



PINN-XAF-MTL: A Multi-Modal Deep Learning Framework Integrating Novel Physics-Informed Cross-Attention Fusion for Predicting Multi-Environment Hybrid Performance

Chirag Radadiya*, Manish Sapovadiya, Ankur Hadiya, Vishal A. Savaliya

Department of Genetics and Plant Breeding, Junagadh Agricultural University, Junagadh, Gujarat, India

*Corresponding Author Email: csradadiya@jau.in

Received: 23 Jul 2026; Received in revised form: 21 Aug 2026; Accepted: 25 Aug 2026; Available online: 31 Aug 2026

©2026 The Author(s). Published by Infogain Publication. This is an open-access article under the CC BY license

(<https://creativecommons.org/licenses/by/4.0/>).

Abstract— Artificial intelligence (AI) is redefining hybrid breeding by shifting the operational paradigm from empirical field screening to predictive representation learning. While next-generation sequencing and high-throughput phenomics have generated massive multi-omics datasets, classical linear models fail to resolve the non-linear epistatic networks, dominance variances, and dynamic Genotype $\times$ Environment ($G \times E$) interactions that govern heterosis and combining ability. This review provides a novel, critical synthesis of AI architectures in hybrid prediction, categorizing deep learning paradigms based on their biological inductive biases. We dissect the transition from single-modality models to intermediate cross-attention multimodal fusion networks that synthesize genomic markers, latent-space phenomics, and temporal environmental covariates. Furthermore, we address critical field bottlenecks—specifically the "Diallel Sparsity Paradox," population structure confounding, data privacy barriers, and the out-of-distribution generalization failures of models facing climate anomalies. To resolve these, we propose the integration of Explainable AI (XAI) for biological validation, Federated Learning frameworks for cross-institutional collaboration, and Physics-Informed Neural Networks (PINNs) that embed deterministic crop physiological equations into deep learning loss functions. Finally, we map out an operational roadmap for Executable Digital Twin (xDT) breeding ecosystems, establishing a definitive blueprint for next-generation, climate-resilient hybrid development.



Keywords— Genomic Prediction, Heterosis, Inductive Bias, Multimodal Fusion, Physics-Informed Neural Networks, Explainable AI, Executable Digital Twin

I. INTRODUCTION

The exploitation of heterosis through hybrid breeding remains our most potent genetic tool for increasing global crop yield ceilings, stability, and biotic/abiotic stress resilience. However, identifying optimal parental combinations from vast germplasm pools presents an astronomical combinatorial challenge. In a standard breeding program utilizing n parental lines, the potential single-cross diallel space expands quadratically as:

$$\text{Number of unique cross combinations} = \frac{n(n-1)}{2}$$

Evaluating even a fraction of these combinations through multi-environment empirical field trials is logistically and financially impossible. This challenge is compounded by the complex genetic architecture of quantitative traits. Hybrid performance is not merely the sum of additive parental effects; it is governed by intricate dominance deviations, higher-order epistatic networks ($A \times A$, $A \times D$, $D \times D$), and highly dynamic Genotype $\times$ Environment ($G \times E$) interactions. Traditional quantitative genetics models, such as Genomic Best Linear Unbiased Prediction (G-BLUP) and its Bayesian variants, rely on explicit linear assumptions. While computationally robust, these classical

frameworks lack the mathematical flexibility to map the non-linear, rugged fitness landscapes characteristic of heterotic expressions.

Concurrently, modern agriculture has entered the era of multi-modal big data. High-density Single Nucleotide Polymorphism (SNP) microarrays and whole-genome resequencing provide ultra-high-dimensional genomic profiles. Simultaneously, high-throughput phenotyping platforms-ranging from Unmanned Aerial Vehicles (UAVs) equipped with hyperspectral and thermal sensors to automated ground rigs-generate high-frequency, time-series phenotypic streams. While these multi-dimensional data assets contain deep predictive signals, they present severe computational challenges: the curse of dimensionality ($p \gg n$), high multicollinearity, and massive data heterogeneity.

To resolve this, Artificial Intelligence (AI), specifically Machine Learning (ML) and Deep Learning (DL), offers a fundamental paradigm shift: predictive representation learning. Instead of manually crafting statistical assumptions, deep neural networks autonomously discover hierarchical, latent representations from unstructured data.

Critically, different AI architectures possess distinct inductive biases-built-in mathematical assumptions that align with specific biological structures. For example, Convolutional Neural Networks (CNNs) mirror the spatial structures of cellular and canopy phenotypes, Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks capture the temporal kinetics of plant development, and Graph Neural Networks (GNNs) map directly onto pedigree lineages and metabolic pathways (Seyum et al., 2022).

This review provides a highly specialized synthesis of the AI-driven hybrid breeding landscape. Moving beyond generic overviews of machine learning in agriculture, we critically evaluate how advanced neural architectures are being engineered to predict heterosis, decipher general and specific combining abilities (GCA/SCA), and model $G \times E$ dynamics. We address major structural bottlenecks, introduce emerging solutions like Physics-Informed Neural Networks (PINNs) and Explainable AI (XAI), and outline an operational blueprint for semi-autonomous, digital-twin breeding pipelines capable of outpacing climate volatility.

II. THE MULTI-MODAL DATA PARADIGM: FROM RAW SEQUENCES TO LATENT SPACE PHENOMICS

2.1 Haplotype-aware and network-embedded genomics

Traditional genomic selection treats SNP markers as independent, decoupled features, ignoring the structural architecture of the genome. Novel AI approaches leverage

whole-genome sequence data by constructing haplotype-aware tokenization. By grouping tightly linked SNPs into functional haplotype blocks, Self-Attention Transformers (such as the *Cropformer* architecture) can process genetic sequences analogously to natural language, capturing long-range linkage disequilibrium and complex epistatic configurations across chromosomes (Deng et al., 2025; Wang et al., 2024).

Furthermore, instead of analyzing genotypes in isolation, modern pipelines embed genomic data into biological networks. By utilizing Graph Convolutional Networks (GCNs), breeders can overlay individual genomic profiles onto known metabolic pathway graphs and protein-protein interaction networks, transforming flat string data into biologically contextualized feature spaces that preserve functional evolutionary relationships.

2.2 Latent space phenomics and spatial-temporal compression

Raw high-throughput phenotyping data is plagued by extreme redundancy and environmental noise. The frontier of phenomic AI lies in Latent Space Phenomics, which moves away from hand-measured agronomic traits (e.g., manual plant height calculations or basic vegetation indices) toward un-engineered, deeply learned representations.

Advanced architectures utilize Variational Autoencoders (VAEs) and Deep CNNs to compress gigabytes of multi-temporal UAV hyperspectral imagery into low-dimensional, high-information latent vectors. These latent features capture complex, sub-visual dynamics of canopy architecture, photosynthetic decline, and cellular stress. By utilizing LSTMs and Temporal Convolutional Networks (TCNs), these spatial latent vectors are tracked across the vegetative and reproductive phases, transforming static phenotypic snapshots into continuous, dynamic curves that reflect the hybrid's real-time growth velocity and plastic responses to diurnal stresses.

2.3 Topological data fusion: early vs. late vs. cross-attention intermediate fusion

The integration of genomic, phenomic, and environmental data streams represents a significant architectural challenge in AI breeding. Literature defines three primary topological fusion strategies:

- **Early Fusion (Concatenation):** Flat features from genomics, phenomics, and weather are directly concatenated into a single input vector. This strategy routinely fails because the massive dimensionality of genomic markers completely drowns out the subtle, time-series phenomic and climate signals.

- Late Fusion (Decision Ensemble): Independent models are trained for each modality (e.g., an ML model for SNPs, a CNN for images), and their final scalar predictions are averaged. While computationally simple, this completely misses cross-modal interactions—such as how a specific genetic combination alters a hybrid’s physiological response to a precise thermal threshold.
- Intermediate Cross-Attention Fusion (The Novel Frontier): Genomic, phenomic, and environmental streams are first passed through independent, modality-specific neural encoders. The resulting deep latent embeddings are then projected into a shared representation space where Cross-Attention Mechanisms mathematically calculate the interactions between modalities. This allows the model to dynamically weight specific genomic regions based on the real-time environmental stress captured by the phenomic sensors, unlocking explicit computational modeling of $G \times E$ mechanisms at the feature level.

2.4 The epigenomic layer and transgenerational stress memory

While the multi-omics framework traditionally unifies structural genomics and downstream transcriptomics, it frequently overlooks the epigenetic layer—a critical determinant of heterotic vigor and phenotypic plasticity. Non-additive gene expression in F_1 hybrids is heavily modulated by alterations in DNA methylation profiles, histone modifications, and small RNA (sRNA) interference patterns that differ markedly from both inbred parents.

Advanced AI pipelines ingest Epigenome-Wide Association Study (EWAS) matrices alongside traditional SNP profiles. By processing chromatin accessibility logs (e.g., ATAC-seq data) through specialized multi-layered Graph Attention Networks (GATs), AI models identify specific genomic loci where parental epigenetic states interact synergistically. Furthermore, embedding these epigenomic features allows

neural models to predict transgenerational stress memory—anticipating whether an F_1 hybrid will inherit a primed, epigenetically regulated tolerance to drought or salinity based on the microenvironmental exposures of its parental inbred lines.

III. ALGORITHMIC INDUCTIVE BIASES AND THE BREEDER'S OPERATIONAL FRAMEWORK

3.1 Mapping biological architecture to neural inductive biases

A major conceptual advancement in breeding AI is matching the mathematical architecture of an algorithm to the biological reality of the trait being modeled. Different algorithms possess unique inductive biases that make them fundamentally suited or unsuited for specific breeding objectives (Lopez et al., 2023).

For instance, Random Forests and Gradient Boosted Machines (XGBoost, LightGBM) feature an inductive bias based on orthogonal hyperplanes (decision trees). This makes them exceptionally robust at handling highly tabular, discrete data matrices and capturing lower-order epistatic interactions without requiring massive datasets. Conversely, the inductive bias of a CNN is spatial invariance and local connectivity, making it the definitive choice for processing raw cellular or canopy image fields. Transformers rely on a permutation-equivariant self-attention bias, allowing them to calculate the global context of an input, which matches the biological reality of long-range epistatic gene interactions across different chromosomes (Deng et al., 2025; Wang et al., 2024).

3.2 The pragmatic breeder's decision matrix

To transition AI from theoretical informatics to field operations, algorithm selection must be governed by a structured framework based on data topology, population size, and computational budgets.

Table 1: Strategic taxonomy of AI optimization in hybrid breeding pipelines

Model Class	Core Inductive Bias	Optimal Data Topology	Target Breeding Metric	Minimum Sample Size (N)	Computational Complexity
Tree Ensembles (RF, XGBoost)	Conditional orthogonal splitting	Tabular GCA/SCA matrices, low-density SNPs	Preliminary cross culling, major gene tracking	$N > 10^2$ lines	Low (CPU optimized)
Deep CNNs	Spatial locality and translation invariance	Multi-spectral imagery, LiDAR point clouds	Automated canopy architecture profiling, disease grading	$N > 10^3$ image sets	Medium-High (GPU mandatory)

LSTMs / TCNs	Temporal continuity and causal ordering	Daily weather logs, multi-temporal UAV sensor traces	Dynamic $G \times E$ stress trajectory modeling, yield stability	$N > 10^1$ environments	Medium (GPU accelerated)
Transformers	Permutation equivariance and global self-attention	Whole-genome haplotypes, multimodal intermediate embeddings	High-order epistasis modeling, zero-shot hybrid performance prediction	$N > 10^4$ crosses	Very High (Distributed GPU clusters)
Graph Neural Networks (GNNs)	Relational homophily and structural topology	Pedigree networks, metabolic pathways, kinships	Heterotic grouping, parental compatibility mapping	$N > 10^3$ nodes	Medium-High (Specialized hardware)
Physics-Informed Networks (PINNs)	Mechanistic constraint adherence	Fused marker data + physiological rate constants	Out-of-distribution climate anomaly generalization	$N > 10^1$ environments	High (Custom loss optimization)

3.3 Physics-informed neural networks (PINNs): resolving biological opacity

A profound limitation of standard deep learning models is their complete ignorance of physical and biological laws; they are purely statistical correlation engines. When exposed to extreme, out-of-distribution (OOD) environmental anomalies caused by climate change (e.g., unprecedented thermal spikes during anthesis), standard networks behave erratically, yielding biologically impossible predictions (Jighly et al., 2023; Technow et al., 2015). Physics-Informed Neural Networks (PINNs) solve this bottleneck by embedding deterministic biological laws directly into the neural network's loss function. Implementations like *DeepBioGS* demonstrate this integration effectively. Instead of optimizing purely for statistical error minimization (L_{data}), the total loss function (L_{total}) is formulated as:

$$L_{total} = L_{data} + \lambda L_{physics}$$

Where λ is a regularization hyperparameter balancing the tasks, and $L_{physics}$ represents the residuals of deterministic crop physiological differential equations, such as the Farquhar-von Caemmerer-Berry (FvCB) model for photosynthetic carbon assimilation or dynamic crop biomass increments (dW/dt) bounded by a dynamic temperature function $\mu(T)$ and maximum potential weight W_{max} :

$$\frac{dW}{dt} = \mu(T) \cdot W \cdot \left(1 - \frac{W}{W_{max}}\right)$$

If the neural network generates a hybrid yield or biomass prediction that violates these biological rate constants or thermodynamic limits, $L_{physics}$ increases dramatically,

penalizing the model (Viet Cuong et al., 2024). This mechanism enforces biological reality within the latent layers, ensuring the network maintains high predictive integrity even when encountering completely novel climate extremes.

IV. ADVANCED APPLICATIONS: RESOLVING THE DIALLEL SPARSITY PARADOX

4.1 Low-rank tensor decomposition and contrastive learning for heterosis

The primary operational obstacle in hybrid breeding is the Diallel Sparsity Paradox: the number of actualized field crosses represents a near-zero percentage of the potential combinatorial space. To predict the millions of untested crosses, advanced AI frameworks deploy Low-Rank Tensor Decomposition.

By treating the breeding pool as a multi-dimensional tensor ($\text{Parent}_A \times \text{Parent}_B \times \text{Environment}$), deep factorization algorithms project parents into low-dimensional latent embedding spaces. The specific combining ability (SCA) of an untested cross is then calculated as the inner product of these latent vectors. Mathematically, the tensor $\mathcal{X}$ is approximated via Canonical Polyadic (CP) decomposition:

$$\mathcal{X} \approx \sum_{r=1}^R a_r \circ b_r \circ c_r$$

Where a_r , b_r , and c_r represent the latent factor vectors for Parent A, Parent B, and the Environment, respectively. If the latent embeddings of a new parental cross display high cosine similarity within a complementary heterotic space,

the model yields a high predicted heterosis vigor without requiring physical field crosses.

Concurrently, Self-Supervised Contrastive Learning is applied to parental genomic sequences. The network is trained to maximize the cosine similarity between parents belonging to complementary heterotic groups while minimizing the similarity across non-complementary pools. By the time the model evaluates a prospective hybrid, it has already mapped the structural genetic distance required to maximize heterotic vigor, effectively shifting the discipline from empirical crossing to *in-silico* hybrid design.

4.2 Multi-task deep learning for GCA and SCA simultaneous optimization

Traditional genomic selection models predict individual traits or combining abilities via separate, isolated formulations. This misses the shared biological pathways

that regulate quantitative traits. Modern architectures utilize Multi-Task Learning (MTL) Networks, such as the *Variational auto-encoder Multi-task Genomic Prediction (VMGP)* framework.

In an MTL framework, a deep neural network splits into parallel prediction heads only at the final layers, while sharing all foundational hidden representations. This architecture forces the model to simultaneously output Maternal GCA, Paternal GCA, Cross-Specific SCA, and ultimate Multi-Environment Grain Yield. Because the hidden layers must simultaneously optimize all four parameters, the model captures deep, generalized biological rules governing gene action, significantly reducing overfitting and boosting predictive accuracy in highly unbalanced, sparse datasets.

V. CROP-SPECIFIC EXECUTION PARADIGMS AND CASE STUDIES




Fig.1: Cross-attention intermediate data fusion and parallelized multi-task neural network topology for simultaneous GCA/SCA optimization.

5.1 Maize: transformer-driven epistatic mapping

In maize (*Zea mays*), where heterotic patterns (e.g., Stiff Synthetic vs. Non-Stiff Synthetic) are distinct but genetically complex, recent pipelines utilize Vision Transformers (ViTs) fused with high-density SNP matrices. By collecting time-series hyperspectral drone imagery from parental inbreds, these models calculate an automated Parental Complementarity Index prior to pollination. This strategy enables the prediction of final kernel weight and stalk lodging resistance across highly divergent latitudinal zones, effectively mapping non-additive genetic networks at scale.

5.2 Rice: multimodal heterosis deciphering

In hybrid rice (*Oryza sativa*), the focus centers on resolving the complex trade-offs between grain quality metrics (amylose content, chalkiness) and heterotic yield advantages. Multi-task deep learning models ingest both genomic sequences and RNA-seq transcriptomic profiles from young panicles. These networks isolate non-additive gene expressions early in development, allowing breeders to cull lines with poor culinary quality or low fertility restoration before executing expensive field expansions.

5.3 High-value vegetables: automated spatial-spectral phenotyping

In high-value hybrid solanaceous and cucurbit crops (tomatoes, peppers, cucumbers), hybrid success is dictated

by strict physical uniformity, shelf-life, and precise disease resistance blocks. Here, deep CNN architectures (e.g., customized Mask R-CNNs) are integrated into automated greenhouse sorting lines and field robotics. These systems perform real-time, non-destructive 3D structural mapping of fruit geometry, color transitions, and canopy development curves, linking these phenotypic features straight to genomic matrices to fast-track parental lines optimized for mechanical harvesting.

VI. CRITICAL BOTTLENECKS AND MITIGATING STRATEGIES

6.1 The out-of-distribution (OOD) generalization failure

The most critical vulnerability of modern AI breeding pipelines is the drop in predictive accuracy when models are deployed in geographic regions or climatic years missing from the training data. If an AI is trained during five years of normal rainfall, it lacks the context to predict hybrid performance during a severe, multi-month megadrought.

Mitigation Strategy: Pipelines must move away from static environmental parameters (e.g., average seasonal temperature) toward Dynamic Environmental Covariates and Photothermal Time-Series (Persa et al., 2022). By tracking daily vapor pressure deficits (VPD), solar radiation

anomalies, and soil water depletion rates through LSTMs, and combining these vectors with Physics-Informed Neural Networks (PINNs), the model learns the physiological stress boundaries of the plant, preventing catastrophic predictive decay during erratic climate shifts.

6.2 Biological blind spots: rare alleles and epistatic noise

Deep learning models are inherently biased toward high-frequency features. Consequently, rare alleles that possess massive phenotypic effects under highly specific environmental stress scenarios are routinely smoothed out as statistical noise by deep neural layers.

Mitigation Strategy: Breeding programs must integrate Biologically Informed Graph Neural Networks. By hardcoding known metabolic pathways, transcription factor bindings, and functional annotation priorities into the graph topology of the network, rare variants occurring within critical catalytic domains are mathematically protected from down-weighting, preserving their predictive value for stress-adaptive traits.

6.3 The explainable AI (XAI) blueprint: building breeder trust

To transform AI from an opaque computational exercise into a trusted field validation engine, frameworks must incorporate post-hoc interpretability layers.

Table 2: Explainable AI (XAI) toolkit for model interpretability

XAI Methodology	Mathematical Mechanics	Direct Translation to Hybrid Breeding Decisions
SHAP (<i>SHapley Additive exPlanations</i>)	Calculates cooperative game-theoretic Shapley values across features.	Isolates the exact quantitative contribution of individual SNP markers or haplotype blocks to specific combining ability (SCA).
LIME (<i>Local Interpretable Model-agnostic Explanations</i>)	Builds local linear surrogate models around a specific deep learning prediction vector.	Explains why a specific hybrid cross was predicted to fail in a localized environment, identifying the exact climatic or phenotypic driver.
Integrated Gradients	Computes the path integral of gradients along the straight line from a baseline blank input to the target input.	Identifies which spectral bands in hyperspectral drone imaging are driving the model's early-season yield forecast.
Attention Map Visualization	Extracts internal weight matrices directly from the multi-head self-attention layers of Transformer models.	Maps the long-range epistatic connections between distant chromosomal regions that regulate heterotic vigor.

6.4 The reproducibility bottleneck: standardizing open-source breeding benchmarks

A significant roadblock to scaling AI in hybrid breeding is the lack of standardized benchmarking protocols. Historically, studies have utilized disparate validation schemes, isolated datasets, and inconsistent hyperparameter tuning budgets, rendering direct algorithmic comparisons nearly impossible. To build industry-wide trust, future

frameworks must adopt FAIR (Findable, Accessible, Interoperable, and Reusable) data principles, validating architectures against open-access foundational multi-omic repositories via unified open-source validation pipelines (such as the *MultiGS* benchmarking suite).

6.5 Confounding population structure and spurious heterotic correlations

A persistent risk in deploying deep representation learning for hybrid prediction is the model's susceptibility to population structure and cryptic kinship. Because deep neural networks excel at extracting high-dimensional patterns, they frequently learn to recognize broad ancestral lineages rather than causal epistatic or dominance configurations. Consequently, a model may show deceptively high prediction accuracy within a specific training family by simply acting as a pedigree classifier, yet fail completely when predicting crosses between novel, distantly related inbred lines.

To decouple true biological heterosis from ancestral confounding factors, future architectures must integrate Adversarial Domain Adaptation (ADA) or causal inference networks. By employing an adversarial loss term that penalizes the network if it can successfully identify the sub-population origin of a parental line, the model is forced to ignore broad lineage background effects and focus exclusively on the universal, functional genomic features that drive heterotic expression and combining ability.

6.6 The privacy-phygital divide: federated learning and resource-constrained "Green AI"

While the theoretical blueprint for AI-driven hybrid prediction relies on massive distributed GPU clusters and multi-omic inputs, operational implementation faces severe infrastructural and socio-political constraints. Public-sector breeding institutions, particularly within developing agricultural economies, operate under restrictive computational budgets that limit the adoption of high-

parameter Transformer networks. To democratize AI breeding, research must pivot toward "Green AI"—utilizing architectural compression techniques such as Knowledge Distillation (KD) to compress complex, multi-layered deep models into lightweight, CPU-optimized student networks without substantial loss in predictive capacity.

Second, private enterprise engagement is stymied by intense intellectual property boundaries. Commercial breeding programs are fundamentally reluctant to pool proprietary parental genotypes and phenotypic training records into centralized cloud frameworks. This roadblock can be circumvented via Federated Learning (FL). FL allows multi-institutional consortia to train a global hybrid prediction network collaboratively. Instead of moving raw multi-omics datasets to a central server, the model is trained locally at each separate institution, and only the encrypted mathematical weight gradients are shared. This protects localized germplasm privacy while enabling the AI to learn from unprecedented, cross-continental dataset volumes.

VII. FUTURE HORIZONS: THE EXECUTABLE DIGITAL TWIN (xDT) ECOSYSTEM

The ultimate evolution of AI in agriculture is the transition from passive digital simulations to integrated Executable Digital Twins (xDTs) driven by hierarchical Agentic AI frameworks. Instead of operating as a static repository for predictive models, an xDT serves as a live, interactive simulation engine that bridges the "phygital" divide between experimental fields and cloud computing architectures (Ghazvini et al., 2024).




Fig.2: Closed-loop architecture of the Executable Digital Twin (xDT) hybrid breeding framework

Under an agentic architecture, independent AI sub-agents—each specializing in discrete operational parameters such as canopy microclimate dynamics, soil nutrient depletion trajectories, or real-time disease progression profiles—cooperate under a centralized, physics-informed controller. By linking continuous IoT soil-moisture sensor streams, localized hyperspectral satellite monitoring, and high-density multi-omics arrays into unified cloud infrastructures, these digital twins continuously run thousands of stochastic "what-if" scenarios across varying regional climate projections daily. If a severe drought event is simulated during the microsporogenesis stage, the xDT automatically reweights the parental cross-selection matrix to prioritize lines with high dynamic stomatal conductance efficiency.

Furthermore, by utilizing Generative Adversarial Networks (GANs) and Diffusion Models designed for sequence generation, AI systems will soon move beyond selecting existing lines to generating synthetic, *in-silico* parental genotypes. These synthetic sequences, engineered to maximize heterotic complementarity while optimizing climate-resilient alleles, can be targeted for empirical creation via multiplexed CRISPR-Cas9 base editing. This convergence of generative AI, advanced representation learning, and precise genome editing will ultimately replace traditional empirical selection pipelines with fully predictive, autonomous crop design frameworks.

VIII. CONCLUSION

Artificial intelligence has advanced far beyond simple linear approximations, introducing an era of deep representation learning capable of decoding the complex, non-linear biological networks that drive heterosis and specific combining ability. By aligning algorithmic inductive biases with biological realities, integrating intermediate cross-attention multimodal architectures, and enforcing biological constraints through Physics-Informed Neural Networks, modern breeding programs can successfully overcome long-standing bottlenecks like the Diallel Sparsity Paradox, population lineage confounding, and climate-driven out-of-distribution failures. To achieve widespread field adoption, the discipline must continue to prioritize model interpretability via Explainable AI, standardizing cross-institutional data streams through secure federated frameworks, and building robust, biologically grounded executable digital twin ecosystems. Ultimately, embedding advanced AI frameworks into hybrid breeding pipelines represents our most scalable computational solution for fast-tracking genetic gains and securing global food systems under a volatile climate future.

REFERENCES

- [1] Deng, P., Liu, K., Zhou, M., Li, M., Yang, R., Cao, C., Wang, M., & Zhang, Z. (2025). DPCformer: An Interpretable Deep Learning Model for Genomic Prediction in Crops. *arXiv*, arXiv:2510.08662. <https://doi.org/10.48550/arxiv.2510.08662>
- [2] Ghazvini, A., Sharef, N. M., Balasundram, S. K., & Lee, L. S. (2024). A Concentration Prediction-Based Crop Digital Twin Using Nutrient Co-Existence and Composition in Regression Algorithms. *Applied Sciences*, 14(8), 3383. <https://doi.org/10.3390/app14083383>
- [3] Montesinos López, O., Mosqueda González, B., Montesinos López, A., & Crossa, J. (2023). Statistical Machine-Learning Methods for Genomic Prediction Using the SKM Library. *Genes*, 14(5), 1003. <https://doi.org/10.3390/genes14051003>
- [4] Persa, R., Graef, G. L., Specht, J. E., Rios, E., Messina, C. D., & Jarquin, D. (2022). Enhancing Genomic Prediction Models for Forecasting Days to Maturity in Soybean Genotypes Using Site-Specific and Cumulative Photoperiod Data. *Agriculture*, 12(4), 545. <https://doi.org/10.3390/agriculture12040545>
- [5] Seyum, E. G., Bille, N. H., Abteu, W. G., Munyengwa, N., Bell, J. M., & Cros, D. (2022). Genomic selection in tropical perennial crops and plantation trees: a review. *Molecular Breeding*, 42(42). <https://doi.org/10.1007/s11032-022-01326-4>
- [6] Technow, F., Messina, C. D., Totir, L. R., & Cooper, M. (2015). Integrating Crop Growth Models with Whole Genome Prediction through Approximate Bayesian Computation. *PLOS ONE*, 10(6), e0130855. <https://doi.org/10.1371/journal.pone.0130855>
- [7] VanRaden, P. M. (2008). Efficient Methods to Compute Genomic Predictions. *Journal of Dairy Science*, 91(11), 4414-4423. <https://doi.org/10.3168/jds.2007-0980>
- [8] Viet Cuong, D., Lalić, B., Petrić, M., Thanh Binh, N., & Roantree, M. (2024). Adapting physics-informed neural networks to improve ODE optimization in mosquito population dynamics. *PLOS ONE*, 19(1), e0315762. <https://doi.org/10.3371/journal.pone.0315762>
- [9] Wang, H., Yan, S., Wang, W., Chen, Y., Hong, J., He, Q., Diao, X., Lin, Y., Chen, Y., Cao, Y., Guo, W., & Fang, W. (2024). Cropformer: An interpretable deep learning framework for crop genomic prediction. *Plant Communications*, 101223. <https://doi.org/10.1016/j.xplc.2024.101223>