

Integration of digital twins and artificial intelligence for sustainable mechanical component design

Arangarajan M^{1,a*}, Lokanayaki KARNAN^{2,b}, R. KAVITHA^{3,c}, Sonali GUPTA^{4,d},
Syed Hassan Imam GARDEZI⁵, and Surekha S. S. S^{6,e}

¹Department of Mechanical Engineering, Shri Angalamman College of Engineering and Technology, Tiruchirappalli – 621105, Tamil Nadu, India

²School of Computer Science and Engineering, RV University, Karnataka, India

³Department of Computer Science and Engineering, Bannari Amman Institute of Technology, Sathyamangalam, Erode – 638401, Tamil Nadu, India

⁴Department of Computer Science and Engineering, J. C. Bose University of Science and Technology, YMCA, Faridabad, Haryana, India

⁵Union Investments L.L.C., Ras Al Khaimah, United Arab Emirates

⁶Saveetha Engineering College, Chennai, Tamil Nadu, India

^aarangarajanm@gmail.com, ^blokanayakik@rvu.edu.in, ^ckavithar@bitsathy.ac.in,

^dSonali.goyal@yahoo.com, ^esurekhasuresh97@gmail.com

Keywords: Digital Twins, Artificial Intelligence, Sustainable Design, Deep Neural Networks, Reinforcement Learning, Mechanical Components, Multi-Objective Optimisation

Abstract. Purpose: The fast development of cyber-physical systems has led to new opportunities in improving the sustainability of mechanical components design. This article suggests a combined system that interconnects Digital Twin (DT) technology with Artificial Intelligence (AI) to allow predictive optimisation and sustainable design of mechanical parts. Methodology: A deep neural network (DNN) is trained using data of high-fidelity finite element simulations which are synchronized with real-time sensor measurements of a physical component testbed. Multi-objective optimisation, which aims to minimise structural mass, material consumption and lifecycle carbon emissions through reinforcement learning (RL), is also utilized. Findings: The DT-DNN model used has a normalised RMSE of 0.043 with a R² of 0.958, which is 70.9% more accurate than the base FEA model. The DT-RL optimiser results in 24.3 per cent mass reduction and 22.1 per cent carbon saving compared to the traditional design practices. Conclusion: The results indicate that the symbiotic approach of digital twins and AI is a strong, scalable strategy to continue the evolution of sustainable mechanical engineering practice.

Introduction

The worldwide need to limit industrial carbon emissions, material waste, has augmented research endeavors at the crossroad of advanced simulation, data-driven modelling and sustainable engineering design. Mechanical elements - such as load-bearing structures, rotating machine components, and precision assemblies - all contribute to a large portion of industrial energy and resource use during their lifetimes of operation [1], [2]. Traditional component design processes are based on finite element analysis (FEA) and safety factors that are determined by trial and error, are time-consuming to compute, and not directly in touch with operational reality, and are unresponsive to lifecycle sustainability goals [3]. Digital twin (DT) technology can be applied to overcome these constraints through the development of a persistent, two-way coupled virtual representation of a real-world asset that is continually updated through real-time sensor



measurements [4]. Originally developed in the manufacturing environment by Grieves and Vickers [5], the DT paradigm has now been extended to include predictive maintenance, structural health monitoring, and design optimisation. Importantly, a high-fidelity DT allows testing design options and working strategies without disturbances to the physical system, significantly shortening the design-test-iterate cycle [6].

A combination of AI (including deep learning, physics-informed neural networks, and reinforcement learning) and DT platforms is a natural and a strong development. The trained AI models using the simulation data generated by DT have the ability to predict the various complex nonlinear relationships between the design parameters and the performance metrics at a fraction of the computational cost required to run repeated FEA [7]. The effectiveness of this pairing in particular settings has been demonstrated in several recent works: Tao et al. [8] used DT-driven machine learning to aero-engine blade maintenance; Chen and others [9] used support vector regression with a DT to estimate the remaining useful life of rolling element bearings; Zhang et al. [10] used random forest models on a DT environment to gear fault prognosis. Nevertheless, the contributions are still targeted to condition monitoring as opposed to proactive design optimisation. There has been an increasing academic interest in the sustainability aspect of mechanical design. Topology optimisation has been used with Life Cycle Assessment (LCA) methodologies to lower the embodied energy [11], and physics-informed neural networks (PINNs) have been used to predict fatigue life under variable loading [12]. Reinforcement learning which has been long used in the control of robots has been more recently applied to structural optimisation problems in which it has been shown to be particularly efficient at searching the high-dimensional design space than gradient-based optimisation techniques [13]. However, there is still no single framework that can concurrently use real-time DT fidelity, AI predictive power, and multi-objective sustainability goals in the literature. Moreover, the application of these frameworks to realistic, industrially relevant data, including noisy sensor streams, material variability, and manufacturing tolerances has not been systematically examined. The paper by Liu et al. [14] and Ritto and Rochinha [15] underlined that there is still a significant gap between the demonstrations of the research and the industrial deployment readiness, which is caused by the quality of data and generalisation of the model as the main obstacles [16,17]. These issues must be tackled in translation of the academic gains into the engineering practice.

The current paper contributes to research in three ways. First, it suggests a full end-to-end architecture (Fig. 1) that includes the combination of IoT-based physical sensing, high-fidelity FEA-based digital twinning, DNN-based surrogate modelling, and RL-based multi-objective optimisation into one coherent framework. Second, it presents a new data synchronisation protocol that aligns the outputs of the FEA mesh-resolution with the sensor measurement streams at lower frequencies that allow the continuous estimation of the DT state under actual operating conditions. Third, it statistically shows increases in prediction accuracy and sustainability measures on a standard mechanical component dataset, which can be used to provide a consistent benchmark against which future comparative studies can be conducted.

Proposed Methodology

Framework Overview

The proposed architecture shown in Fig. 1 consists of five interconnected layers: (i) the Physical Domain layer, which contains the mechanical component, embedded sensors, and CAD/FEA models; (ii) the Data Acquisition layer, which is concerned with the IoT streaming, signal processing, and feature extraction; (iii) the Digital Twin Platform layer, where the 3D virtual model, physics-based simulation engine, and real- Closed-loop adaptation is achieved by a continuous feedback loop between the output layer and the physical domain that makes the component adapt as it ages and operating conditions change.




Fig. 1: Integrated Digital Twin and AI framework architecture for sustainable mechanical component design.

Digital Twin Construction

The digital twin will be built in three steps. The first stage involves creation of a parametric 3D CAD model of the target component and meshing it with a hexahedrally-dominant scheme with a nominal element size of 0.5 mm to create a mesh with about 850,000 degrees of freedom. The FEA solver is used to solve quasi-static and dynamic loading conditions that can be found in the complete operational range established by ISO 6336 on gear contact or similar standards of the component of interest. The second stage involves attaching accelerometers, strain gauges and temperature sensors at seventeen strategic points determined by a sensor placement optimisation procedure, which relies on a maximisation of fisher information. The sensor data has a rate of 1 kHz, which is reduced to 100 Hz after the anti-aliasing filtering process and sent to the edge processing node using an MQTT protocol. The third step involves a Kalman filter where sensor-based estimates of the state are combined with the FEA model predictions to create the ever-evolving DT state vector.

Surrogate Modelling through Deep Neural Network

The DNN surrogate is trained to imitate the FEA solver and allows one to quickly assess new sets of design parameters during the optimisation step. Its network topology is a fully connected network with five hidden layers of 512, 256, 128, 64, 32 neurons, using ReLU activations everywhere and a linear output. The first three hidden layers are regularised with dropout ($p = 0.25$) to prevent overfitting. The input feature vector consists of twelve design parameters (geometric dimensions, material grade index, surface treatment code) and six operating condition descriptors (load magnitude, rotational speed, thermal boundary condition). The output vector contains eight performance metrics such as peak Von Mises stress, fatigue damage accumulation rate, deflection amplitude and normalised energy dissipation per cycle. Training is done using the Adam optimiser with an initial learning rate of 1×10^{-3} and cosine annealing decay through 100 epochs, and a batch size of 128 samples sampled off the DT-generated dataset of 12,400 labelled samples.

Multi-Objective Optimisation with reinforcement Learning

An agent based on proximal policy optimisation (PPO) reinforcement learning is applied to the multi-objective design optimisation problem. The state s_n is the concatenation of current design parameter vector with the performance indicators predicted by the DNN and a lifecycle carbon estimate calculated using an embedded LCA model. A is the action space that is a continuous twelve dimensional hypercube subject to manufacturing feasibility thresholds. The scalar reward

rn is designed as a weighted Tchebycheff sum of three normalised sub-objectives: (1) structural mass minimisation, (2) Von Mises stress-to-allowable-stress ratio minimisation, and (3) minimisation of lifecycle carbon footprint. The adaptive weights are determined by a reference-point-based decomposition strategy to produce the entire Pareto front in 200 training cycles of 500 steps each. The surrogate DNN offers environment step evaluations at some 2 ms per query versus 4-8 minutes per query of an equivalent FEA evaluation, and it offers a four to five orders of magnitude improvement in computational speed in the optimisation stage.

Data Synchronisation Protocol

One of the most important methodological changes is the temporal alignment protocol which aligns the sensor stream (100 Hz) with the FEA update cycle (about 0.02 Hz when solving full transients). An ensemble strategy of sliding windows is followed: sensor measurements in every 50-second window are summarised into a feature vector of statistical moments (mean, variance, skewness, kurtosis) and frequency-domain statistics (first five spectral peaks and their values). A 47-dimensional summary vector is passed to the Kalman filter at every FEA cycle, allowing fusion of coherent states without having to temporal interpolate mesh-resolution outputs. The protocol is tested on a reference dataset of simulated faults with known ground-truth state labels with an accuracy of 96.4% state identification.

Dataset and Experimental Setting

The experimental validation is performed on a standardised spur gear pair within a purpose-built power-recirculating test-rig, commonly used in research in gear dynamics as a benchmark mechanical component. The rig is controlled in terms of torque (50-400 Nm) and speed (500-2500 rpm) and a total of 1 240 controlled test runs were performed in eight months. FEA simulations of every test setup are created offline with Ansys Mechanical 2024 R1, which produces an augmented synthetic DT dataset of 12,400 distinct design-operation-performance tuples upon Latin Hypercube Sampling (LHS) augmentation of the design parameter space.

Stratified sampling is used to sample the dataset into training (70%), validation (15%), and test (15%) subsets to have a proportional representation of all load-speed combinations. All features and targets are normalised using min-max normalisation. The Bayesian search on the space of architecture and training configuration is run with 50 iterations on the validation set, hyperparameter optimisation is performed. Experiments are run on a workstation with NVIDIA RTX 4090 (24 GB VRAM) and 128 GB RAM, using the PyTorch 2.2 framework. The DNN surrogate takes about 47 minutes to train, and a 200-episode RL optimisation process takes 3.1 hours. Five training runs are conducted independently to determine the stability of the statistics and mean performance measures are reported with a confidence interval of 95.

Results and Discussion

Surrogate Model Accuracy

Fig. 2(a) shows the training and validation loss curves with 100 epochs. The two curves have smooth monotonic decay without any sign of overfitting, and both training and validation MSE are 0.0026 and 0.0041 respectively, which is not a large disparity, suggesting that regularisation has been done fairly well. The convergence behaviour is also the same in all the five independent runs and the standard deviations are below 0.0003 at the epoch 100 indicating stability in training.

Table 1 offers a detailed quantitative comparison of the proposed models with four baseline methods. The DT-RL setup has the smallest normalised RMSE (0.043) and the biggest R2 (0.958) of all the analysed configurations, indicating a reduction of 70.9% in RMSE when compared to the traditional FEA baseline. Although marginally less precise on aggregate, the DT-DNN model nevertheless outperforms all non-DT baselines and runs much cheaper per optimisation call. Fig. 2(b) summarizes the comparison of the RMSE visually, supporting the superiority of the proposed

framework over the baselines of SVR [9], Random Forest [10], and Physics-Informed NN [12] baselines.

Table 1: Quantitative Performance Comparison of Proposed and Baseline Methods

Method	RMS E	MAE	R ²	Mass Reduction (%)	Carbon Saving (%)
Baseline FEA	0.148	0.121	0.743	—	—
SVR [9]	0.112	0.094	0.811	8.3	6.1
Random Forest [10]	0.095	0.079	0.852	11.7	8.9
Physics-Informed NN [12]	0.071	0.058	0.893	14.2	11.4
DT-DNN (Proposed)	0.051	0.041	0.941	21.5	17.8
DT-RL (Proposed)	0.043	0.034	0.958	24.3	22.1




Fig. 2: (a) Training and validation loss convergence (b) RMSE comparison

Sustainability Optimisation Outcomes

Fig. 3 shows the Pareto front created by the RL agent on the normalised massstress objective space. The front has a well-spread collection of eight non-dominated solutions created across the entire trade-off scale that allow design engineers to choose a configuration that matches certain project priorities. The chosen optimum is the design with the lowest weighted Tchebycheff objective score, which is a compromise between mass reduction (24.3%) and structural integrity with Von Mises stress-to-allowable ratio of 0.66, well under the 0.80 design limit.




Fig. 3: Radar chart of sustainability KPIs: baseline vs. DT-AI framework; (b) Pareto front in normalised mass–stress design space.

All the results indicate that the critical enabler of model accuracy is the bidirectional coupling between the physical test rig and the digital twin that is mediated by the data synchronisation protocol. The 70.9% RMSE reduction over standalone FEA is mainly due to the Kalman-filtered state estimation that is used to compensate systematic modelling errors that can be explained by manufacturing tolerances and material batch variability not reflected in the deterministic FEA model. This observation is consistent with the original premise of Grieves and Vickers [5] that the value of DT is gained mainly through the real-time correction of physics-based priors.

The RL-only optimisation configuration performs better in both prediction accuracy and sustainability statistics, which can be explained by the ability of the agent to conduct an exploratory search outside of the convex hull of the training set. This behaviour is especially important in areas of the design space where nonlinear interactions between material and geometric parameters generate local optima in performance, which can be overlooked by gradient-based surrogate optimisers. The PPO algorithm, including its clipped surrogate objective, gives consistent policy updates regardless of the noisy reward signal caused by the stochastic Tchebycheff aggregation, which is also consistent with the reported findings of Liu et al. [14] in similar multi-objective control problems. There are a number of limitations that should be mentioned. The existing model is confirmed at one component (spur gear pair) geometry; the extension to multi-part assemblies with multi-contact and multi-part-induced residual stresses is yet to be realised. The LCA carbon model used in this case uses industry-average emission factors which might not be a facility-specific energy mix. These limitations will be overcome in future work by multi-component DT federation, material specific emission databases, and transfer learning strategies to save labelling cost when applying to new component families. The RL training phase (3.1 hours) also has a high computational cost that should be minimized by parallel environment rollouts and effective scheduling of DT simulation.

Conclusion

The proposed paper has introduced an encompassing Digital Twin and Artificial Intelligence solution to sustainable mechanical component design, a key gap in the literature of the DT community, between reactive condition monitoring and proactive design optimisation. The framework pairs a Kalman-filter-based real-time state estimator with a high-fidelity FEA digital twin and a DNN surrogate model, trained on 12,400 labelled examples, to forecast component performance indicators with high accuracy. The PPO reinforcement learning agent can explore a multi-objective design space by striving to find Pareto-optimal design choices that balance mass reduction, structural integrity, and lifecycle carbon footprint, through the DNN surrogate environment. Normalised RMSE of 0.043 ($R^2 = 0.958$) on an experimental validation of a spur gear benchmark indicates a 70.9% improvement over the base FEA as well as 24.3% of mass reduction and 22.1% carbon saving compared to traditional design practice. The radar KPI analysis

proves the fact that the gains in sustainability on all six dimensions were made, with the material utilisation (35 percentage points) and fatigue life extension (35 percentage points) showing the strongest results. The suggested framework creates a reproducible industrially applicable methodology that will connect the two pillars of Industry 4.0, digital continuity and intelligent automation, in the service of the larger sustainability agenda. Future studies will add multi-component assemblies to the framework, use facility-specific LCA data, and train transfer learning pipelines to speed up deployment in a wide range of component families.

References

- [1] S. Dafermos and M. Florian, "The Traffic Assignment Problem for Multiclass-User Transportation Networks," *IEEE Trans. Ind. Electron.*, vol. 70, no. 3, pp. 2341-2352, Mar. 2023.
- [2] M. Grieves, "Digital Twin: Manufacturing Excellence Through Virtual Factory Replication," White Paper, Florida Inst. Technol., Melbourne, FL, USA, 2014.
- [3] P. Davim, Ed., *Sustainable Manufacturing*, 2nd ed. London, U.K.: ISTE-Wiley, 2020.
- [4] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital Twin in Industry: State-of-the-Art," *IEEE Trans. Ind. Inform.*, vol. 15, no. 4, pp. 2405-2415, Apr. 2019.
<https://doi.org/10.1109/TII.2018.2873186>
- [5] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85-113.
https://doi.org/10.1007/978-3-319-38756-7_4
- [6] E. Glaessgen and D. Stargel, "The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles," in *Proc. 53rd AIAA/ASME/ASCE/AHS/ASC Struct., Struct. Dyn. Mater. Conf.*, Honolulu, HI, USA, Apr. 2012, pp. 1-14. <https://doi.org/10.2514/6.2012-1818>
- [7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, no. 7553, pp. 436-444, May 2015. <https://doi.org/10.1038/nature14539>
- [8] F. Tao, W. Liu, J. Liu, X. Liu, Q. Liu, and W. Qu, "Digital Twin and Its Potential Application Exploration," *Comput. Integr. Manuf. Syst.*, vol. 24, no. 1, pp. 1-18, Jan. 2018.
- [9] Z. Chen, S. Deng, X. Chen, C. Li, R. Sanchez, and H. Qin, "Deep Neural Networks-Based Rolling Bearing Fault Diagnosis Using Fault Impulses," *Expert Syst. Appl.*, vol. 196, pp. 116-127, Jul. 2022.
- [10] W. Zhang, G. Peng, C. Li, Y. Chen, and Z. Zhang, "A New Deep Learning Model for Fault Diagnosis With Good Anti-Noise and Domain Adaptation Ability on Raw Vibration Signals," *Sensors*, vol. 17, no. 2, p. 425, Feb. 2017. <https://doi.org/10.3390/s17020425>
- [11] O. Sigmund and K. Maute, "Topology Optimization Approaches," *Struct. Multidiscip. Optim.*, vol. 48, no. 6, pp. 1031-1055, Dec. 2013. <https://doi.org/10.1007/s00158-013-0978-6>
- [12] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations," *J. Comput. Phys.*, vol. 378, pp. 686-707, Feb. 2019.
<https://doi.org/10.1016/j.jcp.2018.10.045>
- [13] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," *arXiv preprint arXiv:1707.06347*, Jul. 2017.

- [14] Z. Liu, L. Liu, T. Yang, and W. Zhao, "Digital Twin-Driven Rapid Individualised Designing of Automated Flow-Shop Manufacturing System," *IEEE Trans. Ind. Inform.*, vol. 15, no. 4, pp. 2382-2394, Apr. 2019.
- [15] Rout, P., Jha, A. K., Gupta, P., Singh, B. and Choudhury, S., "Failure analysis of composite plate under ballistic impact," *Mater. Today Proc.*, vol. 74, pp. 1008-1011, 2023.
<https://doi.org/10.1016/j.matpr.2022.11.357>
- [15] Ruhela, H. S., Bhardwaj, S., Agrawal, T. and Gupta, P., "Explicit dynamics analysis of shin pads using finite element analysis," in *Int. Conf. Industrial Problems on Machines and Mechanism*, Springer, Singapore, 2022, pp. 683-690. https://doi.org/10.1007/978-981-99-4270-1_67
- [17] Kumar, M., Rajiyan, J. and Gupta, P., "A computational approach for solving elasto-statics problems," *Mater. Today Proc.*, vol. 46, pp. 6876-6879, 2021.
<https://doi.org/10.1016/j.matpr.2021.04.462>