Causal World Model for Genomic Prediction

Zhuo Huang[1], Javen Shi[1]
[1]Adelaide University

doi.org/10.25909/34227735

ABSTRACT

Genomic prediction (GP) can be easily misled by correlation due to confounding noise in the real world. Moreover, the dynamics and complexity of natural environments further hinder prediction accuracy and planning. Hence, we aim to develop causal world models (CWM) that can extract and uncover the latent causal factors for GP, meanwhile enabling robust sequential planning under evolving environments to optimise long-horizon benefits.

Research Gap

Current GP progress suffers from four major gaps:

Forward Inference

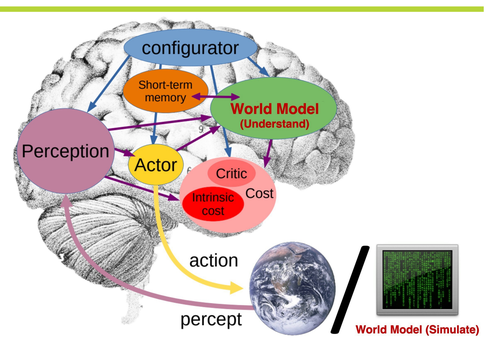
- Insufficient data or conditions to identify the causal factors.
- Limited adaptability across various environments or varieties.

Reverse Inference

- High trial costs for real-world practices.
- Lack of long-horizon planning for long-term benefits.

World Model Background

A world model [1] learns an internal representation of how a system evolves, enabling it to predict the consequences of actions under different conditions.



The framework is formulated as:

$$z_t = h_\theta(G_t, E_t, M_t), \hat{z}_{t+1} = f_\theta(z_t, a_t), \hat{Y}_{t+1} = d_\theta(\hat{z}_{t+1}).$$

In GP, it can represent how genotype, environment, and management jointly influence crop phenotype over time. Further, it supports planning by allowing candidate actions and future outcomes to be evaluated before real-world application.

Our training goals are as follows:

$$\mathcal{L}_{\text{JEPA}} = \mathbb{E}[\|\hat{z}_{t+1}, \text{sg}(z_{t+1})\|_2^2], \mathcal{L}_{\text{GP}} = \mathbb{E}[(\hat{Y}_{t+1} - Y_{t+1})^2]$$

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{JEPA}} + \lambda\mathcal{L}_{\text{GP}} + \mathcal{L}_{\text{reg}}.$$

The three objectives aim for: 1) $\mathcal{L}_{\text{JEPA}}$: fully recover the future representation given its prediction; 2) $\mathcal{L}_{\text{GP}}$: correctly prediction future yield; 3) $\mathcal{L}_{\text{reg}}$: prevent latent collapse using regularization.

Causal World Model

The challenge for CWM [2] in GP is to identify (or partially) causal structure without the influence of environmental noises.

$$\text{ACE}_{i,\theta_1}(g_1, g_0) = \text{ACE}_{i,\theta_2}(g_1, g_0), \forall\theta_2, \theta_2 \in \Theta,$$

$$\text{where ACE}_{i,\theta}(g_1, g_0) = \mathbb{E}[\hat{Y}_\theta|\text{do}(G_t^i) = g_1] - \mathbb{E}[\hat{Y}_\theta|\text{do}(G_t^i) = g_0].$$

We tackle this issue via data management and model training.

Data Management:

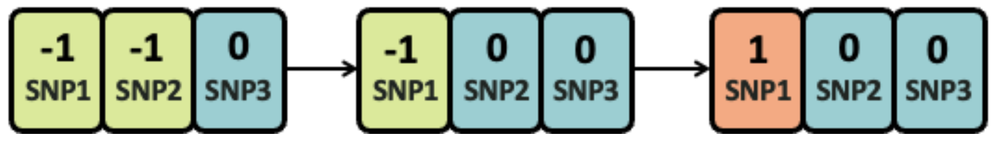
1. LD Pruning. Genotype are high dimensional, while the grain population is relatively limited. Thus, redundant SNPs are removed to avoid computational overhead.
2. APSIM simulation. By acquiring coordinate, sowing/harvest dates, etc, we can simulate weather, temperature, radiation.

Model Training:

Isolate factors via minimum variation. Let $s_t = (G_t, E_t, M_t)$, we prioritise genotype data that satisfy:

$$\|h_\theta(s_t') - h_\theta(s_t)\|_0 < \tau, \text{where } s_t' = s_t + \delta e_i, 0 < |\delta| < \epsilon$$

here τ is the threshold for number of variation factors, e_i is the basis vector for varying the i-th SNP. Intuitively, selecting minimum variation during latent prediction training helps isolating true causal factors. With sufficient variation coverage and invertibility mapping, we can identify the causal structure.



Action Planning

Given a CWM, we aim to optimise long-term farming plan:

$$(M_t^*, \dots, M_{t+H-1}^*) = \arg\max_{M_{t:t+H-1}} \mathbb{E}[\sum_{k=0}^{H-1} \gamma^k R(\hat{Y}_{t+k}, M_{t+k})],$$

here $(M_t^*, \dots, M_{t+H-1}^*)$ is the optimal sequence of management plan over horizon H, $R(\cdot)$ is the cost function with weight γ^k. Existing planning approaches solved the optimisation by cross-entropy methods (CEM) [3] or gradient-based methods [4].

Empirical Results

We evaluate our initial progress on single environment from AGT dataset compared with both classic and deep learning methods.

Within-Environment Comparison

Method	Acc (%)	RMSE	MAE	R^2	Pearson r
CWM	92.23	424.65	356.01	0.3123	0.5420
gBLUP	91.98	453.50	356.01	0.2929	0.5380
BayesC	91.75	453.50	354.38	0.2902	0.5257
DNNGP	91.75	453.41	350.32	0.2810	0.4833
BayesB	91.21	454.18	341.41	0.2740	0.4677
GPTransformer	91.11	454.34	333.84	0.2734	0.3574
BayesA	91.04	450.58	345.36	0.2623	0.4152
RR-BLUP	90.91	457.90	350.79	0.2557	0.4582
Bayesian LASSO	89.81	456.19	356.29	0.2451	0.4561
DeepGS	88.97	467.63	381.68	0.2394	0.4424

Within-Environment vs. Cross-Environment Evaluation

Acc (%)	01	02	03	04	05	07	08	09	10	11	12
In	93.41	89.22	92.83	92.58	92.05	90.82	92.27	92.08	93.99	94.81	90.47
Out	64.25	87.56	87.74	90.34	77.65	87.24	70.67	57.54	71.12	66.96	73.64

References

[1] Ha, David, and Jürgen Schmidhuber. "World models." arXiv preprint arXiv:1803.10122 2.3 (2018): 440.

[2] Schölkopf, Bernhard, et al. "Toward causal representation learning." Proceedings of the IEEE 109.5 (2021): 612-634.

[3] Rubinstein, Reuven. "The cross-entropy method for combinatorial and continuous optimization." Methodology and computing in applied probability 1.2 (1999): 127-190.

[4] Ratliff, Nathan, et al. "CHOMP: Gradient optimization techniques for efficient motion planning." 2009 IEEE international conference on robotics and automation.

The support of the Grains Research and Development Corporation (GRDC) provided through the Analytics for the Australian Grains Industry (AAGI) project is acknowledged.