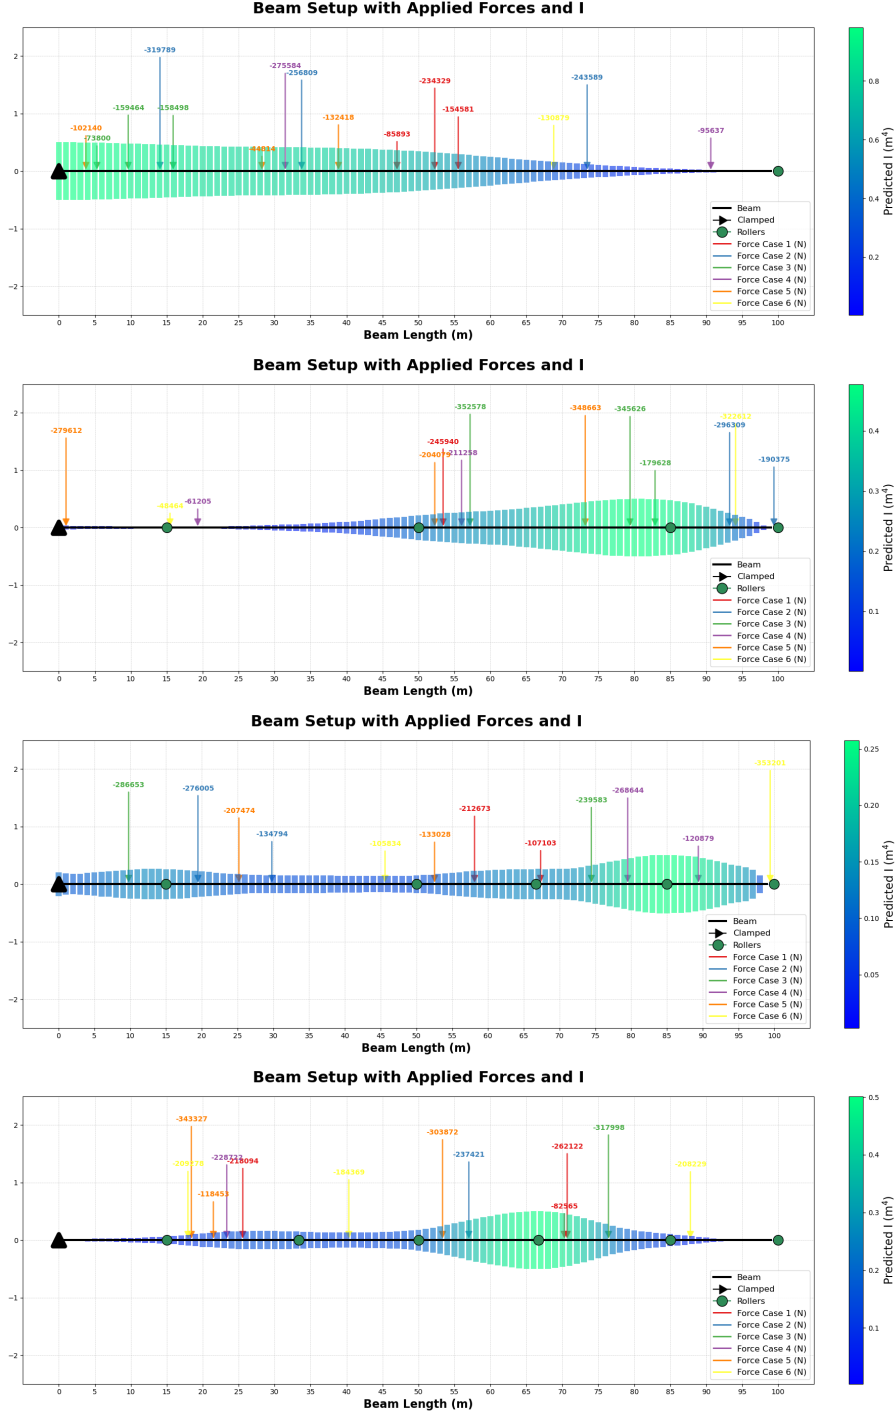


Figure 7: Transformer-Diffusion model predictions of optimized I for various random load and support conditions. Training was carried out using randomized 500k data (including roller number, locations, span widths and force locations/magnitudes). Test case predictions include a *clamped* left end and various support conditions increasing from the top row to the bottom row.

an $R^2 > 0.92$). These observations support the feasibility of the diffusion-transformer model for predicting optimized structures subject to arbitrary support and loading conditions when data is randomized. This is a desired feature as we aim to train a model with as little bias towards a given suite of support or loading conditions as possible. Lastly, this supports the potential for generalization of the transformer-diffusion model

for handling more complex structures, such as 2D/3D frames.

5.3. Computational Considerations

For the previous cases studies, a local machine with an NVIDIA RTX 4080, Core i9-14900K CPU and 64 GB DDR6 RAM was utilized. For these problems, this is more than sufficient to quickly (less than 8 seconds per epoch for a total of 200 or less epochs while employing the Transformer-Diffusion Model at default settings) train the models. Upon further testing, it was found that similar spec laptops with an NVIDIA A2000 or RTX 4060 were also adequate, albeit at approximately half the speed per epoch. In general, for the case studies and more broadly, it was found that the major computing bottlenecks include (a) the number of CPUs or GPUs available for generating training data and (b) the available VRAM when training models with StructDataHeavy. For example, when trialing the multicore data generator with the local machine previously mentioned, generating 2,000,000 data samples took over 24 hours (hence why this was done on a remote HPC). In the case of limited VRAM (RTX 4060 machine), training issues (e.g. 'out of memory' warnings) may be mitigated by using smaller batch sizes (while accepting higher noise in the gradients during training) and overall increased training time. Lastly, for machines with limited RAM, it is advisable to load massive data in chunks rather than the default setting (noting that StructDataHeavy is > 24 GB).

5.4. Discussion

5.4.1. Potential for Extended Capabilities: Simultaneous Load and Support Optimization

OpenPyStruct's current v1 framework handles multiple load and support conditions (although the default configuration is meant to have consistent supports and randomized force locations), enabling multi-load case conceptual structural optimization. Building upon this foundation, there exists significant potential to extend the toolkit to perform simultaneous optimization of both loading conditions and support configurations. This dual optimization approach can lead to the discovery of more resilient and efficient structural designs by considering the interplay between applied forces and support placements holistically.

By integrating advanced multi-objective optimization algorithms, OpenPyStruct could balance conflicting objectives such as minimizing material usage while maximizing structural integrity under diverse loading scenarios. Leveraging machine learning models, particularly those capable of capturing complex dependencies like Transformer-Diffusion models, the framework can predict optimized configurations that simultaneously address multiple loads and supports. Additionally, incorporating reinforcement learning techniques could enable the system to iteratively learn and adapt support strategies in response to, for example, dynamic loading conditions. Such an extension could streamline design workflows by reducing the need for sequential optimizations and facilitate the exploration of novel structural configurations that might be overlooked in traditional optimization paradigms.

Future developments could include user-friendly interfaces for defining complex load and support interactions, real-time optimization feedback, and enhanced visualization tools — potentially integrated with CAD — to intuitively present the interplay between optimized loads and supports. Given that OpenPyStruct is developed in CPython, the authors aim to explore its integration into CAD environments that incorporate visual programming interfaces, such as Rhinoceros & Grasshopper, Revit & Dynamo as well as Blender & Sverchok. To facilitate this, OpenPyStruct could be packaged as a core library that is referenced in each of the visual programming interfaces and accessed via standardized API calls. Such an integration would expose OpenPyStruct with both the design community and practitioners, which in turn could contribute to the project with real-world test cases.

While OpenPyStruct, in principle, is ready to be used as an aid for teaching, we currently forego communities without at least an intermediate set of Python skills. Integrating with the user interfaces discussed above would allow OpenPyStruct to be used in more architectural communities. Such advancements would help position OpenPyStruct as an integrable tool capable of tackling a suite of practical structural optimization challenges.

5.4.2. Potential for Extended Capabilities: Integration of Novel Standardization

Standardization plays a pivotal role in ensuring consistency, interoperability, and scalability within structural optimization frameworks. OpenPyStruct stands to benefit from the integration of novel standardization

protocols, which can enhance its robustness and facilitate broader adoption across diverse design applications. One avenue for novel standardization is the adoption of industry-recognized data schemes for structural properties, load conditions, and support configurations. Implementing standardized input and output formats would simplify data exchange between `OpenPyStruct` and other engineering software, enabling seamless integration into existing design pipelines. Additionally, embracing standardized application programming interfaces (APIs) and modular architectures can promote extensibility, allowing users to incorporate custom modules or integrate with other machine learning frameworks without significant overhead.

Furthermore, integrating standards related to model validation and benchmarking can enhance the credibility and reliability of `OpenPyStruct`'s optimization outcomes. Establishing standardized test cases and performance metrics would facilitate comparative studies, enabling users to assess model performance consistently and identify areas for improvement. Collaborative efforts to define and adhere to such standards may also be used to drive the development of best practices within the open-source structural community.

5.4.3. Potential for Extended Capabilities: Integration of Novel Losses

The optimization efficacy of `OpenPyStruct` is inherently tied to the design and implementation of its loss functions, which guide the model towards desirable structural configurations. Integrating novel loss functions presents a promising avenue to enhance the framework's ability to capture intricate design nuances, enforce multifaceted constraints, and prioritize specific performance criteria. By diversifying the loss landscape, `OpenPyStruct` can achieve more nuanced and context-aware optimizations, tailored to the unique demands of various engineering scenarios.

One potential innovation is the incorporation of multi-objective loss functions that simultaneously address multiple performance metrics, such as enforcing stability, reducing material usage, and enhancing structural stability. Techniques like Pareto optimization can be employed to balance these competing objectives, enabling the discovery of optimal trade-offs that align with real-world structural engineering priorities. Additionally, incorporating uncertainty-aware loss functions can bolster the framework's resilience to variability in load conditions and material properties, fostering designs that maintain performance under unforeseen circumstances.

Another promising direction is the integration of physics-informed loss components that encapsulate higher-order mechanical principles or emergent phenomena not currently captured by existing loss terms. For instance, losses that penalize excessive vibrations, thermal stresses, or dynamic load responses can enrich the optimization process, leading to more comprehensive and reliable structural designs. Moreover, leveraging adversarial loss functions, inspired by generative adversarial networks (GANs), might enhance the model's ability to generate realistic and robust structural configurations by challenging it to outperform a discriminator network trained to identify suboptimal designs.

Lastly, adaptive loss functions that evolve during the training process based on performance feedback can offer dynamic optimization pathways, possibly enabling `OpenPyStruct` to refine its focus as the design evolves. Such adaptability can be beneficial in complex optimization landscapes where static loss functions may struggle to navigate local minima or saddle points. By experimenting with and integrating these novel loss paradigms, `OpenPyStruct` has potential to significantly elevate its optimization capabilities, delivering more sophisticated and contextually appropriate structural solutions.

5.4.4. Potential for Extended Capabilities: Simplification Using Standard Shapes and Inclusion of Environmental Effects

Simplified structural modeling approaches are often necessary in engineering practice where clarity, speed, and code compliance play critical roles. Rather than optimizing arbitrary or continuous cross-sectional properties, practitioners typically select from a discrete set of standardized sections (e.g., American Wide Flange (W) shapes, British Universal Beams (UB), etc.) that already meet design and manufacturing constraints. In the case of localized ML-based conceptual design of steel members integrating the influence zone concept [5], models used in [1, 2] have demonstrated promise via incorporating discrete members. Indeed, this may simplify the decision space (and ill-posedness of the design problem) and ensure that resulting structures are more readily constructible in some scenarios. Moreover, standardized shapes come with comprehensive tabulated design properties, proven fabrication methods, and guidelines that reduce the risk of design errors. In this regard, hybrid methods incorporated in `OpenPyStruct` with discrete selection of recognized section shapes have promise for synergizing automated optimization and practical engineering specifications.

In addition to geometric and material properties, continuous structural designs often need to account for environmental effects such as temperature fluctuations and wind loads. Temperature changes, in particular, can lead to deformation of elements that induce significant stress in statically indeterminate systems or systems with constrained supports. Therefore, when training data-driven structural models, introducing features related to temperature ranges or thermal load distributions can improve the fidelity and designability of the resultant design recommendations. It may be beneficial to represent such environmental boundary conditions explicitly in the optimization framework – either in physics constraints or as part of the input features for data-driven models – thereby ensuring that predicted designs include these real-world variations.

Combining standardized shapes and environmental modeling could expand the applicability of computational frameworks proposed here beyond purely idealized scenarios. This integrated approach could allow designers to more reliably and transparently assess design trade-offs involving discrete section selection, temperature-induced effects, and other environmental considerations. Indeed, future versions of `OpenPyStruct` should likely incorporate industry-standard steel or concrete shape libraries alongside data describing average regional temperature swings, enhancing the realism of automated ML-based design.

6. Conclusion and Future Work

6.1. Summary of Key Contributions

In our present contribution, we introduced an open-source toolkit `OpenPyStruct` (Toolkit: <https://github.com/dsmly6/OpenPyStruct.git>, Data Repository: <https://zenodo.org/records/14742438>), designed to bridge physics-based structural optimization with conventional and modern machine learning techniques. `OpenPyStruct` integrates `OpenSeesPy` for precise structural forward simulations and `PyTorch` for efficient gradient-based optimization, facilitating the optimization of structural properties under diverse loading and support conditions. Key contributions of this project include:

1. **Comprehensive Framework:** Development of a modular architecture that supports single and multi-load case optimizations, enabling versatile applications in structural engineering.
2. **Scalable Data Generation:** Implementation of single-core and multi-core data generators, alongside GPU-accelerated optimizers, allowing for the creation of extensive datasets essential for training high-fidelity machine learning models.
3. **Modern Machine Learning Integration:** Incorporation of various ML models, including Feedforward Neural Networks, Physics-Informed Neural Networks, and Transformer-Diffusion architectures, to enhance predictive capabilities and optimization performance.
4. **Flexible Optimization Tools:** The suite of customizable loss functions and design parameters, empowering users to tailor the optimization process to specific engineering requirements and constraints.
5. **Robust Visualization and Analysis:** Development of intuitive visualization modules that aid in the interpretation of structural behavior and optimization outcomes, fostering a deeper understanding of design impacts.

6.2. Future Directions for `OpenPyStruct` and Machine Learning Structural Optimization

Looking ahead, there are several promising avenues to further enhance `OpenPyStruct` and the broader domain of machine learning-driven structural optimization:

- **Enhanced Multi-Objective Optimization:** Expanding the framework to support simultaneous optimization of multiple (non-physics based) objectives, such as cost, weight, and structural resilience, will allow for more comprehensive and balanced design solutions.
- **Integration of Advanced Standardization Protocols:** Adopting and contributing to emerging industry standards for data representation and model interoperability can improve consistency, facilitate collaboration, and promote broader adoption of `OpenPyStruct` across different open-source engineering applications.

- **Development of Novel Loss Functions:** Exploring and incorporating innovative loss functions that capture higher-order mechanical principles or specific engineering constraints can refine optimization outcomes, ensuring designs are not only efficient but also robust and reliable under various conditions.
- **Incorporation of Uncertainty Quantification:** Integrating techniques to account for uncertainties in material properties, load conditions, and environmental factors can enhance the reliability and safety of optimized structural designs.
- **Expansion to 2D and 3D:** As the availability of powerful ML approaches increases, even beyond the robust transformer-diffusion model, the ability for future `OpenPyStruct` ML design optimizers to handle evermore diverse and massive data sets (and their resultant increasing degree of ill posedness) will also increase. As this field progresses, we aim to expand the toolkit to handle progressively more difficult conceptual design tasks.
- **Simplification Using Standard Shapes and Inclusion of Environmental Effects:** Incorporating standard shapes (e.g., *W* sections) and environmental effects (e.g., temperature) into `OpenPyStruct` may enhance practicality, realism, and compliance in many scenarios, paving the way for more robust and constructible designs.
- **User-Friendly Interfaces:** Improving the accessibility of `OpenPyStruct` through intuitive user interfaces will lower the barrier to entry, enabling a wider range of users to engage the toolkit effectively.
- **Collaborative and Community-Driven Development:** Encouraging community contributions and fostering collaborative development efforts can accelerate the evolution of `OpenPyStruct`, ensuring it remains at the cutting edge of structural optimization technology.

By pursuing these directions, `OpenPyStruct` aims to continuously evolve through community contributions, where we mean to meaningfully contribute to addressing some of the complex challenges of structural optimization and design. Thus, we aim to refine the toolkit based on user feedback and emerging research/industry trends to support the structural engineering community with an evermore robust, efficient, and innovative optimization toolkit.

7. Data Availability Statement

Code and data are openly available at the respective locations: `OpenPyStruct` Toolkit: <https://github.com/dsmyl6/OpenPyStruct.git>, Data Repository: <https://zenodo.org/records/14742438> .

8. Statements and Declarations

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The authors have no competing interests to declare that are relevant to the content of this article.

9. CRediT Author Statement

Danny Smyl: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization, Supervision, Project administration; **Bozhou Zhuang:** Methodology, Software, Validation, Writing - Original Draft, Writing - Review & Editing, **Sam Rigby:** Validation, Writing - Review & Editing; **Edvard Bruun:** Validation, GitHub implementation, Writing - Review & Editing; **Brandon Jones:** Validation, Writing - Review & Editing; **Patrick Kasther:** Validation, Writing - Review & Editing; **Iris Tien:** Validation, Writing - Review & Editing **Adrien Gallet:** Software, Validation, Writing - Review & Editing.

10. Appendix A: Summary of Variables and Functions

This section provides a summary of the key variables and functions employed in `OpenPyStruct` as follows:

Variables and Descriptions	
E	Young's modulus
G	Shear modulus
L	Total length of the beam
N	Number of nodes in the beam model
M	Number of elements in the beam model
I_0	Initial guess for the second moments of area for each beam element
α_{moment}	Weighting coefficient controlling the bending loss term in the total objective
α_{shear}	Weighting coefficient controlling the shear loss term in the total objective
I_e	Second moment of area for the e -th beam element, the primary design variable
M_e	Bending moment at the e -th beam element (defined at end nodes)
V_e	Shear force at the e -th beam element (defined at end nodes)
A_e	Cross-sectional area of the e -th beam element
$\mathcal{L}_{\text{primary}}$	Primary loss function minimizing the sum of second moments of area I_e
$\mathcal{L}_{\text{bending}}$	Bending loss function penalizing bending energy
$\mathcal{L}_{\text{shear}}$	Shear loss function penalizing shear energy
$\mathcal{L}_{\text{total}}$	Total loss function combining primary, bending, and shear loss terms
Functions and Descriptions	
<code>BeamOptimization()</code>	Main function coordinating the beam optimization procedure
<code>ops.analyze(1)</code>	Executes a static analysis on the beam model for one iteration using <code>OpenSeesPy</code>
<code>total_loss.backward()</code>	Computes the gradients of the total loss function with respect to the design parameters
<code>optimizer.step()</code>	Updates the design parameters (second moments of area I_e) based on the gradients
<code>scheduler.step()</code>	Adjusts the learning rate according to the predefined schedule to improve convergence

Table 4: Summary of Key Non-Visualization Variables and Functions used in the Beam Optimization Algorithm.

Variables and Descriptions	
<code>n_cases</code>	Number of sub-cases per sample
<code>nelem (d_{elem})</code>	Number of elements in the model
<code>hidden_units</code>	Number of hidden units per perceptron layer
<code>dropout_rate</code>	Dropout rate for regularization
<code>num_epochs</code>	Maximum number of training epochs
<code>batch_size (B)</code>	Batch size for training data
<code>learning_rate l_r</code>	Learning rate for optimization
<code>weight_decay (λ)</code>	L_2 regularization coefficient for the optimizer
<code>initial_alpha (α_0)</code>	Initial weighting coefficient for loss terms
<code>box_constraint_coeff (β)</code>	Coefficient for penalty constraints on predictions
<code>diffusion_T (T)</code>	Number of steps in the diffusion process
<code>dim_feedforward</code>	Dimension of feedforward layers in Transformer
<code>num_heads</code>	Number of attention heads in Transformer layers
<code>sigma_0 (σ_0)</code>	Initial noise level for inputs
<code>gamma_noise (γ_{noise})</code>	Decay rate for noise during training
<code>gamma (γ)</code>	Learning rate decay factor for the scheduler
<code>c</code>	Scaling factor for unifying across the output space
<code>deflection_dim (d_δ)</code>	Dimension of deflections in the output
<code>rotation_dim (d_θ)</code>	Dimension of rotations in the output
<code>output_dim (d_{output})</code>	Total output dimension: $d_{\text{output}} = d_{\text{elem}} + d_\delta + d_\theta$
Functions and Descriptions	
<code>pad_sequences()</code>	Pads sequences to a uniform length for batch processing
<code>unify_label_with_c()</code>	Aggregates labels via $\mu + c \cdot \sigma$, where μ is the mean and σ is the std
<code>fit_transform_3d()</code>	Applies scaling to 3D arrays using a <code>StandardScaler</code>
<code>merge_sub_features()</code>	Concatenates multiple feature arrays into a single array
<code>scale_user_inputs()</code>	Prepares and scales user-provided input features for modeling
<code>ResidualBlock</code>	Implements residual connections with normalization and convolutions
<code>FNNWithResidual</code>	Feedforward neural network with residual blocks and normalization
<code>TrainableL1L2Loss</code>	Custom loss combining L_1 and L_2 terms weighted by α
<code>CompositeLoss</code>	Combines loss terms for I , deflections, and rotations
<code>DiffusionModule</code>	Implements noise addition and removal using a perception
<code>PositionalEncoding</code>	Adds positional encodings
<code>TransformerWithDiffusion</code>	Transformer-based model incorporating a diffusion module.

Table 5: Summary of Key Non-Visualization Variables, Mathematics, and Functions from the Modules.

11. Appendix B: Classical Optimization Pseudocode

Algorithm 2: Beam Optimization Algorithm using OpenSeesPy and PyTorch

Data: Material properties E, G ; cross-sectional area A ; beam length L ; number of nodes N ; number of elements M ; initial guess for I_0 ; optimization parameters (learning rate, momentum, weight decay) **Result:** Optimized second moments of area I_{opt}

1 Initialization

2 | Set initial guess I_0 for each element; Initialize learning rate, momentum, and regularization parameters

3 Beam Structure Setup

4 | Compute node positions based on N and L
5 | Define roller supports and point forces
6 | Specify boundary conditions (pin and roller supports)
7 | Apply uniformly distributed load (UDL) and point forces

8 OpenSees Model Setup

9 | **for** $e = 1$ **to** M **do**
10 | | Define beam nodes and apply boundary conditions
11 | | Define element properties with updated second moments of area I_e
12 | | Configure static analysis in `OpenSeesPy`

13 Loss Function Definition

14 | **for** $e = 1$ **to** end **do**

15 | | Compute bending energy loss:

$$\mathcal{L}_{\text{bending}} = \sum_{e=1}^M \frac{M_e^2}{2E I_e}$$

16 | | Compute shear energy loss:

$$\mathcal{L}_{\text{shear}} = \sum_{e=1}^M \frac{V_e^2}{G A_e}$$

17 | | Compute primary loss:

$$\mathcal{L}_{\text{primary}} = \sum_{e=1}^M I_e$$

18 | | Compute total loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{primary}} + \alpha_{\text{moment}} \mathcal{L}_{\text{bending}} + \alpha_{\text{shear}} \mathcal{L}_{\text{shear}}$$

19 Optimization Loop

20 | **for** $epoch = 1$ **to** max_epochs **do**
21 | | `optimizer.zero_grad()`
22 | | Rebuild OpenSees model with current I_{tensor}
23 | | `ops.analyze(1)`
24 | | `total_loss.backward()`
25 | | `optimizer.step()`
26 | | `scheduler.step()`
27 | | Clamp I_{tensor} to avoid negative second moments of area
28 | | **if** $total_loss \leq best_loss - tolerance$ **then**
29 | | | $best_loss = total_loss$
30 | | | $no_improvement_epochs = 0$
31 | | | **else**
32 | | | | **if** $no_improvement_epochs \geq patience$ **then**
33 | | | | | `break`
34 | | | | $no_improvement_epochs = no_improvement_epochs + 1$

35 **return** optimized second moments of area I_{opt}

12. Appendix C: ML Optimization Pseudocode (Transformer-Diffusion Example)

Algorithm 3: Machine Learned Multi Load Case Optimizer

Data: Input features: Roller locations, Force positions, Force values, Node positions; Ground truth: second moments of area

Result: Predicted second moments of area for structural optimization

- 1 **Initialize**
 - 2 Define hyperparameters: $n_{\text{cases}}, n_{\text{elem}},$ dropout rate, learning rate;
 - 3 Configure diffusion parameters: T, α_τ ;
 - 4 Initialize device (CPU or GPU);
 - 5 **Input Preprocessing**
 - 6 Load and normalize input data;
 - 7 Pad sequences to align dimensions;
 - 8 Aggregate labels using mean and standard deviation with coefficient c ;
 - 9 Split data into training and validation sets;
 - 10 Standardize input features and targets;
 - 11 **Model Definition**
 - 12 Define **Diffusion Module** for noise injection and removal:

$$\mathbf{x}_t = \sqrt{a_\tau} \mathbf{x} + \sqrt{1 - a_\tau} \epsilon$$

Predict noise ϵ and recover signal:

$$\hat{\mathbf{x}} = \frac{\mathbf{x}_t - \sqrt{1 - a_\tau} \hat{\epsilon}}{\sqrt{a_\tau}}$$
 - 13 Implement **Transformer Encoder** with positional encoding:

$$\mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right), \quad \mathbf{X}_{\text{encoded}} = \mathbf{X} + \mathbf{P}$$
 - 14 Define fully connected network for final prediction;
 - 15 **Training**
 - 16 Use Adam optimizer with weight decay and learning rate scheduling:

$$l_s = l_0 \cdot \gamma^s$$

Train for n_{epochs} using a custom loss function combining L_1 and L_2 losses, and penalties:

$$\mathcal{L} = \alpha \cdot L_1(y_i, \hat{y}_i) + (1 - \alpha) \cdot L_2(y_i, \hat{y}_i) + \beta \cdot P(\hat{y}_i) + \lambda \cdot L_2(W)$$

Apply early stopping based on validation loss;
 - 17 **Evaluation**
 - 18 Load the best-performing model;
 - 19 Evaluate predictions on the validation set;
 - 20 Compute performance metrics (e.g., R^2 score);
 - 21 Visualize predictions alongside ground truth second moments of area;
-

References

- [1] A. Gallet, Machine learning for structural design models from the inverse problem perspective, Ph.D. thesis, University of Sheffield (2024).
- [2] A. Gallet, S. Rigby, T. Tallman, X. Kong, I. Hajirasouliha, A. Liew, D. Liu, L. Chen, A. Hauptmann, D. Smyl, Structural engineering from an inverse problems perspective, Proceedings of the Royal Society A 478 (2257) (2022) 20210526.
- [3] E. Turco, Tools for the numerical solution of inverse problems in structural mechanics: review and research perspectives, European Journal of Environmental and Civil Engineering 21 (5) (2017) 509–554.
- [4] J. L. Mueller, S. Siltanen, Linear and nonlinear inverse problems with practical applications, SIAM, 2012.

- [5] A. Gallet, A. Liew, I. Hajirasouliha, D. Smyl, Influence zones of continuous beam systems, in: Structures, Vol. 68, Elsevier, 2024, p. 107069.
- [6] F. Almeida, A. Awruch, Design optimization of composite laminated structures using genetic algorithms and finite element analysis, Composite structures 88 (3) (2009) 443–454.
- [7] V. V. Toropov, A. Filatov, A. Polynkin, Multiparameter structural optimization using fem and multi-point explicit approximations, Structural optimization 6 (1993) 7–14.
- [8] J. Kudela, R. Matousek, Recent advances and applications of surrogate models for finite element method computations: a review, Soft Computing 26 (24) (2022) 13709–13733.
- [9] H. E. Fairclough, M. Gilbert, Layout optimization of long-span structures subject to self-weight and multiple load-cases, Structural and Multidisciplinary Optimization 65 (7) (2022) 197.
- [10] H.-T. Thai, Machine learning for structural engineering: A state-of-the-art review, in: Structures, Vol. 38, Elsevier, 2022, pp. 448–491.
- [11] X. Wang, B. Zhuang, D. Smyl, H. Zhou, M. Naser, Machine learning for design, optimization and assessment of steel-concrete composite structures: A review, Engineering Structures (2025).
- [12] K. A. James, J. S. Hansen, J. R. Martins, Structural topology optimization for multiple load cases using a dynamic aggregation technique, Engineering Optimization 41 (12) (2009) 1103–1118.
- [13] Y. Zhou, H. Zhan, W. Zhang, J. Zhu, J. Bai, Q. Wang, Y. Gu, A new data-driven topology optimization framework for structural optimization, Computers & Structures 239 (2020) 106310.
- [14] L. Zheng, S. Kumar, D. M. Kochmann, Data-driven topology optimization of spinodoid metamaterials with seamlessly tunable anisotropy, Computer Methods in Applied Mechanics and Engineering 383 (2021) 113894.
- [15] V.-N. Hoang, N.-L. Nguyen, D. Q. Tran, Q.-V. Vu, H. Nguyen-Xuan, Data-driven geometry-based topology optimization, Structural and Multidisciplinary Optimization 65 (2) (2022) 69.
- [16] H. Jeong, J. Bai, C. P. Batuwatta-Gamage, C. Rathnayaka, Y. Zhou, Y. Gu, A physics-informed neural network-based topology optimization (pinnto) framework for structural optimization, Engineering Structures 278 (2023) 115484.
- [17] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, Nature Reviews Physics 3 (6) (2021) 422–440.
- [18] A. Gallet, D. Smyl, Design of continuous structures using physics informed neural networks, EngrXivPreprint (2023). doi:<https://doi.org/10.31224/4242>.
URL <https://engrxiv.org/preprint/view/4242>
- [19] M. Zhu, F. McKenna, M. H. Scott, Openseespy: Python library for the opensees finite element framework, SoftwareX 7 (2018) 6–11.
- [20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, Advances in neural information processing systems 32 (2019).