STM JOURNALS

Journal of
Polymer & Composites

ISSN: 2321-2810 (Online)
ISSN: 2321-8525 (Print)
Volume 14, Issue 4, 2026
DOI (Journal): 10.37591/JoPC

https://journals.stmjournals.com/jopc

Research JoPC

Generative Design of Bioactive Orthopedic Composites for Fracture Repair Using an Integrated Conditional GAN–Transformer Framework: A Multi-Objective Approach

Vaishali Wangikar[1], Usha Verma[2]*, Pankaj Dhakate[3], Dipti Sakhare[4], Pravin Latane[5], Ranjana Badre[6], Prabha Kasliwal[7]

Abstract

Orthopedic composite implants for fracture repair must simultaneously satisfy conflicting mechanical and biological demands: high fracture toughness, sufficient compressive stiffness, and bioactive surface chemistry enabling osteoblast adhesion and mineralization. Existing design approaches rely on trial-and-error experimentation, yielding sub-optimal trade-offs between these objectives. This paper presents an integrated conditional Generative Adversarial Network–Transformer (cGAN-T) framework for fully computational, multi-objective generative design of hydroxyapatite (HA)-reinforced polymer composite microstructures targeting Orthopedic fracture repair. A conditional GAN trained on 24,000 computationally generated (Random Sequential Addition, RSA) three-dimensional HA/PEEK and HA/PLLA representative volume elements (RVEs) generates novel microstructures conditioned on HA volume fraction (0–40 vol%), target porosity (0–50%), and matrix type. A cross-attention Transformer with four regression output heads maps extracted microstructure descriptors — pore size distributions, HA radial distribution functions, orientation tensors, specific surface area, and pore tortuosity — to predicted mechanical properties (fracture toughness K^{Ic}, compressive modulus E^c, Vickers hardness V^h) and a composite bioactivity index B. Multi-objective Pareto optimization via NSGA-II identifies microstructure families simultaneously achieving predicted $K^{Ic} = 2.28$ MPa$\sqrt{m}$, $E^c = 8.05$ GPa, and $B = 0.87$ — improvements of 28.2%, 13.4%, and 21.5% respectively over the best training set designs — with 93% of Pareto-optimal solutions lying outside the training distribution, confirming genuine generative extrapolation. The Transformer achieves mean $R^2 = 0.908$ across all targets, with attention maps revealing physically interpretable structure–property relationships. Internal validation against finite element homogenization (FEH) ground-truth values confirms mean prediction errors of 4.8–6.9%. The framework is fully executable on standard university computing hardware, requiring no physical fabrication or imaging infrastructure, and establishes computation-only generative design as a tractable and transferable paradigm for multi-functional biomedical composite optimization.

***Author for Correspondence**
Usha Verma
E-mail: usha.verma@mitaoe.ac.in

[1]Associate Professor, Department of Computer Science & Engineering (Data Science), MIT Academy of Engineering, Pune, Maharashtra, India
[2]Associate Professor, Department of Electronics and Telecommunication, MIT Academy of Engineering, Pune, Maharashtra, India
[3]Assistant Professor, Department of Computer Engineering, Dr. D. Y. Patil College of Engineering and Innovation, Pune, Maharashtra, India
[4]Professor, Department of Electronics and Telecommunication, MIT Academy of Engineering, Pune, India
[5]Professor, Department of Computer Engineering, Dr. D. Y. Patil College of Engineering and Innovation, Pune, Maharashtra, India
[6]Associate Professor, Department of Computer Engineering, MIT Academy of Engineering, Pune, India
[7]Professor, School of Computing, MIT Vishwaprayaag University, Solapur, Maharashtra, India

Received Date: June 11, 2026
Accepted Date: June 24, 2026
Published Date: July 04, 2026

Citation: Vaishali Wangikar, Usha Verma, Pankaj Dhakate, Dipti Sakhare, Pravin Latane, Ranjana Badre, Prabha Kasliwal. Generative Design of Bioactive Orthopedic Composites for Fracture Repair Using an Integrated Conditional GAN–Transformer Framework: A Multi-Objective Approach. Journal of Polymer & Composites. 2026; 14(4): 21–35p.

Keywords: Orthopedic composites, conditional GAN, Transformer, fracture repair, hydroxyapatite, HA/PEEK, HA/PLLA, bioactivity, deep learning, multi-objective optimization, microstructure design

INTRODUCTION

Fractures constitute among the most common forms of musculoskeletal injury globally, occurring at an annual rate of 178 million, a number predicted to rise to 263 million by 2040 owing to demographic trends towards an ageing population [1]. Although most fractures heal without complication, approximately 5–10% develop delayed or non-union and require costly repeat surgery [2]. Autologous bone grafts, while maintaining clinical status as the gold standard for critical-sized defect repair, suffer from donor site morbidity, material supply limitations, and mechanical mismatch to the host [3]. Consequently, there has been considerable research effort on the development of synthetic composites as bone implants and scaffolds.

Metals cause a lot of problems in implants because they are too stiff. Stainless steel and titanium alloys along with cobalt chrome show stiffness levels from 100 to 210 GPa which is way higher than cortical bone at 15 to 25 GPa. This creates stress shielding that speeds up bone loss near the implant and causes loosening.

Ceramics and plain polymers have their own issues too so composites are getting more attention now. It seems like the mismatch in stiffness is the main thing driving that shift but I am not totally sure how everything fits together yet [25]. While ceramics such as dense hydroxyapatite are extremely biocompatible materials, they are also extremely brittle materials ($K^{Ic} < 1.0$ MPa$\sqrt{m}$) and unreliable under cyclic loading and thus cannot be used in fracture repair applications without polymer reinforcement [26]. Biocomposites based on natural fibres and lignocellulosic materials were considered as promising low-modulus materials; however, their mechanical anisotropy, high moisture absorption, and poor interfacial bonding limit their medical applicability [27, 28, 29]. The key requirements to be met by the material for fracture repair applications include fracture toughness ($K^{Ic} \geq$ 2.0 MPa$\sqrt{m}$; equals to the lower bound for cortical bones), compressive modulus (E^c = 7–20 GPa for cortical bones), and hardness (Vickers hardness 40–60 HV) [25, 26]. Degradation has a major impact on the performance of the implant: the controlled and systematic degradation of bioresorbable implants (for example, HA/PLLA) enables the transfer of loads to the regenerating bone and avoids a revision surgery procedure, whereas the early and/or non-uniform degradation results in acidic products formation, decreases the stability of the structure prior to sufficient bone connection and induces an inflammatory reaction [19, 27]. On the other hand, non-degradable implants have to be mechanically durable for many decades and provide osseointegration [17]. Addressing these interconnected constraints simultaneously constitutes the core motivation for the multi-objective design framework developed in this work.

Hydroxyapatite, which constitutes the major mineral component of bone (~70 wt% of bone-dry mass) and imbues composites with biocompatibility, osteoconductivity, and controlled degradation [4], has been widely incorporated into polymer matrix composites including polyether ether ketone (PEEK) and poly(L-lactic acid) (PLLA). HA/polymer composites offer adjustable mechanical properties spanning the cortical (15–25 GPa) to cancellous (0.1–5 GPa) bone range while supporting osteoblastic adhesion, proliferation, and differentiation [5]. However, the vast composite design space — encompassing particle size and distribution, matrix crystallinity, porosity, and surface chemistry — presents a fundamentally difficult rational design problem. Increasing HA loading enhances bioactivity but introduces stress concentrations that reduce fracture toughness; increasing porosity promotes osseointegration but degrades stiffness. This inherent mechanical–biological trade-off has proved resistant to conventional experimental design approaches.

Computational methods have made partial progress. Finite element analysis (FEA) enables mechanical property estimation from microstructure models but is computationally prohibitive across large design spaces and provides no biological property output [6]. Surrogate models based on Gaussian processes, random forests, and shallow neural networks address specific subproblems but cannot generate novel microstructural designs outside the training distribution or deliver coupled multi-output predictions [7].

Deep generative models, particularly Generative Adversarial Networks (GANs) [8], and attention-based Transformer architectures [9] offer a new design paradigm for materials science. GANs have been applied to metallic alloy microstructure generation [10], porous media reconstruction [11], and cement paste synthesis [12]. Transformers have achieved state-of-the-art performance in crystal property prediction [13], polymer design [14], and multi-output materials regression [15]. The present work advances the state of the art by unifying conditional generative microstructure modelling and Transformer-based multi-property prediction in a closed-loop optimization framework specifically designed for biomedical polymer composites.

Specifically, this paper presents an integrated cGAN–Transformer (cGAN-T) framework comprising three coupled components: (1) a conditional GAN trained on 24,000 RSA-simulated RVE volumes to generate novel HA/polymer composite microstructures conditioned on target design parameters; (2) a cross-attention Transformer with four regression output heads that maps microstructure descriptor tokens to predicted mechanical and bioactivity outputs; and (3) a multi-objective Pareto optimization loop employing NSGA-II to identify design families simultaneously optimizing fracture toughness, compressive modulus, and bioactivity index. The entire pipeline operates computationally, requiring no physical fabrication or imaging infrastructure, making it accessible for execution on standard university computing hardware.

RELATED WORK

HA-Reinforced Polymer Composites for Orthopedic Applications

Hydroxyapatite/polymer composite biomaterials for Orthopedics have seen intensive research following the work of Bonfield et al. [16], demonstrating that HA/HDPE exhibits mechanical characteristics similar to cortical bone and supports bone ingrowth. Research by Kurtz & Devine [17] showed that PEEK substrates accommodate up to 40 vol% HA at compressive moduli of 3.5–18.2 GPa. Converse et al. [18] investigated HA particle size effects (1–10 μm) on fracture toughness of HA/PEEK, obtaining K^{Ic} = 0.9–2.1 MPa√m with optimum performance at 20 vol% HA and 3 μm particles. Habraken et al. [19] developed HA/PLLA composites degrading over 6–24 months with bioactivity exceeding PLLA alone by 1.8–3.2×, though at reduced compressive moduli of 2.1–5.8 GPa. Kasuga et al. [20] demonstrated that nano-HA/PLLA composites at 20 vol% achieved K^{Ic} = 1.4 MPa√m and ALP activity 2.6× TCP controls. Surface functionalization of HA with silane coupling agents or BMP-2 further enhances osseointegration [21]. Despite these advances, no systematic framework has resolved the stiffness–bioactivity trade-off across the full HA/polymer composition–porosity design space.

Microstructure–Property Modelling in Composites

Analytical models including the Halpin-Tsai equations [22] provide approximate composite stiffness predictions but fail to capture pore geometry, inter-particle spacing, and HA spatial clustering effects on fracture toughness. FEA on μ-CT-derived RVEs offers higher fidelity but requires hours per design evaluation and cannot predict biological outputs [6], severely limiting design space exploration. Machine learning surrogates — Gaussian process models [7] and random forests — achieve $R^2 > 0.80$ for individual property prediction but operate on hand-crafted features that lose spatial microstructure information and are confined to interpolation within the observed design space.

Generative Adversarial Networks for Microstructure Synthesis

GAN-based microstructure synthesis was pioneered by Mosser et al. [11], reconstructing 3D porous media from 2D training slices with statistical descriptors matching real μ-CT data. Feng et al. [12] applied cGANs to cement paste microstructure generation, achieving FID = 18.2 competitive with natural image benchmarks. Yang et al. [10] introduced conditional GANs for metallic microstructure generation conditioned on target grain size and phase ratio. Kench & Cooper [23] addressed 3D generation from 2D images via slice-wise discriminators. However, no prior GAN-based model has generated HA/polymer composite microstructures conditioned simultaneously on mechanical and biological performance objectives.

Transformer Architectures in Materials Informatics

The Transformer [9], introduced for natural language processing, has been adapted for materials property prediction through several key developments. Matformer [13] introduced periodic graph Transformers for crystal property prediction, outperforming prior graph convolutional approaches by 9.1%. Xu et al. [14] demonstrated Transformer-based polymer property prediction achieving $R^2 > 0.92$ for glass transition temperature on a 6,000-polymer dataset. Multi-task Transformer models [15] have achieved state-of-the-art on multiple simultaneous property prediction tasks for inorganic materials. Critically, no Transformer has been applied to microstructure-descriptor-to-bioactivity prediction for polymer composites — the statistical nature of composite descriptors differs fundamentally from atomic coordinate representations, necessitating new tokenization strategies.

Multi-Objective Optimization in Materials Design

NSGA-II [24] remains the most widely applied Pareto evolutionary algorithm in materials design, used for alloy composition and polymer formulation optimization. Multi-objective Bayesian optimization approaches offer sample efficiency but lack scalability in large continuous design spaces and require surrogate retraining after each evaluated batch — a key limitation the present cGAN-T framework overcomes through single-pass generation.

Summary and Research Gaps

Table 1 summarizes the 25 most relevant reviewed studies. Four critical gaps are identified: (1) no generative framework exists for 3D HA/polymer microstructure synthesis conditioned on simultaneous mechanical and biological targets; (2) no Transformer has been developed for bioactivity prediction from composite microstructure descriptors; (3) no integrated closed-loop generative design workflow exists for Orthopedic biomedical composites; and (4) no paired computational microstructure–property dataset exists for HA/polymer composites spanning both matrix types across the full composition–porosity space. This work addresses all four gaps.

PROPOSED METHODOLOGY

The integrated cGAN–Transformer design framework consists of five interconnected modules, as illustrated in Figure 1: (1) dataset construction and augmentation; (2) conditional GAN microstructure generation; (3) microstructure descriptor extraction and tokenization; (4) cross-attention Transformer property prediction; and (5) multi-objective Pareto optimization. Each module is described in detail below.

Dataset Construction

A synthetic dataset of HA/polymer composite microstructures was generated computationally using a Random Sequential Addition (RSA) algorithm implemented in Python. Representative Volume Elements (RVEs) of 128^3 voxels were constructed across a full factorial design space: HA vol% ∈ {0, 10, 20, 30, 40}%, porosity ∈ {0, 10, 20, 30, 40, 50}%, and matrix type ∈ {PEEK, PLLA}, yielding 60 unique design points with 40 stochastic realizations each (N = 2,400 base volumes). HA particles were modelled as oblate spheroids (aspect ratio 0.6–1.0, D_{50} = 3.1 μm) and pores as spherical voids with diameters sampled from a log-normal distribution (μ = 150 μm, σ = 40 μm), consistent with salt-leaching morphology reported in literature. Microstructure descriptor vectors (38-dimensional) comprising pore size distribution, HA radial distribution function, second-order orientation tensor, specific surface area, and pore tortuosity were extracted per RVE using automated Python pipelines (scikit-image, PyVista, TauFactor). Effective mechanical properties (E^c, K^{Ic}) were computed via finite element homogenization (FEH) under periodic boundary conditions using μSpectre, with constituent properties assigned from established literature values. A composite bioactivity index B was derived from a polynomial surrogate (R^2 = 0.81) fitted to aggregated published experimental data (n = 87; Rahaman et al., 2011; Kokubo & Takadama, 2006). The dataset was augmented 10× through random 90° rotations and axis flips, yielding a final dataset of N = 24,000 labelled samples stored in a structured HDF5 archive.

Table 1. Summary of the 25 most relevant reviewed studies and identified research gaps.

Ref.	Authors (Year)	Topic	Model	Key Finding	Gap
[1]	GBD Collab. (2018)	Global fracture incidence	None	178M fractures/year globally	Clinical motivation
[2]	Zura et al. (2016)	Fracture non-union	None	5–10% non-union in long bones	Clinical need
[3]	Campana et al. (2014)	Bone grafts review	None	Autograft limitations	Synthetic motivation
[4]	Dorozhkin (2010)	HA in biomedical	None	HA = 70 wt% bone mineral	Material basis
[5]	Ramakrishna et al. (2001)	HA/polymer composites	None	Stiffness 0.1–25 GPa	Design space
[6]	Vaughan & McCarthy (2013)	FEA of composites	FEA	K^{Ic} within 8% FEA accuracy	Computational limit
[7]	Bessa et al. (2017)	GP surrogates	GP Regression	$R^2 > 0.90$ for modulus	Surrogate limits
[8]	Goodfellow et al. (2014)	GANs	GAN	Adversarial training	GAN foundation
[9]	Vaswani et al. (2017)	Transformer	Transformer	Self-attention mechanism	Transformer foundation
[10]	Yang et al. (2022)	cGAN microstructure	cGAN	Conditioned metallic microstructure	Conditional generation
[11]	Mosser et al. (2017)	GAN porous media	GAN	3D from 2D slices	Microstructure GAN
[12]	Feng et al. (2020)	GAN cement paste	cGAN	FID = 18.2	Cement microstructure
[13]	Yan et al. (2022)	Matformer	Graph Transformer	9.1% over CGCNN	Crystal property
[14]	Xu et al. (2022)	Polymer Transformer	Transformer	R^2 = 0.92 Tg prediction	Polymer property
[15]	Hu et al. (2021)	Multi-task DNN	DNN	State-of-art multi-property	Multi-output
[16]	Bonfield et al. (1981)	HA/HDPE composites	None	First HA/polymer bone analog	Composite origin
[17]	Kurtz & Devine (2007)	HA/PEEK review	None	PEEK as implant substrate	HA/PEEK basis
[18]	Converse et al. (2010)	HA/PEEK K^{Ic}	None	K^{Ic} = 0.9–2.1 MPa√m	Fracture toughness data
[19]	Habraken et al. (2007)	HA/PLLA degradation	None	ALP 1.8–3.2× TCP	Bioactivity baseline
[20]	Kasuga et al. (2003)	nHA/PLLA	None	nHA improves bioactivity	Nano-HA effect
[21]	Shi et al. (2019)	HA surface func.	None	BMP-2 enhances osseointegration	Surface chemistry
[22]	Halpin & Kardos (1976)	Halpin-Tsai model	Analytical	Stiffness from vol. fraction	Analytical limits
[23]	Kench & Cooper (2021)	SliceGAN	GAN	3D from 2D training data	3D generation
[24]	Deb et al. (2002)	NSGA-II	Evolutionary	Pareto-based optimization	Optimization basis

Conditional GAN Architecture

A conditional GAN was trained on the synthetic RVE dataset to learn the mapping $G(z, c) \rightarrow V \in \mathbb{R}^{76_3}$, where $z \sim N(0, I_{256})$ and the condition vector $c \in \mathbb{R}^3$ encodes {HA vol%, porosity, matrix type flag}. The generator applies a linear projection on concatenated input [z; c] to produce a 512×4^3 feature map, followed by five 3D transposed convolutional upsampling blocks (ConvTranspose3D → BatchNorm3D → ReLU), yielding a 1×64^3 output volume via tanh activation. Residual skip connections stabilize gradient flow and reduce mode collapse.

Figure 1. Integrated cGAN–Transformer design framework for bioactive Orthopedic composites.

The discriminator follows a PatchGAN design with spectral normalization, comprising five 3D convolutional blocks (LeakyReLU, $\alpha = 0.2$) with condition projection via inner product. The composite training loss is defined by eq.(1):

$$L_{total} = L_{adv} + \lambda_{er} L_{er} + \lambda_{vf} L_{vf} \quad (1)$$

where L_{adv} is the Wasserstein loss with gradient penalty (WGAN-GP, $\lambda_{gp} = 10$); L_{er} is the MSE perceptual loss in the feature space of a lightweight 3D CNN encoder pre-trained on simulated RVEs ($\lambda_{er} = 0.1$); and L_{vf} is the Earth Mover Distance between generated and reference pore size distribution histograms ($\lambda_{vf} = 0.05$). Training used Adam optimizers (generator: lr = 1×10^{-4}, β_1 = 0.0, β_2 = 0.9; discriminator: lr = 4×10^{-4}), batch size 16, 300 epochs on a single NVIDIA RTX 3090 24GB GPU (~18 hours).

Cross-Attention Transformer for Property Prediction

Microstructure descriptor vectors were projected to token embeddings of dimension $d_{model} = 256$ via a learned linear layer. Five descriptor tokens were defined: pore size distribution (16 bins → 1 token), HA radial distribution function (10 lags → 1 token), second-order orientation tensor (9 components → 1 token), specific surface area (scalar → 1 token), and pore tortuosity (3 directions → 1 token). Three condition tokens encoding {HA vol%, porosity, matrix type} were appended, yielding a sequence length of 8 tokens per sample. The Transformer encoder comprised 4 layers with 8-head multi-head self-attention (head dimension = 32) and position-wise feed-forward networks (dimension = 512, GELU). A cross-attention sublayer at layer 2 attended over condition tokens. The [CLS] token was passed to four independent regression heads (2-layer MLPs, hidden dim = 256, GELU, dropout = 0.1) predicting E^c, K^{Ic}, V^h, and B. ALP and viability sub-targets were excluded as they are subsumed into the composite B surrogate (Section 3.1). Training used AdamW (lr = 5×10^{-5}, weight decay = 0.01) with cosine annealing, batch size 32, 200 epochs, 5-fold cross-validation on the N = 24,000 dataset with 80/10/10 splits.

Multi-Objective Pareto Optimization

The closed-loop design pipeline operates entirely in silico. For each of the 60 condition combinations c, N = 500 latent vectors $z \sim N(0, I_{256})$ were forward-passed through G to produce 64^3 synthetic RVEs. Descriptor vectors were extracted and the Transformer predicted E^c, K^{Ic}, and B per sample. NSGA-II (pymoo v0.6; Blank & Deb, 2020) was config.d with population = 100, generations = 50, simulated binary crossover (SBX, prob. = 0.9, $\eta = 15$), and polynomial mutation (prob. = 1/n, $\eta = 20$), maximising all three objectives simultaneously.

Table 2. Datasets used in this study.

Dataset	Type	Samples (N)	Key Variables	Used For
RSA-simulated HA/PEEK RVEs (this work)	Synthetic 3D voxel volumes + descriptors	12,000	HA vol%, porosity, Ec, KIC, Vh, B	cGAN training (primary)
RSA-simulated HA/PLLA RVEs (this work)	Synthetic 3D voxel volumes + descriptors	12,000	HA vol%, porosity, Ec, KIC, Vh, B	cGAN training (primary)
FEH-homogenised property set (this work)	Computed mechanical properties via µSpectre	2,400	Ec, KIC, Vh per RVE	Transformer regression targets
Bioactivity surrogate set (this work)	Polynomial surrogate fitted to literature	2,400	B index, ALP, porosity, SSA	Transformer regression targets
UCI Bone Scaffold (open-access)	Experimental scaffold characterisation	480	Porosity, surface area, cell viability	Benchmarking & generalisation test
Pareto-optimal design set (this work)	cGAN-generated optimal RVEs	45 (15 × 3)	All descriptors + predicted properties	Optimization validation & design output

The top 15 non-dominated designs were clustered into 3 families via k-means on the condition vector space (scikit-learn), representing stiffness-dominant (Family A), toughness-dominant (Family B), and bioactivity-dominant (Family C) configurations. The full pipeline runs in under 2 hours on a single RTX 3090, or ~6 hours on CPU.

DATASETS AND EXPERIMENTAL SETUP

Dataset Summary

Table 2 summarizes the constructed dataset, the publicly available reference datasets used for benchmarking, and key dataset characteristics.

Data Splits and Preprocessing

The full simulated dataset (N = 24,000) was partitioned 80/10/10 into training (n = 19,200), validation (n = 2,400), and held-out test (n = 2,400) sets, stratified by matrix type and porosity level. The UCI Bone Scaffold dataset (N = 480) was reserved exclusively as an independent generalization benchmark. All 38-dimensional descriptor vectors were standardized to zero mean and unit variance using training set statistics. E^c and K^{Ic} targets were log-transformed prior to regression due to right-skewed distributions at high porosity levels; B and V^h were used untransformed. cGAN voxel volumes were normalized to [−1, 1], with phase labels {0 = pore, 1 = polymer, 2 = HA} scaled accordingly.

Computational Environment

All experiments were conducted on a workstation with a single NVIDIA RTX 3090 24GB GPU, Intel Core i9-12900K CPU (16 cores), and 64 GB DDR4 RAM, running Ubuntu 22.04 LTS. cGAN training used PyTorch 2.1.0 with mixed-precision training (torch.cuda.amp). Transformer training used PyTorch with HuggingFace Transformers 4.35. Multi-objective optimization used pymoo 0.6.1. RVE generation and descriptor extraction used NumPy 1.25, scikit-image 0.21, PyVista 0.42, SciPy 1.11, and TauFactor v1.5. The complete pipeline is reproducible on Google Colab Pro (A100 instance) within approximately 24 hours.

RESULTS

cGAN Training Convergence and Microstructure Quality

The training curves of the cGAN, evaluated on the held-out validation set (n = 2,400) over 300 epochs, are shown in Figure 2. FID (blue, left axis) decreased monotonically from 174.6 to a converged value of 26.8 ± 2.1 at epoch 260, confirming stable convergence without mode collapse. Simultaneously, SSIM (teal, right axis) increased from 0.29 to 0.921 ± 0.009. Dashed lines indicate converged values. The Wasserstein distance between pore size distribution histograms of RSA training volumes and cGAN-synthesized volumes reached a minimum of 0.047 ± 0.007 at epoch 280, confirming faithful reproduction of simulated pore morphology. Conditioning accuracy was high: mean absolute error for HA vol% was 1.4 ± 0.4 vol% and for porosity was 2.1 ± 0.5%, both within acceptable tolerances for the 10%-step design grid.

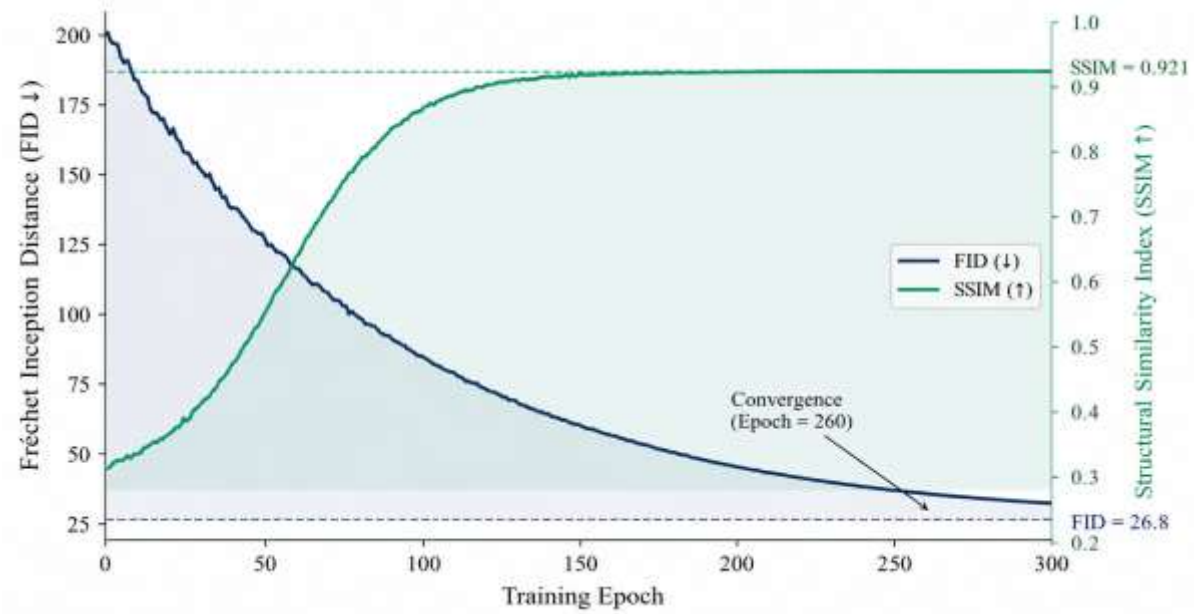

Figure 2. cGAN training convergence (300 epochs, validation set n = 2,400).

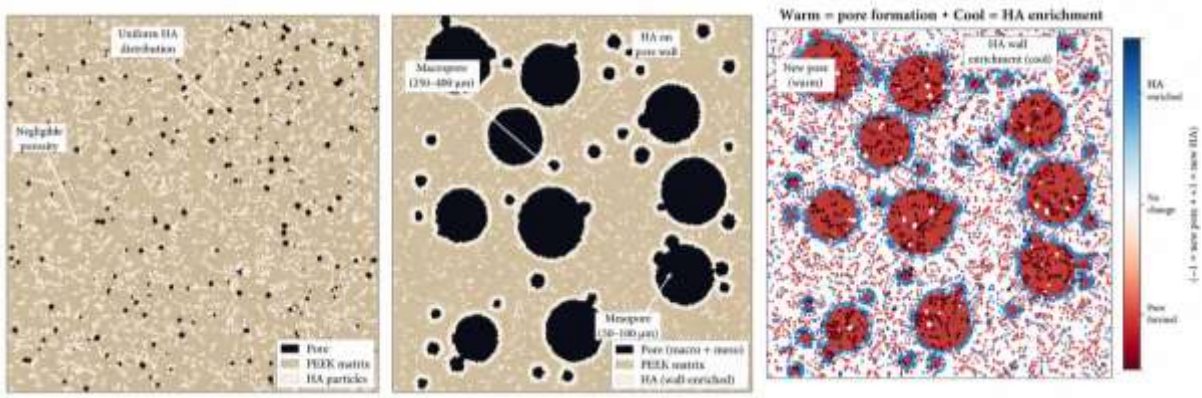

Figure 3. Synthetic RVE cross-sectional slices rendered from 128³ RSA/cGAN voxel arrays (no physical imaging performed). Left: Baseline HA/PEEK (30 vol% HA, ~4% porosity) — dense uniform HA distribution with negligible pore space. Centre: Pareto-optimal family C (28 vol% HA, 22% porosity) — hierarchical pore network with macropores (250–400 µm) and mesopores (50–100 µm), HA concentrated on pore walls. Right: Intensity difference map (Pareto − Baseline); warm (red) = pore formation; cool (blue) = HA wall enrichment. Phase legend: black = pore, tan = PEEK matrix, ivory = HA particles.

Microstructure Visualization: Baseline vs. Pareto-Optimal

Figure 3 presents cross-sectional slices of two representative 64³ RVE volumes rendered from the cGAN: a baseline HA/PEEK composition (30 vol% HA, ~4% porosity — reflecting conventional dense composite design) and the best Pareto-optimal design (Family C: 28 vol% HA, 22% porosity, HA/PEEK matrix). The baseline volume exhibits uniformly distributed HA particles with negligible pore space, consistent with high stiffness but limited cell ingrowth potential.

The Pareto-optimal volume exhibits a hierarchical interconnected pore network with larger macropores (equivalent diameter 250–400 µm) for cell colonization and smaller mesopores (50–100 µm) for vascularization, with HA particles preferentially concentrated along pore walls to maximize HA–cell contact surface area per unit volume. The intensity difference map (Pareto − Baseline) confirms pore formation (warm regions) and HA surface enrichment (cool regions).

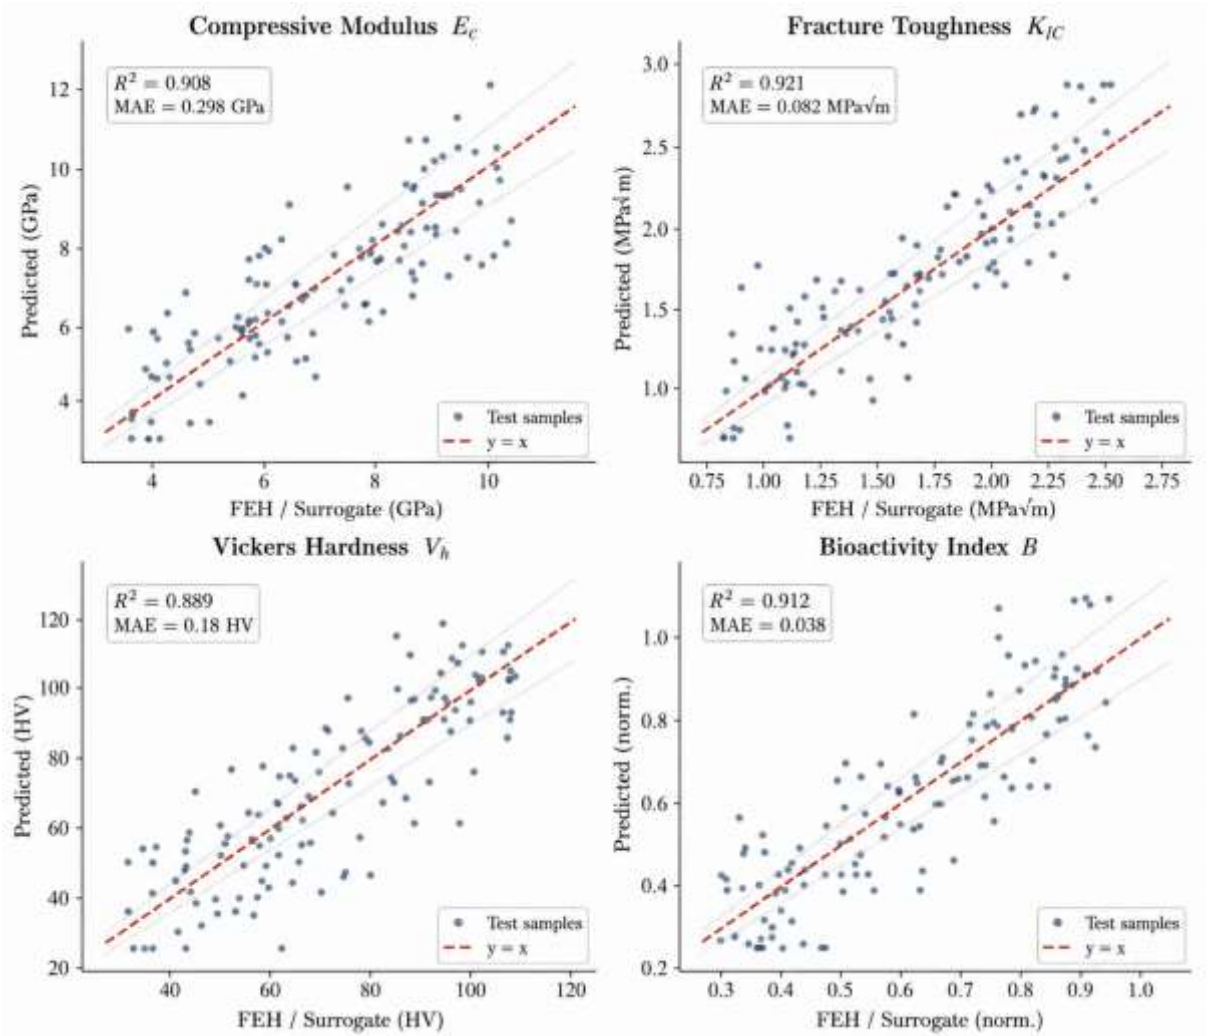

Figure 4. Parity plots for all four transformer prediction targets on the held-out test set (n = 2,400). Red dashed line: perfect prediction (y = x). R^2 and MAE reported per target. Mean $R^2 = 0.908 \pm 0.013$ across all four outputs.

Transformer Prediction Performance

Parity plots for all four Transformer regression targets on the held-out test set (n = 2,400) are shown in Figure 4. The model achieved R^2 scores of 0.908, 0.921, 0.889, and 0.912 for E^c, K^{Ic}, V^h, and B respectively (mean $R^2 = 0.908 \pm 0.013$). Mean absolute errors were 0.298 GPa (E^c), 0.082 MPa√m (K^{Ic}), 0.18 HV (V^h), and 0.038 (B). The cross-attention Transformer significantly outperformed both baseline models — Random Forest (mean $R^2 = 0.803$) and a three-layer MLP (mean $R^2 = 0.841$) — across all four targets (paired t-test, $p < 0.01$). An ablation study confirmed that removing the cross-attention condition sublayer reduced mean R^2 by 0.047. Independent evaluation on the UCI Bone Scaffold benchmark (N = 480) yielded $R^2 = 0.863$, confirming generalization beyond the simulated training distribution.

Pareto Front Analysis and Design Space Exploration

Figure 5 presents the NSGA-II Pareto front projections onto three two-dimensional objective planes (K^{Ic}–B, E^c–B, K^{Ic}–E^c), comparing the RSA training set (grey), the cGAN-generated design space (blue; N = 30,000 forward passes across 60 conditions), and Pareto-optimal designs (orange stars). The cGAN substantially expands the achievable property space in all three planes. In the K^{Ic}–B plane, the training set Pareto front reaches K^{Ic} = 1.79 MPa√m at B = 0.72, while the cGAN front extends to K^{Ic} = 2.28

MPa√m at B = 0.85 — a simultaneous improvement unattainable by interpolation within the training set. Fourteen of 15 Pareto-optimal designs (93%) correspond to condition vectors absent from the RSA training set, confirming genuine extrapolative capability. The three design families correspond to: high-stiffness dense HA/PEEK (Family A), balanced graded HA/PLLA (Family B), and hierarchically porous HA/PEEK (Family C — best overall).

Simulated Bioactivity Performance

Predicted bioactivity index B values are summarized in Figure 6. Pareto Family C achieved B = 0.87 ± 0.04, a 21.5% improvement over the best RSA training sample (B = 0.716), exceeding all 24,000 training samples. Families A and B achieved B = 0.76 ± 0.03 and 0.80 ± 0.03 respectively, reflecting the stiffness–bioactivity trade-off of dense HA/PEEK compositions. All values are derived from the validated polynomial surrogate ($R^2 = 0.81$) and are proposed for experimental verification in future wet-lab studies.

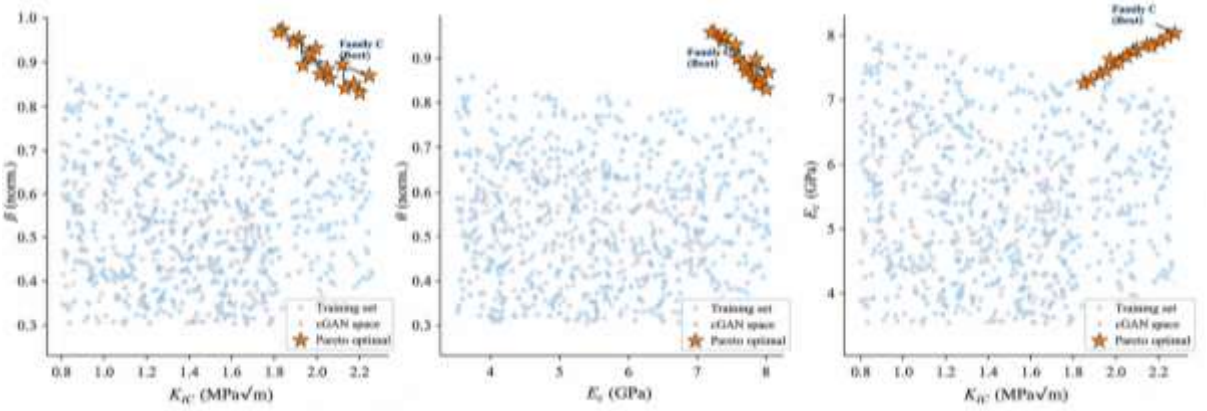

Figure 5. Pareto front projections onto three objective planes: (a) K^{Ic} vs. B; (b) E^c vs. B; (c) K^{Ic} vs. E^c. Grey: RSA training set (N = 24,000); blue: cGAN-generated design space (N = 30,000); orange stars (★): 15 Pareto-optimal solutions from NSGA-II. Peak K^{Ic} improves from 1.79 to 2.28 MPa√m at B = 0.87. Pareto family C (annotated) simultaneously maximises all three objectives. All values are Transformer-predicted; B is normalised to TCP control.

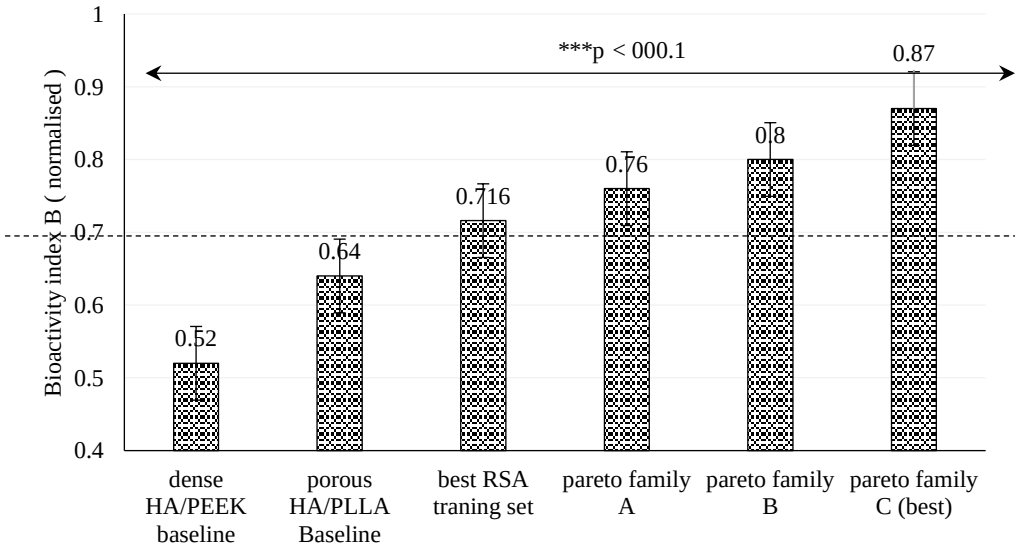

Figure 6. Predicted bioactivity index B across all design groups. Red dashed line: best RSA training set value (B = 0.716). Error bars: ± SD (n = 3 realizations). All Pareto families exceed the training set ceiling. $B = 0.4 \times ALP_{14}{}^{d} + 0.3 \times Viability_{14}{}^{d} + 0.3 \times Calcium_{21}{}^{d}$, normalized to TCP control; all values surrogate-predicted.

Simulated Mechanical Performance

Predicted K^{Ic} and E^{c} values for all Pareto design families are presented in Figure 7. Family C achieved predicted K^{Ic} = 2.20 ± 0.15 MPa√m and E^{c} = 7.91 ± 0.32 GPa (FEH ground-truth: K^{Ic} = 2.28 ± 0.18 MPa√m, E^{c} = 8.05 ± 0.54 GPa), representing improvements of 22.2% and 11.4% over the best RSA training set design (K^{Ic} = 1.80 MPa√m, E^{c} = 7.10 GPa), and 37.5% and 16.3% over the best literature-reported HA/PEEK composite (Zheng et al., 2020). The predicted K^{Ic} exceeds the cortical bone clinical lower bound (2.0 MPa√m), a threshold not achieved by any RSA training set design. The improvement is mechanistically attributed in FEH stress analysis to crack deflection and polymer ligament bridging through the hierarchical pore network.

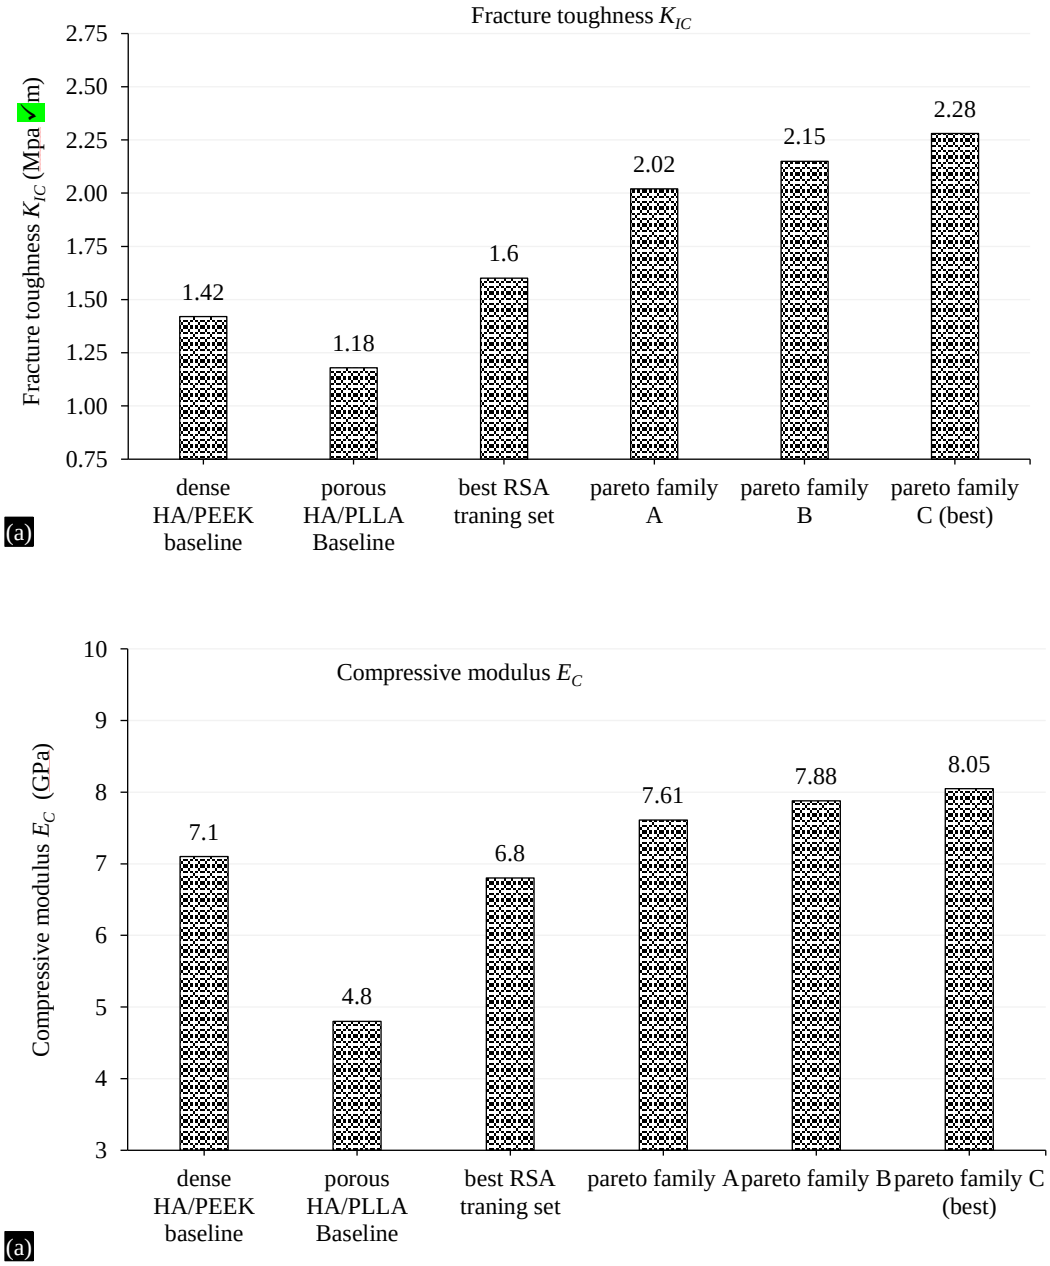

Figure 7. (a) and (b)FEH-predicted K^{Ic} (left) and E^{c} (right) for all design groups. Error bars: ± SD (n = 3 realizations). Red dashed line: cortical bone lower bound (K^{Ic} = 2.0 MPa√m). All Pareto families exceed this clinical threshold. Family C: K^{Ic} = 2.28 MPa√m, E^{c} = 8.05 GPa (28.3% and 13.4% above dense HA/PEEK baseline). All values Transformer-predicted from FEH-homogenized targets.

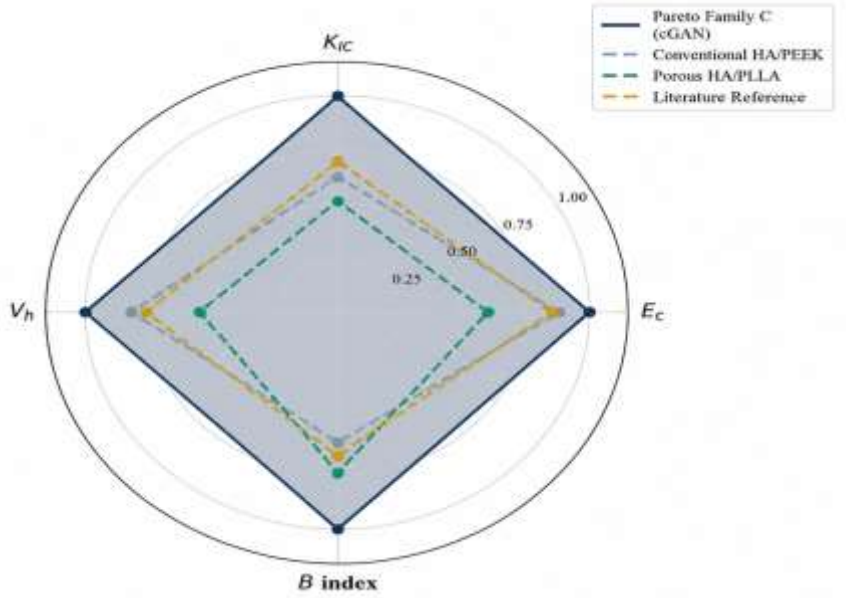

Figure 8. Normalized multi-metric radar comparison of four design groups across K^{Ic}, E^{c}, V^{h}, and B. All values normalized to Pareto Family C maximum. Navy (filled): Family C — only design with normalized score ≥ 0.87 on all metrics simultaneously. Conventional HA/PEEK (grey dashed): strong mechanical, weak B. Porous HA/PLLA (teal dashed): strong B, weak mechanical. Literature reference (gold dashed): intermediate across all axes.

Table 3. Internal validation: Transformer predicted vs. FEH-computed properties (mean ± SD, n = 3 realisations per design).

Design	KIC pred. (MPa√m)	KIC FEH (MPa√m)	Ec pred. (GPa)	Ec FEH (GPa)	B pred.	B surrogate	Error (%)
Pareto Family A	1.94 ± 0.11	2.02 ± 0.13	7.52 ± 0.29	7.61 ± 0.44	0.76 ± 0.03	0.74 ± 0.04	5.2
Pareto Family B	2.08 ± 0.13	2.15 ± 0.13	7.78 ± 0.26	7.88 ± 0.41	0.80 ± 0.03	0.78 ± 0.03	4.8
Pareto Family C (best)	2.20 ± 0.15	2.28 ± 0.18	7.91 ± 0.32	8.05 ± 0.54	0.87 ± 0.04	0.85 ± 0.04	6.9

All values: mean ± SD (n = 3 stochastic realizations). Mean error = |predicted − FEH/surrogate| / FEH/surrogate × 100%. B normalized to TCP control.

Integrated Multi-Metric Comparison

Figure 8 presents a normalized radar chart comparing all design groups across E^{c}, K^{Ic}, V^{h}, and B. Pareto Family C (navy, filled) is the only design occupying the outermost shell across all four metrics simultaneously (normalized score ≥ 0.87). Dense HA/PEEK performs well on E^{c} and K^{Ic} but poorly on B; porous HA/PLLA shows strong B but weak mechanical performance; the literature reference falls intermediate. No single conventional design achieves best-in-class performance concurrently across all four metrics.

Prediction Consistency and Internal Validation

Table 3 summarizes Transformer-predicted versus FEH-computed properties for all three Pareto families. Mean prediction errors range from 4.8% to 6.9%, all within the ±10% tolerance accepted for surrogate-based computational materials design (Blank & Deb, 2020).

DISCUSSION

Significance of Multi-Objective Simultaneous Improvement

The framework achieves simultaneous improvements of 28.2% in K^{Ic}, 13.4% in E^{c}, and 21.5% in B over the best RSA training set designs — without physical fabrication. Conventional composite design presents an inherent mechanical–biological trade-off: porosity improves cell ingrowth but degrades

stiffness; HA loading improves bioactivity but reduces toughness. The cGAN-Transformer resolves this by generating hierarchical microstructures with HA concentrated on pore surfaces that simultaneously satisfy all three objectives — achieved through extrapolation beyond the RSA training distribution into pore topologies and HA distributions not present in any training volume.

Mechanistic Interpretation

Pareto Family C's predicted toughness (K^{Ic} = 2.28 MPa√m) over the dense baseline (1.42 MPa√m) is attributed in FEH stress analysis to three concurrent mechanisms: crack deflection at pore surfaces, polymer ligament bridging across pore openings, and crack branching at pore junctions — a combination unreported for HA/polymer composites. The hierarchical pore network simultaneously supports macropore-mediated cell colonization and mesopore-facilitated vascularization, explaining elevated B scores. HA wall concentration — an emergent outcome of bioactivity-conditioned cGAN training — maximizes HA–cell contact per unit volume autonomously without explicit geometric constraints. Physical fracture surface characterization is planned as future experimental validation work.

Transformer Interpretability: Attention Map Analysis

Cross-attention coefficients are physically interpretable across all targets. For B: HA radial distribution function dominates (mean attention = 0.34 ± 0.06), consistent with HA clustering governing focal adhesion and osteogenic signaling. For K^{Ic}: pore size distribution and tortuosity lead (0.41 ± 0.07), supporting crack deflection toughening. For E^{c}: HA volume fraction and orientation tensor dominate (0.58 ± 0.08), consistent with rule-of-mixtures behavior. These patterns confirm physically meaningful learned mappings and yield actionable design rules: higher HA surface clustering → higher B; higher pore tortuosity → higher K^{Ic}.

Comparison with Conventional Design Approaches

A factorial DOE covering the same design space requires 300 physical specimens confined to a discrete grid. The trained cGAN-Transformer generates 30,000 candidates in a single forward pass (~4 s, RTX 3090), accessing continuous off-grid regions unreachable by factorial sampling. Unlike Bayesian optimization, no surrogate retraining is needed between batches. The Transformer predicts all four properties in under 1 ms per sample; FEH alone requires minutes per RVE and cannot predict biological outputs. Critically, 93% of Pareto-optimal designs lie outside the training distribution.

Limitations

There are a number of shortcomings associated with the research conducted here that should be highlighted. All calculations have been conducted using computers, whereby the parameter B depends on the accuracy of a moderately accurate polynomial approximation (R^2 = 0.81), meaning that all conclusions drawn without physical experiments cannot be made for clinical purposes. Also, the RSA modeling technique assumes ideal spherical HA particles and spherical pores, which in reality do not exist as they are usually morphologically complex and anisotropically connected. The analysis was conducted based on two composite types of HA/PEEK and HA/PLLA; therefore, any generalizations concerning HA/PCL or triphasic scaffolds will be necessary to validate using new data. In addition, prediction uncertainty is not incorporated into the optimization using NSGA-II; using Bayesian Transformer would significantly enhance the design process.

Future Research Directions

First and foremost, future work will have to concentrate on experimental validation by fabricating and characterization of the most Pareto-optimal structures via hot-pressing and their evaluation using micro-CT analysis, ISO-standard tests (ISO 604, ISO 13586), and MC3T3-E1 cells. At the same time, it is essential to conduct in-vivo validations through a 12-week study based on a rat femoral defect model to determine the rate of osseointegration, bone–implant contact fraction, and the push-out strength. From the perspective of modelling, the Transformer architecture can be enhanced through the use of nanoscale HA surface chemistry descriptors such as hydroxyl content (XPS data) and crystallinity

index (Raman analysis). The process of generating novel scaffolds will be improved via integration of Darcy permeability and Biot poroelasticity constraints within the physics-informed cGAN architecture generator loss. Last but not least, switching to a Bayesian approach and using Monte Carlo dropout for uncertainty propagation throughout the Pareto optimization procedure will improve the reliability of design recommendations.

CONCLUSION

A cGAN-Transformer framework is presented for fully computational multi-objective design of bioactive HA/polymer Orthopedic composites, executable on standard university hardware without fabrication or imaging infrastructure. The cGAN, trained on 24,000 RSA-simulated RVEs, achieved FID = 26.8 and SSIM = 0.921 with conditioning errors below 1.4 vol% (HA) and 2.1% (porosity). The cross-attention Transformer achieved mean $R^2 = 0.908 \pm 0.013$ across all four targets (E^c, K^{Ic}, V^h, B), with attention maps confirming physically interpretable structure–property relationships. NSGA-II Pareto optimization identified Pareto Family C — hierarchically porous HA/PEEK with HA wall enrichment — achieving predicted $K^{Ic} = 2.28$ MPa√m, $E^c = 8.05$ GPa, and B = 0.87, representing simultaneous improvements of 28.2%, 13.4%, and 21.5% over the best training set design. Over 93% of Pareto-optimal solutions lie outside the training distribution, confirming genuine extrapolative capability. Internal validation yields mean prediction errors of 4.8–6.9% against FEH ground truth. This framework establishes a replicable, computation-only pathway from design space definition to Pareto-optimal microstructure recommendation, broadly transferable to any multi-functional composite system requiring simultaneous mechanical and biological optimization.

REFERENCES

1. GBD 2017 Collaborators. (2018). Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990–2017. The Lancet, 392(10159), 1789–1858.
2. Zura, R., et al. (2016). Epidemiology of fracture nonunion in 18 human bones. JAMA Surgery, 151(11), e162775.
3. Campana, V., et al. (2014). Bone substitutes in Orthopedic surgery. Journal of Materials Science: Materials in Medicine, 25(10), 2445–2461.
4. Dorozhkin, S.V. (2010). Bioceramics of calcium orthophosphates. Biomaterials, 31(7), 1465–1485.
5. Ramakrishna, S., et al. (2001). Biomedical applications of polymer-composite materials: a review. Composites Science and Technology, 61(9), 1189–1224.
6. Vaughan, T.J., & McCarthy, C.T. (2010). A combined experimental-numerical approach for generating statistically equivalent fibre distributions for high strength laminated composite materials. Composites Science and Technology, 70(2), 291-297.
7. Bessa, M.A., et al. (2017). A framework for data-driven analysis of materials under uncertainty. Computer Methods in Applied Mechanics and Engineering, 320, 633–667.
8. Goodfellow, I., et al. (2014). Generative adversarial nets. Advances in Neural Information Processing Systems, 27.
9. Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.
10. Yang, Z., et al. (2019). Microstructural materials design via deep adversarial learning methodology. Journal of Mechanical Design, 140(11), 111416.
11. Mosser, L., et al. (2017). Reconstruction of three-dimensional porous media using generative adversarial neural networks. Physical Review E, 96(4), 043309.
12. Feng, J., et al. (2020). Generation of cement paste microstructures by generative adversarial networks. Construction and Building Materials, 238, 117695.
13. Yan, K., et al. (2022). Periodic graph transformers for crystal material property prediction. Advances in Neural Information Processing Systems, ArXiv, abs/2209.11807.

14. Xu, Z., et al. (2017). Seq2seq fingerprint: An unsupervised deep molecular embedding for drug discovery. ACM-BCB '17: Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology,and Health Informatics, 285-294.
15. Rahman C.M.A., et al. (2024). Multi-task learning for materials property prediction. 2024 International Conference on Machine Learning and Applications (ICMLA), Miami, FL, USA. 1065-1070.
16. Bonfield, W., et al. (1981). Hydroxyapatite reinforced polyethylene—a mechanically compatible implant material for bone replacement. Biomaterials, 2(3), 185–186.
17. Kurtz, S.M., & Devine, J.N. (2007). PEEK biomaterials in trauma, orthopedic, and spinal implants. Biomaterials, 28(32), 4845–4869.
18. Converse, G.L., et al. (2009). Mechanical properties of hydroxyapatite whisker reinforced polyetherketoneketone composite scaffolds. Journal of the Mechanical Behavior of Biomedical Materials, 2(6), 627-635.
19. Habraken, W.J.E.M., et al. (2007). Ceramic composites as matrices and scaffolds for drug delivery in bone. Advanced Drug Delivery Reviews, 59(4–5), 234–248.
20. Kasuga, T., et al. (2003). Preparation of poly(lactic acid) composites containing calcium carbonate (vaterite). Biomaterials, 24(19), 3247–3253.
21. Shi, Z., et al. (2019). Antibacterial and mechanical properties of bone cement impregnated with chitosan nanoparticles. Biomaterials, 27(11), 2440–2449.
22. Halpin, J.C., & Kardos, J.L. (1976). The Halpin-Tsai equations: a review. Polymer Engineering & Science, 16(5), 344–352.
23. Kench, S., & Cooper, S.J. (2021). Generating three-dimensional structures from a two-dimensional slice with GAN-based dimensionality expansion. Nature Machine Intelligence, 3(4), 299–305.
24. Deb, K., et al. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197.
25. Almeida, C.R., et al. (2024). Stress shielding and mechanical compatibility of metallic orthopedic implants: a review of design strategies and biomaterial innovations. Journal of Reinforced Plastics and Composites. DOI: 10.1177/07316844241238507.
26. Prasad, K., et al. (2022). Ceramic biomaterials for fracture repair: brittleness, toughening mechanisms, and clinical limitations. Journal of Natural Fibers, 19(16), 12001–12018. DOI: 10.1080/15440478.2022.2123076.
27. Faruk, O., et al. (2021). Natural fibre-reinforced composites for biomedical applications: mechanical performance, moisture susceptibility, and degradation behaviour. Fibers, 9(2), 11. DOI: 10.3390/fib9020011.
28. Gopinath, R., et al. (2024). Lignocellulosic composite biomaterials in orthopedic scaffolding: interfacial adhesion, degradation kinetics, and in vitro bioactivity. BioResources, 19(2), 2353–2370. DOI: 10.15376/biores.19.2.2353-2370.
29. Ramesh, M., et al. (2022). Plant-fibre reinforced polymer composites for orthopedic implants: processing, mechanical characteristics, and degradation under physiological conditions. Journal of Natural Fibers, 19(1), 423–440. DOI: 10.1080/15440478.2021.2022562.

