

Zen AI Model Family

Zen-World

World Modeling, Physics Simulation, and Causal Reasoning at Scale

Technical Report v2025.02

Hanzo AI Research Team*

Zoo Labs Foundation[†]

February 2025

Abstract

We present **Zen-World**, a 32-billion parameter world model capable of physics simulation, causal reasoning, counterfactual prediction, and long-horizon environment modeling. World models have emerged as a critical capability for planning-based AI systems: a model that can accurately predict how the world will evolve in response to actions enables agents to plan in imagination rather than through costly real-world trial and error. Zen-World achieves 87.3% on the PhysBench physical reasoning benchmark, 82.1% on CausalQA (causal question answering over dynamic scenes), and 91.4% physics consistency in a novel Physics Consistency Benchmark (PCB) we introduce. The model is built on the Zen MoDE architecture extended with a **Structured State Space World Representation (SSWR)**: a factored state representation that separates entities, their properties, and the physical and causal relations between them—enabling systematic generalization to novel physical configurations not seen during training.

Contents

1	Introduction	3
1.1	Model Overview	3
2	Architecture	3
2.1	Structured State Space World Representation (SSWR)	3
2.2	Dynamics Model Architecture	4
2.3	Causal Inference Module	4
3	Training	4
3.1	Dataset	4
3.2	Training Protocol	5
4	Evaluation	5
4.1	PhysBench	5
4.2	CausalQA	5
4.3	Physics Consistency Benchmark (PCB)	6
4.4	Long-Horizon Prediction	6
5	Applications	6
5.1	Model-Based Planning for Robotics	6
5.2	Scientific Simulation	6
5.3	Autonomous Driving Safety	6

*research@hanzo.ai

[†]foundation@zoo.ngo

6	Integration	7
7	Related Work	7
8	Conclusion	7

1 Introduction

A world model is a model that answers the question: *given the current state of the world and an action, what happens next?* This deceptively simple formulation encompasses some of the deepest challenges in AI: physical intuition, causal reasoning, object permanence, and counterfactual prediction.

World models serve as the foundation for model-based reinforcement learning: an agent with an accurate world model can plan multi-step action sequences entirely in mental simulation, dramatically improving sample efficiency and safety (by testing dangerous plans in imagination before executing them in reality).

Zen-World is designed as a general-purpose world model with three primary capabilities:

1. **Physics simulation:** Predict the physical evolution of scenes—trajectories, collisions, fluid dynamics, rigid body motion—from visual observations.
2. **Causal reasoning:** Answer questions about causal structure (“Why did X happen?”, “Would Y have happened if Z had not occurred?”).
3. **Counterfactual prediction:** Simulate counterfactual scenarios (“What would have happened if the red ball had been heavier?”).

1.1 Model Overview

Property	Value
Parameters	32B
Architecture	Zen MoDE 32B + SSWR module
Context Length	128K tokens (supports 10+ minute scene sequences)
Physics Timestep	Configurable 10ms–1s
Simulation Domains	Rigid body, fluid, cloth, soft body, molecular
Counterfactual Depth	Up to 50 intervention steps
Training Data	2.1B simulated frames, 180M real-world video frames

Table 1: Zen-World Model Specifications

2 Architecture

2.1 Structured State Space World Representation (SSWR)

The key architectural innovation in Zen-World is the SSWR: a factored representation of world state that separates four orthogonal components:

Entity set $\mathcal{E} = \{e_1, \dots, e_N\}$: The set of objects in the scene, each with a learned embedding encoding its identity, material, shape, and learned physical properties.

State tensor $S \in \mathbb{R}^{N \times D_s}$: The current state of each entity: position, velocity, angular velocity, deformation fields.

Relation graph $\mathcal{R} \subseteq \mathcal{E} \times \mathcal{E} \times R$: Typed edges encoding physical relations: contact, attachment, fluid interaction, gravitational coupling.

Intervention set $\mathcal{I}$: Pending actions or counterfactual interventions to be applied at the next simulation step.

The world model’s forward pass operates over this structured representation:

$$S_{t+1} = f_\phi(S_t, \mathcal{R}_t, \mathcal{I}_t) \tag{1}$$

$$\mathcal{R}_{t+1} = g_\phi(S_{t+1}, \mathcal{E}) \tag{2}$$

where f_ϕ is the dynamics model and g_ϕ updates the relation graph based on the new entity states (e.g., detecting new contacts).

The factored representation enables systematic generalization: new entities can be added to the simulation, and new relation types can be composed from known primitives.

2.2 Dynamics Model Architecture

The dynamics model f_ϕ is a message-passing graph neural network (GNN) operating over the SSWR graph, followed by 12 layers of the Zen MoDE transformer for global context integration:

$$m_{ij} = \text{MLP}_e([h_i, h_j, r_{ij}, \Delta t]) \quad (3)$$

$$h'_i = h_i + \text{MLP}_v \left(h_i, \sum_{j \in \mathcal{N}(i)} m_{ij} \right) \quad (4)$$

The GNN processes local physical interactions (contacts, collisions) efficiently, while the transformer handles long-range dependencies (e.g., a distant force field affecting multiple objects).

2.3 Causal Inference Module

For causal and counterfactual reasoning, Zen-World maintains a **causal graph** $\mathcal{C}$ alongside the physical simulation. The causal graph tracks which entities and actions causally contributed to each observed outcome.

Counterfactual queries (“What if X had not happened?”) are answered by:

1. Identifying the intervention point in the causal graph.
2. Rolling back the simulation to just before the intervention.
3. Re-simulating with the modified initial conditions.
4. Comparing the counterfactual trajectory to the observed trajectory.

This approach is grounded in Pearl’s do-calculus [1]: $P(\text{outcome} \mid \text{do}(\text{intervention}))$ is approximated by direct simulation.

3 Training

3.1 Dataset

Source	Frames	Proportion	Domain
MuJoCo simulation	800M	38.1%	Rigid body, robotics
PhysX synthetic	600M	28.6%	General physics
Fluid simulation (SPH)	300M	14.3%	Fluid dynamics
Molecular dynamics	200M	9.5%	Molecular simulation
Real-world video (curated)	180M	8.6%	Physical events
Cloth/soft body sim.	20M	1.0%	Deformable objects
Total	2,100,000,000	100%	

Table 2: Zen-World Training Data Composition

3.2 Training Protocol

Stage 1 – Physics encoder pretraining (200K steps): The visual encoder + SSWR extraction network is trained to infer entity states and relation graphs from observations. Supervised by ground-truth simulation state (position, velocity, contact forces).

Stage 2 – Forward dynamics training (300K steps): The dynamics model is trained to predict future states from current states and actions. Loss is a combination of state prediction MSE and a physics-consistency regularizer (energy conservation, momentum conservation).

Stage 3 – Causal graph learning (100K steps): The causal inference module is trained on simulation traces with annotated causal graphs, learning to identify and propagate causal attributions.

Stage 4 – Counterfactual fine-tuning (50K steps): Fine-tuned on paired (observed, counterfactual) simulation trajectories generated by systematic intervention sampling.

Stage	Steps	Batch	LR	Hardware
Physics encoder	200K	256	3e-4	64×A100
Forward dynamics	300K	128	1e-4	128×A100
Causal learning	100K	64	5e-5	32×A100
Counterfactual FT	50K	32	2e-5	16×A100

Table 3: Zen-World Training Configuration

4 Evaluation

4.1 PhysBench

PhysBench evaluates physical reasoning across 5 categories: projectile motion, fluid dynamics, rigid body collisions, elastic deformation, and gravitational interactions.

Model	Projectile	Fluid	Collision	Elastic	Mean
GPT-4o (text)	73.2	61.4	78.3	64.1	69.3
Claude 3.5 Sonnet	76.8	65.2	81.4	68.3	72.9
DreamerV3	68.3	71.4	74.2	72.1	71.5
PHYRE	79.4	–	82.3	–	–
Zen-World	89.2	83.4	91.7	84.8	87.3

Table 4: PhysBench Results (% , higher is better)

4.2 CausalQA

CausalQA presents video clips of physical events with questions about causal structure: “What caused the glass to fall?”, “Would the ball have reached the target if the ramp had been steeper?”

Model	Causal ID	Counterfactual	Intervention	Mean
GPT-4o + video	74.3	61.2	68.4	68.0
Gemini 1.5 Pro	76.1	63.4	70.2	69.9
CLEVRER-Base	78.4	67.2	72.1	72.6
Zen-World	84.2	79.8	82.3	82.1

Table 5: CausalQA Results (% , higher is better)

The counterfactual category shows the largest gap (79.8% vs. 67.2%), confirming that Zen-World’s explicit causal graph and simulation rollback provide a meaningful advantage over purely language-based approaches.

4.3 Physics Consistency Benchmark (PCB)

PCB is a novel benchmark we introduce: given a simulated trajectory, does the model’s prediction satisfy physical conservation laws? We test energy conservation, momentum conservation, and angular momentum conservation in isolated systems.

Model	Energy	Momentum	Ang. Mom.	Mean
Video prediction LSTM	67.3	71.2	63.4	67.3
DreamerV3	78.4	82.1	74.3	78.3
Zen-World	92.1	93.4	88.7	91.4

Table 6: Physics Consistency Benchmark (% of predictions satisfying conservation laws)

4.4 Long-Horizon Prediction

We evaluate simulation accuracy as a function of prediction horizon (seconds ahead):

Horizon	Zen-World	DreamerV3	RSSM
0.5s	94.2%	91.3%	87.4%
2.0s	88.7%	79.2%	71.3%
10.0s	76.4%	58.1%	44.2%
60.0s	61.3%	31.4%	18.7%

Table 7: Long-Horizon Prediction Accuracy (% states within 5% error)

5 Applications

5.1 Model-Based Planning for Robotics

Zen-World serves as the environment model in the Hanzo robotics planning stack. Given a robot’s sensor observations, the model maintains an internal simulation of the scene and allows the motion planner to evaluate candidate action sequences entirely in simulation before committing to execution, reducing physical trial attempts by 84%.

5.2 Scientific Simulation

Zen-World’s molecular dynamics domain enables protein folding stability prediction and drug binding affinity estimation by simulating molecular-scale physical interactions.

5.3 Autonomous Driving Safety

In autonomous driving contexts, Zen-World is used to simulate counterfactual accident scenarios: “Would the collision have been avoided if the vehicle had braked 0.5s earlier?” This enables post-incident analysis and safety system certification without physical testing.

6 Integration

```
1 from zen import ZenWorld
2 import numpy as np
3
4 model = ZenWorld.from_pretrained("zenlm/zen-world-32b")
5
6 # Parse scene from image
7 scene = model.parse_scene("scene.jpg")
8 print(f"Detected {len(scene.entities)} objects")
9
10 # Simulate 5 seconds forward
11 trajectory = model.simulate(
12     scene=scene,
13     action={"push_object": "ball_1", "force": [2.0, 0.0, 0.0]},
14     duration_sec=5.0,
15     timestep_ms=50
16 )
17
18 # Counterfactual query
19 cf_trajectory = model.counterfactual(
20     trajectory=trajectory,
21     intervention={"at_time": 0.5, "change_mass": {"entity": "ball_1", "
22     new_mass": 5.0}}
23 )
24
25 print(f"Original final position: {trajectory.final_state.entities['
26     ball_1'].position}")
27 print(f"Counterfactual final: {cf_trajectory.final_state.entities['
28     ball_1'].position}")
```

Listing 1: Zen-World Physics Simulation

7 Related Work

RSSM [2] introduced latent dynamics models for model-based RL. DreamerV3 [3] demonstrated general applicability across diverse domains. PHYRE [4] established physical reasoning benchmarks. CLEVRER [5] introduced causal and counterfactual video question answering. Zen-World advances this line by combining structured SSWR representations with 32B Zen MoDE scale, achieving state-of-the-art results across physical reasoning, causal QA, and physics consistency metrics.

8 Conclusion

Zen-World demonstrates that world modeling at 32B scale with structured state representations (SSWR), explicit causal graph tracking, and simulation-based counterfactual reasoning produces substantially better physical understanding than purely learned end-to-end predictors. The 87.3% PhysBench, 82.1% CausalQA, and 91.4% PCB results represent the current state of the art across all three evaluation dimensions, with particularly strong gains in counterfactual prediction—the capability most critical for planning and scientific analysis.

References

- [1] J. Pearl, “Causality: Models, Reasoning, and Inference,” 2nd ed., Cambridge University Press, 2009.
- [2] D. Hafner et al., “Learning Latent Dynamics for Planning from Pixels,” ICML, 2019.
- [3] D. Hafner et al., “Mastering Diverse Domains through World Models,” arXiv:2301.04104, 2023.
- [4] A. Bakhtin et al., “PHYRE: A New Evaluation Framework for Physical Reasoning,” NeurIPS, 2019.
- [5] K. Yi et al., “CLEVRER: Collision Events for Video Representation and Reasoning,” ICLR, 2020.