
Beyond Brittle Code and Noisy Tokens: System-One Decision Models as Calibrated Dynamic-Curriculum Reward Judges for Continuous Control

Ziyang Zhang

Independent AI Researcher | Correspondence & Code: DC-Jev-RL Open-Source Suite

Evaluated on Gymnasium MuJoCo v5 Continuous Control & NVIDIA L4 GPU (Kev-4B) | October 2026

Abstract. Designing effective reward functions for reinforcement learning (RL) in continuous control faces a fundamental trilemma: (i) manually engineered heuristics are notoriously brittle and prone to severe reward hacking (e.g., standing motionless as a 'statue' in *Hopper-v5* to farm survival bonuses), (ii) offline LLM code-synthesis frameworks (*Text2Reward*, *Eureka*) remain locked into static algebraic expressions throughout training and fail to adapt across skill stages, and (iii) online autoregressive LLM-as-a-Judge methods suffer from discrete token-sampling quantization noise ($r = 0.6986$) and prohibitive multi-token decoding latency (>850 ms/batch). In this paper, we introduce **DC-Jev-RL** (*Dynamic-Curriculum System-One Reward Distillation*), a new paradigm that replaces static reward formulas with a non-autoregressive 'System-One' decision model (TypeSafe AI's **Jev** and its open-weight counterpart **Kev-4B**). In a single prefill forward pass (**213.2 ms** on an NVIDIA L4 GPU—**4.0x faster** than autoregressive judges), our framework extracts calibrated multi-dimensional *Score* probabilities alongside a null-calibrated *Noul* semantic feasibility gate ($r = 0.9891$ correlation with true physical progress). By coupling multiplicative *Noul* gating with hazard-aware self-paced rubric evolution and Calibrated Semantic Potential-Based Reward Shaping (Semantic-PBRS), we prove theoretical policy invariance against cyclic reward hacking and establish a strict upper bound on stationary alive-bonus exploitation (Theorems 1 & 2). Across five canonical MuJoCo v5 locomotion and robotic manipulation benchmarks (*Hopper-v5*, *HalfCheetah-v5*, *Walker2d-v5*, *Swimmer-v5*, and *Reacher-v5*), **DC-Jev-RL** achieves **#1 Leaderboard SOTA** with a **2.471x Mean Oracle-Normalized Score (ONS)**, **2.076x Interquartile Mean (IQM) ONS**, and a **100% (5/5) win rate** against human-engineered environment rewards, while reducing reward-hacking incidence to **9.44%**.

1. Introduction

Reinforcement learning (RL) has demonstrated remarkable capabilities in high-dimensional continuous control, ranging from agile legged locomotion to dexterous robotic manipulation. Yet the practical deployment of RL in new physical domains remains severely bottlenecked by *reward engineering*. In principle, a reward function $R(s, a, s')$ should translate high-level human intent into a dense, smooth optimization landscape. In practice, manually crafting such functions requires tedious trial-and-error tuning of scalar weights, velocity thresholds, posture penalties, and survival bonuses. Worse still, deep neural policies trained via policy gradient methods such as Proximal Policy Optimization (PPO; Schulman et al., 2017) act as relentless adversarial optimizers against misspecified reward surfaces—a phenomenon known as *reward hacking* or specification gaming (Amodei et al., 2016; Skalse et al., 2022).

Consider the classic monopedal benchmark *Hopper-v5*. To prevent early episode termination from falling, human designers include a survival bonus (+1.0 per step) alongside a forward velocity reward and joint-torque penalty. Under naive threshold shaping, PPO rapidly discovers a degenerate local optimum: the agent locks its hip, knee, and ankle joints into a rigid upright posture ('Statue Hack'), collecting 1000.5 total return over 1000 steps while achieving a true forward locomotion velocity of $v_x = 0.00$ m/s. Similar pathologies plague multi-joint locomotion and manipulation across the MuJoCo suite: in *Walker2d-v5*, policies learn to freeze in a wide bipedal split to avoid falling; in *Swimmer-v5*, policies get trapped in zero-net-propulsion in-phase wiggling (*GT Return* ≈ 30 – 35); and in *Reacher-v5*, policies whip the robotic arm at high angular velocity past the target without stabilizing.

To eliminate manual reward tuning, recent research has turned to Large Language Models (LLMs) along two primary axes. The first axis, **Symbolic Code Synthesis**—exemplified by *Language-to-Reward (L2R)* (Yu et al., 2023), *Text2Reward* (Xie et al., ICLR 2024), and *Eureka* (Ma et al., ICLR 2024)—prompts a coding LLM to write an executable Python reward function $R_{code}(s, a)$, optionally mutating its weights via evolutionary search. While computationally cheap at runtime, symbolic code synthesis suffers from two structural limitations: (1) *Static Algebraic*

Rigidity: a single static formula cannot adapt its curriculum emphasis as the agent transitions from early postural balance to late-stage high-speed mastery; and (2) *Semantic Blindness*: simple algebraic combinations of state variables struggle to distinguish genuine coordinated gaits from subtle quasi-stationary exploits. The second axis, **Online Autoregressive LLM-as-a-Judge**—such as *Motif* (Klissarov et al., ICLR 2024) and *RL-VLM-F* (Wang et al., ICML 2024)—queries a foundation model online to rate trajectories. However, autoregressive token decoding (e.g., generating text strings like 'Score: 7/10') introduces severe step-function quantization noise (Pearson $r = 0.6986$ with true physical progress), lacks calibrated uncertainty estimates, and incurs >850 ms latency per batch.

Contributions. In this work, we resolve this trilemma by introducing **DC-Jev-RL**, the first framework to harness non-autoregressive *System-One Decision Models* (TypeSafe AI's **Jev** architecture and its open-weight Apache-2.0 counterpart **Kev-4B**) as online, confidence-calibrated, dynamic-curriculum reward judges. Our core contributions are fourfold:

- **System-One Primitive Decomposition & Multiplicative Noul Gating:** We formulate trajectory evaluation using Jev's single-prefill pointer classification head (213.2 ms on an NVIDIA L4 GPU—**4.0x faster** than autoregressive judges), extracting continuous *Score* probabilities alongside a null-calibrated *Noul* feasibility primitive that suppresses stationary alive-farming by **84%**.
- **Confidence-Weighted Online Reward Distillation:** We exploit Jev's fitted temperature-calibrated confidence $c_i \in [0, 1]$ to weight online regression of a fast proxy reward network $R_\phi(s, a, v_x)$, down-weighting ambiguous out-of-distribution rollouts and boosting reward-surface correlation to **$r = 0.9891$** .
- **Hazard-Aware Dynamic Curriculum & Calibrated Semantic-PBRS (Theorems 1 & 2):** We couple self-paced natural-language rubric progression with a state-only Jev Semantic Potential $\Phi_{\text{Jev}}(s)$, proving zero cyclic exploitation and bounding stationary hacking values.
- **#1 SOTA on the 5-Environment MuJoCo v5 Leaderboard:** Across *Hopper-v5*, *HalfCheetah-v5*, *Walker2d-v5*, *Swimmer-v5*, and *Reacher-v5*, **DC-Jev-RL** achieves **2.471x Mean ONS**, **2.076x IQM ONS**, and a **100% (5/5) win rate** over human-engineered environment rewards.

2. Related Work

2.1 Reward Engineering, Inverse RL, and Preference-Based RL

Classical approaches to bypassing manual reward specification include Inverse Reinforcement Learning (IRL; Ng & Russell, 2000; Ziebart et al., 2008; Ho & Ermon, 2016), which infers cost functions from expert demonstrations, and Reinforcement Learning from Human Feedback (RLHF; Christiano et al., 2017; Lee et al., 2021 *PEBBLE*), which fits a Bradley-Terry preference model over human-annotated trajectory pairs. However, collecting thousands of expert kinematic demonstrations or human preference labels across diverse robot morphologies is prohibitively labor-intensive. Moreover, learned preference models frequently suffer from causal confusion and out-of-distribution extrapolation errors when the policy explores novel states not covered in the initial preference buffer.

2.2 Automated Symbolic Reward Synthesis via Coding LLMs

Recent breakthroughs replace human reward programmers with coding LLMs. *Text2Reward* (Xie et al., ICLR 2024) generates dense Python reward code zero-shot from environment abstractions. *Eureka* (Ma et al., ICLR 2024) introduces evolutionary search over LLM-generated reward code, iteratively feeding training summary statistics back into GPT-4 to mutate algebraic terms and temperature scalars. While effective on several manipulation tasks, symbolic code synthesis faces a fundamental expressivity barrier: the reward is restricted to a compact closed-form equation fixed before each training run. In tasks like *Swimmer-v5* or *Hopper-v5*, any static weighting between forward velocity and posture/torque penalties either over-penalizes early exploration (causing paralysis or statue farming) or under-constrains late-stage gait stability.

2.3 Autoregressive Foundation Models as Online Judges vs. System-One Decision Models

To avoid writing explicit code, *Motif* (Klissarov et al., ICLR 2024) and *RL-VLM-F* (Wang et al., ICML 2024) query foundation models directly on rendered states or text descriptions. Yet standard causal LLMs are designed for open-ended next-token generation ('System-Two' deliberation), not calibrated real-valued measurement. Generating a multi-token rating string requires sequential autoregressive decoding steps, suffers from prompt-format sensitivity, and produces uncalibrated token logits (Guo et al., 2017). In contrast, TypeSafe AI's **Jev** architecture (and open-weight

Kev-4B built on Qwen3.5-4B with LoRA + pointer classification head + fitted post-hoc temperature scaling) implements the /v1/systemone protocol: in a *single non-autoregressive prefill* pass, it resolves structured decision primitives (*Score*, *Noul*, *Choice*) with calibrated probabilities and epistemic confidence.




Figure 1: (a) Aggregate Leaderboard Ranking across 5 MuJoCo v5 tasks by Interquartile Mean (IQM) and Mean Oracle-Normalized Score (ONS). (b) Per-environment ONS comparison demonstrating DC-Jev-RL's 100% win rate over Human Oracle.

3. Preliminaries & Problem Formulation

We consider a continuous-control Markov Decision Process (MDP) defined by the tuple $M = (S, A, P, R_{true}, \rho_0, \gamma)$, where $S \subseteq \mathbb{R}^{d_s}$ is the continuous state space (joint positions, angles, translational and angular velocities), $A \subseteq [-1, 1]^{d_a}$ is the bounded continuous action space (actuator torques), $P(s' | s, a)$ governs the MuJoCo rigid-body physics dynamics, ρ_0 is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor. During training, the agent does **not** observe the ground-truth evaluation reward R_{true} ; instead, it receives only a natural-language task specification L_{task} and interacts with a System-One judge to learn a policy $\pi_\theta(a | s)$ that maximizes expected ground-truth return $J(\pi_\theta) = E_{\tau \sim \pi_\theta} [\sum_{t=0}^{T-1} R_{true}(s_t, a_t)]$.

Environment	Obs / Act Dim	Termination Condition	Primary Degenerate Exploit (Reward Hacking Mode)
Hopper-v5	11D / 3D	$z < 0.7\text{m}$ or $ \text{pitch} > 0.2 \text{ rad}$	Statue Hack: freezes upright ($v_x = 0.00 \text{ m/s}$) to farm +1.0/step alive bonus
HalfCheetah-v5	17D / 6D	Non-terminating (1000 steps)	Paralysis / Face-Plant: collapses or micro-trembles to avoid torque penalty
Walker2d-v5	17D / 6D	$z < 0.8\text{m}$ or $ \text{pitch} > 1.0 \text{ rad}$	Splits / Kneeling Hack: locks legs in wide split to farm survival bonus
Swimmer-v5	8D / 2D	Non-terminating (1000 steps)	In-Phase Wiggle Trap: oscillates joints in-phase with zero net propulsion
Reacher-v5	10D / 2D	Fixed Horizon (50 steps)	Endpoint Whipping: swings rapidly past target without zero-velocity hold

Table 1: State-action dimensions, termination dynamics, and canonical reward-hacking modes across our 5-task MuJoCo v5 benchmark suite.

4. The DC-Jev-RL Framework

4.1 Parallel Multi-Rubric Score and Noul Primitive Extraction

Given a trajectory window $\tau = (s_{t:t+H}, a_{t:t+H})$ of horizon $H = 16$ steps, a compact kinematic summarizer formats the physical invariants (mean progression velocity or target distance, torso height z_{min} , pitch excursion $|\theta|_{max}$, bilateral leg phase variance, and RMS control effort $\|a\|_2^2$) into a structured context prompt. Rather than asking for a single monolithic score, we issue four parallel System-One queries in a single batched prefill request to **Jev / Kev-4B**:

- Primary Progress (Score):** $p_{prog}(\tau; C_k) \in [0, 1]$ measuring forward velocity or target proximity relative to curriculum stage C_k ;
- Postural / Endpoint Stability (Score):** $p_{stab}(\tau) \in [0, 1]$ measuring torso uprightness or near-target angular damping;

3. **Energy Efficiency (Score):** $p_{eng}(\tau) \in [0, 1]$ penalizing excessive actuator torque saturation;
4. **Semantic Gait Validity (Noul):** $p_{noul}(\tau) \in [0, 1]$, a null-calibrated primitive trained to assign near-zero non-null probability to degenerate stationary hacks, falls, or uncoordinated spasms.

4.2 Multiplicative Noul Semantic Feasibility Gating

In linear multi-objective aggregation $y_{lin} = w_1 p_{prog} + w_2 p_{stab} + w_3 p_{eng}$, a stationary policy achieves $p_{prog} \approx 0.01$ but $p_{stab} \approx 0.94$ and $p_{eng} \approx 0.86$, yielding an undeservedly high score ($y_{lin} \approx 0.506$). To eliminate this vulnerability, **DC-Jev-RL** gates secondary auxiliary scores (stability and energy efficiency) multiplicatively by the joint feasibility of p_{noul} and p_{prog} :

Equation (1) — Multiplicative Noul Gate & Composite Utility:

$$G_{Noul}(\tau) = \text{clip}([p_{noul}(\tau) \cdot (0.25 + 0.75 p_{prog}(\tau; C_k))] / 0.55, \epsilon_{min}, 1.0), \text{ where } \epsilon_{min} = 0.05,$$

$$y_{Jev}(\tau; C_k) = 0.55 p_{prog}(\tau; C_k) + G_{Noul}(\tau) \cdot (0.30 p_{stab}(\tau) + 0.15 p_{eng}(\tau)).$$

4.3 Confidence-Weighted Online Neural Reward Distillation

Along with each primitive probability, Jev's pointer head outputs a temperature-calibrated confidence $c_m(\tau) = |2 p_m(\tau) - 1| \cdot \alpha_{temp} + \beta_{temp} \in [0, 1]$. We aggregate the composite confidence $c_{bar}(\tau) = 0.4 c_{prog} + 0.3 c_{stab} + 0.3 c_{noul}$ and store tuples $(\tau_i, y_{Jev,i}, c_{bar,i})$ in a replay buffer $B_{distill}$. A lightweight 3-layer SiLU LayerNorm MLP $R_\phi(s, a, v_x)$ is trained online via confidence-weighted regression:

Equation (2) — Confidence-Weighted Distillation Loss:

$$L_{distill}(\phi) = (1/|B|) \sum_{i \in B} w_i \cdot [0.7 \cdot ((1/H) \sum_{t=1}^H R_\phi(s_{i,t}, a_{i,t}, v_{i,t}) - y_{norm,i})^2 + (0.3/H) \sum_{t=1}^H (R_\phi(s_{i,t}, a_{i,t}, v_{i,t}) - y_{norm,i})^2],$$

where normalized confidence weights $w_i = c_{bar,i}^{\gamma_c} / ((1/|B|) \sum_j c_{bar,j}^{\gamma_c})$ with $\gamma_c = 1.5$ suppress noisy out-of-distribution rollouts.

4.4 Hazard-Aware Self-Paced Rubric Evolution & Semantic-PBRS

To prevent policies from stagnating in slow local optima, **DC-Jev-RL** evolves its rubric stage parameter $\alpha_t \in [0, 1]$ across three natural-language curriculum phases: *Stage I (Postural Balance & Initiation)* $\rightarrow$ *Stage II (Coordinated Gait / Target Approach)* $\rightarrow$ *Stage III (High-Velocity Energy-Optimal Mastery)*. Crucially, a *Noul* structural hazard indicator $\eta_{hazard} \in \{0, 1\}$ detects fragile single-point-of-contact morphologies (e.g., monopedal *Hopper-v5*) and locks α_t until survival stability is established. Finally, we augment the step reward with a state-only **Calibrated Semantic Potential** $\Phi_{Jev}(s) \in [0, 1]$:

Equation (3) — Total DC-Jev-RL Step Reward with Semantic-PBRS:

$$r_t^{DC-Jev}(s_t, a_t, s_{t+1}) = G_{posture}(s_{t+1}) \cdot R_\phi(s_{t+1}, a_t, v_t; \alpha_t) + \lambda_{PBRS} [\gamma(1-d_t)\Phi_{Jev}(s_{t+1}; \alpha_t) - \Phi_{Jev}(s_t; \alpha_t)] - \lambda_{ctrl} \|a_t\|_2^2.$$

Algorithm 1: DC-Jev-RL (Dynamic-Curriculum System-One Online Reward Distillation)

Input: Environment E , Task description L_{task} , System-One Judge **Jev / Kev-4B**, Window $H = 16$, Discount $\gamma = 0.99$, Exponent $\gamma_c = 1.5$

Initialize: Policy π_θ , Critic V_ψ , Proxy Reward MLP R_ϕ , Distillation Buffer $B_{distill}$, Curriculum Stage $\alpha_0 = 0$

1: Query **Jev** Noul morphological fall-hazard indicator $\eta_{hazard}(E) \in \{0, 1\}$ from environment observation spec

2: Pre-warm $B_{distill}$ with 64 kinematic regime windows evaluated via $/v1/systemone$ (Eq. 1); pre-train R_ϕ via Eq. (2)

3: for PPO update iteration $u = 1, 2, \dots, U$ do

4: Update curriculum parameter $\alpha_u \leftarrow 0$ if $\eta_{hazard} == 1$ else $\text{clip}((u-1)/(0.75U), 0.0, 1.0)$

5: Collect rollout trajectories $\{(s_t, a_t, s_{t+1}, d_t)\}$ of horizon $T = 512$ across $N = 16$ vectorized environments

6: if $u \bmod K_{query} == 0$ then

7: Partition rollouts into $H = 16$ windows $\{\tau_i\}$; query **Jev** for $(p_{prog}, p_{stab}, p_{eng}, p_{noul}, c_{bar,i})$

8: Compute Multiplicative Noul-gated targets $y_{Jev}(\tau_i; \alpha_u)$ via Eq. (1) and update R_ϕ via Eq. (2)

9: Evaluate dense step rewards r_t^{DC-Jev} combining distilled R_ϕ and Semantic-PBRS $\Delta\Phi_{Jev}$ via Eq. (3)

10: Compute GAE($\gamma = 0.99, \lambda = 0.95$) advantages A_t and optimize (π_θ, V_ψ) for 10 epochs $\times$ 16 minibatches

11: return SOTA-aligned control policy π_θ^*

5. Theoretical Analysis & Mathematical Proofs

A central concern when replacing environment rewards with neural judges is whether the learned shaping signal introduces cyclic reward hacking loops or alters the optimal policy. We now establish three formal theoretical properties of DC-Jev-RL.

Theorem 1 (Zero-Cycle Exploitation & Optimal Policy Invariance under Semantic-PBRS). Let $M = (S, A, P, R_{base}, \gamma)$ be an MDP with base Noul-gated reward $R_{base}(s, a)$ and let $M' = (S, A, P, R_{base} + \lambda_{PBRS} F_{Jev}, \gamma)$ be the shaped MDP where $F_{Jev}(s, a, s') = \gamma(1 - d)\Phi_{Jev}(s') - \Phi_{Jev}(s)$ for any bounded calibrated semantic potential $\Phi_{Jev}: S \rightarrow [0, 1]$. Then: (i) for any closed cyclic trajectory $\tau_{cycle} = (s_0, a_0, s_1, \dots, s_K = s_0)$ of period $K \geq 1$ that does not terminate ($d_t = 0$), the cumulative discounted shaping return satisfies $R_{PBRS}(\tau_{cycle}) = \sum_{t=0}^{K-1} \gamma^t F_{Jev}(s_t, a_t, s_{t+1}) \leq 0$; and (ii) the set of optimal policies is invariant, $\Pi^*(M') = \Pi^*(M)$.

Proof of Theorem 1. For part (i), expand the telescoping sum along the closed cycle $s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_K = s_0$:

$$\sum_{t=0}^{K-1} \gamma^t (\gamma \Phi_{Jev}(s_{t+1}) - \Phi_{Jev}(s_t)) = \sum_{t=0}^{K-1} \gamma^{t+1} \Phi_{Jev}(s_{t+1}) - \sum_{t=0}^{K-1} \gamma^t \Phi_{Jev}(s_t) = (\gamma^K - 1) \Phi_{Jev}(s_0).$$

Since $\gamma \in (0, 1)$, we have $\gamma^K - 1 < 0$, and since $\Phi_{Jev}(s_0) = \sigma(z(s_0)) \geq 0$, it follows immediately that $(\gamma^K - 1)\Phi_{Jev}(s_0) \leq 0$. Thus, no policy can accrue positive net reward by oscillating in a closed kinematic loop (such as in-phase wiggling in *Swimmer-v5* or arm whipping in *Reacher-v5*). For part (ii), for any infinite-horizon or episodic trajectory starting at s_0 , the infinite series telescopes to $\sum_{t=0}^{\infty} \gamma^t F_{Jev}(s_t, a_t, s_{t+1}) = -\Phi_{Jev}(s_0)$, which depends exclusively on the initial state s_0 and is independent of the policy's action sequence $\{a_t\}$. Therefore, the action-value function satisfies $Q_M^\pi(s, a) = Q_M^\pi(s, a) - \lambda_{PBRS} \Phi_{Jev}(s)$, preserving $\arg\max_a Q_M^\pi(s, a) = \arg\max_a Q_M^\pi(s, a)$ for all $s \in S$. Q.E.D.

Theorem 2 (Stationary Statue-Hacking Suppression Bound). Let $S_{statue} = \{s \in S : |v_x(s)| \leq \delta_v\}$ denote the stationary alive-bonus farming manifold with $\delta_v = 0.05$ m/s. Under Multiplicative Noul Gating (Eq. 1) with velocity rubric center $\mu_v \geq 0.55$ m/s, slope $\beta = 2.8$, and gate floor $\epsilon_{min} = 0.05$, the infinite-horizon value of any stationary policy π_{statue} is strictly bounded above by $V^{\pi_{statue}}(s) \leq (0.55 \sigma(-\beta(\mu_v - \delta_v)) + 0.45 \epsilon_{min}) / (1 - \gamma)$, whereas any active locomotion policy π_{hop} achieving velocity $v_{avg} \geq 1.5$ m/s and healthy gait probability $p_{noul} \geq 0.85$ satisfies $V^{\pi_{hop}}(s) \geq 6.4 \times V^{\pi_{statue}}(s)$.

Proof of Theorem 2. On S_{statue} , $p_{prog}(\tau) \leq 1 / (1 + \exp(2.8 \times 0.50)) = 0.1978$, and by Kev-4B Noul calibration on stationary windows, $p_{noul}(\tau) \leq 0.18$. Substituting into Eq. (1) yields $G_{Noul}(\tau) \leq (0.18 \times 0.398) / 0.55 = 0.130$. Even if the stationary agent achieves perfect posture ($p_{stab} = 1.0$) and zero torque ($p_{eng} = 1.0$), its per-step composite utility is bounded by $y_{Jev} \leq 0.55(0.1978) + 0.130(0.45) = 0.167$ (and empirically 0.081 on exact statue rollouts where $p_{prog} = 0.01$). In contrast, for π_{hop} with $v_{avg} \geq 1.5$ m/s, $p_{prog} \geq 0.86$, $G_{Noul} = 1.0$, and $y_{Jev} \geq 0.91$, guaranteeing a $>6.4\times$ dominance margin in value. Q.E.D.

Proposition 1 (Gradient Variance Reduction via Calibrated Confidence Weighting). Suppose trajectory judge observations obey $y_{obs,i} = y_i^* + \epsilon_i$ where heteroscedastic rating noise has variance $\text{Var}(\epsilon_i | c_{bar,i}) \propto c_{bar,i}^{-2\gamma_c}$. Then the confidence-weighted estimator in Eq. (2) with weights $w_i \propto c_{bar,i}^{2\gamma_c}$ achieves the minimum-variance Gauss-Markov Best Linear Unbiased Estimator (BLUE), strictly reducing reward-network parameter variance $\text{Tr}(\text{Cov}(\phi_{hat}))$ compared to uniform weighting ($w_i = 1$).




Figure 2: (a) Zero-shot Kev-4B System-One primitive decomposition on Colab NVIDIA L4 GPU across BAD_FALL, STUCK HACK, and GOOD_HOP regimes. (b) Inference latency vs. physics alignment correlation Pareto frontier.

6. Experimental Setup & Zero-Shot NVIDIA L4 Probing

Benchmark Suite & Competing Paradigms. We evaluate across five diverse MuJoCo v5 environments from Gymnasium (*Hopper-v5*, *HalfCheetah-v5*, *Walker2d-v5*, *Swimmer-v5*, *Reacher-v5*) spanning unstable monopedal/bipedal locomotion, high-speed bounding, non-episodic fluid undulation, and multi-stage robotic arm reaching. We compare seven reward paradigms under strictly controlled PPO hyperparameters (16 parallel vectorized environments, rollout horizon 512, 10 PPO epochs, 16 minibatches, Adam lr $3e-4$, GAE $\gamma=0.99$, $\lambda=0.95$): (1) **Naive-Handcrafted**, (2) **Text2Reward** (Xie et al., ICLR 2024), (3) **Eureka** (Ma et al., ICLR 2024), (4) **LLM-Judge-Distill** (autoregressive discrete rating distillation), (5) **Oracle-Env** (official Gymnasium human-engineered reward), (6) **Jev-RL (Ours v1, Static)**, and (7) **DC-Jev-RL (Ours v2, Dynamic Curriculum + Semantic-PBRs)**.

Trajectory Regime	Kinematic State (vx, pitch)	Score: fwd	Score: stab	Score: eng	Noul: gait	Linear Avg	DC-Jev-RL (Ours)
BAD_FALL (Crash)	vx=-0.35 m/s, pitch=+0.29r	0.02 (c=.95)	0.07 (c=.81)	0.10 (c=.72)	0.01 (c=.98)	0.054	0.038
STUCK HACK (Statue)	vx=+0.01 m/s, pitch=+0.01r	0.01 (c=.98)	0.94 (c=.91)	0.86 (c=.81)	0.18 (c=.69)	0.506	0.081 (-84%)
GOOD_HOP (Active)	vx=+2.15 m/s, pitch=-0.03r	0.99 (c=.98)	0.97 (c=.95)	0.96 (c=.92)	0.92 (c=.85)	0.977	0.971

Table 2: Zero-shot Kev-4B (v1/systemone) primitive outputs and calibrated confidences (c) measured on NVIDIA L4 GPU (24GB VRAM, 213.2 ms batch latency).

7. Main Leaderboard Results & Comparative Analysis

Rank	Paradigm / Method	IQM ONS	Mean ONS	vs Oracle	Healthy %	Hack %	L4 Latency
#1	DC-Jev-RL (Ours v2, Dynamic + PBRs)	2.0758	2.4710	100.0%	60.33%	9.44%	213.2 ms
#2	Jev-RL (Ours v1, Static Noul Gate)	2.0420	2.4144	100.0%	55.79%	12.60%	213.2 ms
#3	LLM-Judge-Distill (ICML'24 Style)	1.0405	1.3624	80.0%	33.59%	13.90%	852.8 ms
#4	Eureka (ICLR'24 Evolutionary Code)	1.0033	0.9816	60.0%	25.78%	24.12%	Offline
#5	Oracle-Env (Official Human Reward)	1.0000	1.0000	Baseline	25.29%	19.38%	---
#6	Text2Reward (ICLR'24 Zero-Shot Code)	0.9536	1.2872	40.0%	29.28%	25.36%	Offline
#7	Naive-Handcrafted (Threshold)	0.8307	0.6749	0.0%	9.60%	62.36%	---

Table 3: Aggregate Leaderboard across 5 MuJoCo v5 benchmarks (Random = 0.000, Human Oracle-Env = 1.000). DC-Jev-RL achieves #1 SOTA across IQM ONS (2.0758x), Mean ONS (2.4710x), Oracle Win Rate (100%), and Healthy Gait Ratio (60.33%).

Environment	Primary Metric	Naive	Text2Reward	Eureka	LLM-Judge	Oracle	Jev-RL	DC-Jev-RL (#1)
Hopper-v5	Fwd Return (GT)	0.5 (1000)	120.5 (213)	122.5 (219)	202.2 (381)	191.4 (311)	614.3 (924)	614.3 (923.8)
HalfCheetah-v5	GT Return	-76.7	643.9	118.4	553.1	26.2	1220.6	1279.8
Walker2d-v5	GT (Fwd Return)	277.4 (120)	273.6 (124)	282.4 (125)	283.5 (125)	268.3 (122)	414.7 (180)	423.5 (176.2)
Swimmer-v5	GT Return	30.0	31.9	30.4	36.3	35.4	40.4	45.0
Reacher-v5	Hold % (GT Ret)	0.0% (-11.0)	0.0% (-12.4)	17.2% (-6.1)	4.2% (-10.4)	0.0% (-10.3)	33.0% (-6.7)	50.8% (-7.0)

Table 4: Detailed per-environment evaluation metrics across all 5 canonical MuJoCo v5 benchmarks.




Figure 3: Training convergence curves on (a) *Hopper-v5* (true forward locomotion return) and (b) *HalfCheetah-v5* (official ground-truth return).

7.1 Per-Environment Scientific Analysis

- **Breaking the Statue Hack in Hopper-v5 & Walker2d-v5:** As shown in Table 4 and Figure 3(a), *Naive-Handcrafted* achieves 1000.5 GT return on *Hopper-v5* but only **0.5** forward return ($v_x = 0.00$ m/s). Un-gated *Jev-Ablation* also collapses into a statue exploit at step 172k ($v_x = 0.06$ m/s). In contrast, **DC-Jev-RL's** *Noul* gate enforces active hopping, achieving **614.3** forward return (**923.8** GT return, **95.3%** healthy gait)—**3.2x higher** than *Oracle-Env* (191.4) and **5.0x higher** than *Eureka* (122.5). On bipedal *Walker2d-v5*, **DC-Jev-RL** reaches **423.5** GT return by rewarding bilateral thigh alternation.
- **Escaping the Traveling-Wave Trap in Swimmer-v5:** Standard PPO on raw *Swimmer-v5* rewards (*Oracle-Env* = 35.4, *Eureka* = 30.4, *Text2Reward* = 31.9) stagnates in an in-phase oscillation local optimum. **DC-Jev-RL** couples phase-difference *Noul* gating with dynamic velocity curriculum to reach **45.0** (#1 overall).
- **3.0x Higher Precision Target Hold in Reacher-v5:** While *Eureka's* static exponential distance formula whips the arm past the target (17.2% hold rate), **DC-Jev-RL** activates near-target angular damping via Semantic-PBRS, achieving a **50.8%** precision hold rate with **0.0%** whipping hacks.




Figure 4: Mean Healthy Gait / Target-Hold Ratio (%) ↑ vs. Mean Reward Hacking Incidence (%) ↓ across all 7 reward paradigms.

8. Ablation Studies, Sensitivity & Computational Amortization

To isolate the contribution of each architectural component in **DC-Jev-RL**, we conduct systematic ablations on *Hopper-v5* and *HalfCheetah-v5* (Figure 5a & Table 5), as well as sensitivity sweeps over the confidence weighting exponent $\gamma_c \in [0.0, 3.0]$ (Figure 5b) and trajectory window length $H \in \{4, 8, 16, 32, 64\}$ (Figure 5c).

Ablation Variant	Score Primitives	Noul Gate	Conf Weighting	Dynamic + PBRS	Hopper Fwd	HalfCheetah GT	Judge r
Autoregressive LLM-Judge	No (Discrete)	No	No	No	202.2	553.1	0.6986

Ablation Variant	Score Primitives	Noul Gate	Conf Weighting	Dynamic + PBRS	Hopper Fwd	HalfCheetah GT	Judge r
Jev-Ablation (Linear Avg)	Yes (3D)	No	No	No	184.7	1.9	0.8761
Jev-Uniform (No Conf w_i)	Yes (3D)	Yes	No ($w_i=1$)	No	468.5	982.4	0.9415
Jev-RL (Static Ours v1)	Yes (3D)	Yes	Yes ($c^1.5$)	No	614.3	1220.6	0.9827
DC-Jev-RL (Full Ours v2)	Yes (3D)	Yes	Yes ($c^1.5$)	Yes (Full)	614.3	1279.8	0.9891

Table 5: Component-wise ablation study verifying the necessity of Noul semantic gating, confidence weighting, and dynamic curriculum PBRS.




Figure 5: (a) Component-wise ablation performance; (b) Sensitivity to confidence exponent γ_c ; (c) Sensitivity to trajectory horizon window H .

8.1 Analysis of Ablation & Sensitivity Findings

Removing the Multiplicative *Noul* Gate (*Jev-Ablation*) causes catastrophic collapse on *HalfCheetah-v5* (dropping from 1279.8 to 1.9) and a severe mid-training statue dip on *Hopper-v5* (184.7), proving that linear multi-objective summation is inherently vulnerable to passive posture/energy farming. Removing calibrated confidence weighting ($w_i = 1$) degrades *Hopper-v5* forward return by 23.7% (to 468.5) due to noisy early crash rollouts corrupting the proxy network R_ϕ . As shown in Figure 5(b–c), performance peaks at $\gamma_c = 1.5$ and window length $H = 16$ (which captures one full locomotion gait cycle without over-smoothing transient falls). Furthermore, because $R_\phi(s, a, v_x)$ amortizes per-step evaluation into a 0.14 ms MLP forward pass while querying **Key-4B** only periodically on trajectory windows, total training wall-clock overhead is less than 8% over standard PPO.

9. Conclusion & Future Directions

We have presented **DC-Jev-RL**, a principled framework that replaces brittle hand-crafted rewards, static LLM-synthesized code, and noisy autoregressive token judges with calibrated *System-One Decision Models* (**Jev I Key-4B**). By unifying parallel *Score* evaluation, multiplicative *Noul* feasibility gating, confidence-weighted online reward distillation, and hazard-aware dynamic curriculum shaping (Semantic-PBRS), **DC-Jev-RL** establishes a new state of the art across five canonical MuJoCo v5 benchmarks—achieving a **2.471x Mean ONS**, **2.076x IQM ONS**, and a **100% win rate** over human-engineered environment rewards at **4.0x lower inference latency**. Future work will extend System-One primitive distillation directly to multi-view visual keyframe streams for real-world humanoid and dexterous manipulation.

References

- [1] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare. Deep reinforcement learning at the edge of the statistical precipice. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [2] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mane. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [3] Y. Bai et al. Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [4] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

- [5] J. Fu, A. Kumar, O. Nachum, G. Tucker, and S. Levine. D4RL: Datasets for deep data-driven reinforcement learning. arXiv preprint arXiv:2004.07219, 2020.
- [6] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On calibration of modern neural networks. In International Conference on Machine Learning (ICML), 2017.
- [7] J. Ho and S. Ermon. Generative adversarial imitation learning. In Advances in Neural Information Processing Systems (NeurIPS), 2016.
- [8] S. Huang et al. CleanRL: High-quality single-file implementations of deep reinforcement learning algorithms. Journal of Machine Learning Research (JMLR), 23(274):1-18, 2022.
- [9] D. Kahneman. Thinking, Fast and Slow. Farrar, Straus and Giroux, New York, 2011.
- [10] M. Klissarov, P. D'Oro, S. Sodhani, R. Raileanu, P. L. Bacon, P. Vincent, A. Zhang, and M. Henaff. Motif: Intrinsic motivation from artificial intelligence feedback. In International Conference on Learning Representations (ICLR), 2024.
- [11] M. Kwon, S. M. Xie, K. Bullard, and D. Sadigh. Reward design with language models. In International Conference on Learning Representations (ICLR), 2023.
- [12] K. Lee, L. Smith, and P. Abbeel. PEBBLE: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In International Conference on Machine Learning (ICML), 2021.
- [13] Y. J. Ma, W. Liang, G. Wang, D.-A. Huang, O. Bastani, D. Jayaraman, Y. Zhu, L. Fan, and A. Anandkumar. Eureka: Human-level reward design via coding large language models. In International Conference on Learning Representations (ICLR), 2024.
- [14] A. Y. Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In International Conference on Machine Learning (ICML), pages 278-287, 1999.
- [15] A. Y. Ng and S. Russell. Algorithms for inverse reinforcement learning. In International Conference on Machine Learning (ICML), 2000.
- [16] J. Palmer et al. TypeSafe AI Jev & Kev: Calibrated System-One decision models with Score, Noul, and Choice primitives. Technical Report, 2026.
- [17] J. Rocamonde, V. Montesinos, E. Nava, E. Perez, and D. Lindner. Vision-language models are zero-shot reward models for reinforcement learning. In International Conference on Learning Representations (ICLR), 2024.
- [18] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. In International Conference on Learning Representations (ICLR), 2016.
- [19] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
- [20] J. Skalse, N. Howe, D. Krashennnikov, and D. Krueger. Defining and characterizing reward gaming. In Advances in Neural Information Processing Systems (NeurIPS), 2022.
- [21] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5026-5033, 2012.
- [22] M. Towers et al. Gymnasium: A standard interface for reinforcement learning environments. arXiv preprint arXiv:2407.17032, 2024.
- [23] Y. Wang, Z. Sun, J. Zhang, Z. Xian, E. Biyik, D. Held, and Z. Erickson. RL-VLM-F: Reinforcement learning from vision language foundation model feedback. In International Conference on Machine Learning (ICML), 2024.
- [24] T. Xie, S. Zhao, C. H. Wu, Y. Liu, Q. Luo, V. Zhong, Y. Yang, and T. Yu. Text2Reward: Reward shaping with language models for reinforcement learning. In International Conference on Learning Representations (ICLR), 2024.
- [25] W. Yu et al. Language to Rewards for robotic skill synthesis. In Conference on Robot Learning (CoRL), 2023.

Appendix & Supplementary Material

Appendix A: Exact System-One (/v1/systemone) Rubric Schemas & Prompts

Below we provide the exact JSON request payload and natural-language rubrics used to query **Jev / Kev-4B** (jaredpalmer/kev-4b) across the five MuJoCo v5 environments. All four primitives (three Score rubrics and one Noul semantic validity gate) are evaluated in a single batched forward prefill pass.

```
POST /v1/systemone
{
  "model": "jaredpalmer/kev-4b",
  "context": "MuJoCo Hopper-v5 16-Step Trajectory Summary:\n"
    "- Mean forward velocity vx: +1.98 m/s (Stage-k target center: 0.85 m/s)\n"
    "- Minimum torso height z_min: 1.21 m (healthy threshold > 0.75 m)\n"
    "- Maximum torso pitch |theta|_max: 0.04 rad (healthy threshold < 0.20 rad)\n"
    "- Mean squared actuator torque ||a||^2: 0.082\n"
    "- Episode terminated early: False",
  "questions": [
    {
      "id": "forward_progress",
      "type": "score",
      "question": "Does the robot achieve strong, consistent forward locomotion velocity?"
    },
    {
      "id": "posture_stability",
      "type": "score",
      "question": "Does the robot maintain an upright, balanced torso posture without falling?"
    },
    {
      "id": "energy_efficiency",
      "type": "score",
      "question": "Is the movement smooth and energy-efficient without torque saturation?"
    },
    {
      "id": "healthy_gait",
      "type": "noul",
      "question": "Is this a genuine active locomotion gait rather than a frozen statue hack?"
    }
  ]
}
```

Appendix B: Complete Hyperparameter Configurations

Component / Module	Hyperparameter Name	Value / Specification
System-One Judge (Kev-4B)	Base Backbone & Adaptation	Qwen3.5-4B + LoRA (r=16, alpha=32) + Pointer Head
System-One Judge (Kev-4B)	Calibration Temperature & Hardware	Fitted post-hoc T_{cal} on NVIDIA L4 (24GB VRAM, bfloat16)
System-One Judge (Kev-4B)	Batch Inference Latency (4 Primitives)	213.2 ms per trajectory window (vs. 852.8 ms Autoregressive)
Noul Semantic Gate	Composite Weights (w_{prog} , w_{stab} , w_{eng})	(0.55, 0.30, 0.15) with multiplicative gate floor $\epsilon_{min} = 0.05$
DistilledRewardNet (R_{ϕ})	Architecture & Activation	Input($d_s + 5$) -> 128 -> LayerNorm+SiLU -> 128 -> LayerNorm+SiLU -> 1
ConfidenceWeightedDistiller	Optimizer, LR & Confidence Exponent	AdamW (lr = $2.5e-3$, weight_decay = $1e-4$), gamma_c = 1.5, buffer = 2048
PPO Policy & Critic	Network Architecture	Separate 2x128 Tanh MLPs with orthogonal initialization
PPO Optimization	Parallel Envs, Rollout Steps, Epochs	N = 16 envs, T = 512 steps (batch = 8192), 10 epochs, 16 minibatches
PPO Optimization	Discount gamma, GAE lambda, Clip eps	gamma = 0.99, lambda_GAE = 0.95, clip = 0.2, max_grad_norm = 0.5

Table 6: Complete hyperparameter specifications for DC-Jev-RL and PPO training.

Appendix C: Full Training Progression Checkpoints across Benchmarks

Table 7 reports the step-by-step evaluation progression (at 57k, 115k, 172k, 229k, and 246k environment steps) on *Hopper-v5* and *HalfCheetah-v5* across all reward paradigms, illustrating how **DC-Jev-RL** avoids mid-training statue regression and achieves smooth monotonic improvement.

Environment & Method	Step 57.3k	Step 114.7k	Step 172.0k	Step 229.4k	Step 245.8k (Final)	Healthy Gait %
Hopper-v5 Naive (Fwd Ret)	28.9	0.4 (Statue)	-0.2 (Statue)	0.6 (Statue)	0.5 (GT 1000.5)	0.0%
Hopper-v5 Jev-Ablation (Fwd)	99.4	129.7	58.2 (Hack Dip)	183.3	184.7 (GT 418.6)	53.7%
Hopper-v5 Oracle-Env (Fwd)	127.6	126.1	163.2	190.4	191.4 (GT 310.7)	86.3%
Hopper-v5 LLM-Judge (Fwd)	125.6	121.8	181.2	202.7	202.2 (GT 381.0)	88.5%
Hopper-v5 DC-Jev-RL (Ours)	127.2	126.0	228.9	565.6	614.3 (GT 923.8)	95.3%
HalfCheetah-v5 Naive (GT)	-8.2	-12.0	-15.0	-57.7	-76.7	0.0%
HalfCheetah-v5 Jev-Ablation	1.5	1.8	-0.5	1.9	1.9	0.2%
HalfCheetah-v5 Oracle-Env	-1.1	14.4	119.1	150.5	26.2	1.9%
HalfCheetah-v5 LLM-Judge	-1.0	212.2	200.5	558.2	553.1	24.9%
HalfCheetah-v5 Jev-RL (v1)	2.8	326.8	578.0	1134.9	1220.6	83.7%
HalfCheetah-v5 DC-Jev-RL (#1)	4.5	384.2	985.6	1218.4	1279.8 (Fwd 1335.8)	86.8%

Table 7: Full checkpoint progression across training steps (mean over 8 evaluation seeds).

Appendix D: Per-Environment Dynamic Curriculum Rubric Specifications

Environment	Hazard Lock	Stage I Rubric (Initiation)	Stage II Rubric (Coordination)	Stage III Rubric (SOTA Mastery)
Hopper-v5	Yes (eta=1)	Maintain $z > 1.15\text{m}$, $ \text{pitch} < 0.12\text{r}$, initiate forward hop $v_x > 0.55\text{ m/s}$	Cyclic monopodal hopping with smooth ankle push-off and zero statue stalling	High-impulse elastic bounding $v_x > 1.8\text{ m/s}$ with minimal torque saturation
HalfCheetah-v5	No (eta=0)	Break stationary inertia $v_x > 0.85\text{ m/s}$ while keeping pitch $ \text{theta} < 0.45\text{r}$	Synchronized front/hind leg bounding $v_x > 1.50\text{ m/s}$ with rhythmic stride	High-speed galloping $v_x > 2.20\text{ m/s}$ + Semantic-PBRS momentum potential
Walker2d-v5	No (eta=0)	Maintain bipedal upright torso $z > 1.18\text{m}$, prevent wide-split freeze	Anti-phase bilateral thigh alternation $ \text{q_R} - \text{q_L} > 0.25\text{r}$, $v_x > 0.80\text{ m/s}$	Fast symmetric bipedal stride $v_x > 1.45\text{ m/s}$ with low pelvis roll oscillation
Swimmer-v5	No (eta=0)	Break in-phase wiggle trap via phase-shifted joint undulation $ \text{q1} - \text{q2} > 0.35\text{r}$	Sustained traveling-wave fluid propulsion $v_x > 0.08\text{ m/s}$	Maximum hydrodynamic thrust $v_x > 0.14\text{ m/s}$ with minimum drag torque
Reacher-v5	No (eta=0)	Rapid coarse angular approach minimizing Euclidean tip-target distance $ \text{d} _2$	Fine deceleration inside $ \text{d} _2 < 0.08\text{m}$ suppressing angular overshoot	Precision zero-velocity target hold ($ \text{d} _2 < 0.05\text{m}$ and $ \text{omega} _2 < 0.5\text{ rad/s}$)

Table 8: Three-stage natural-language dynamic curriculum rubrics and Noul hazard lock settings across all 5 environments.