

MH-FLOCKE: Biologically Grounded Embodied Cognition Through a 15-Step Closed-Loop Architecture for Quadruped Locomotion Learning

Revised March 2026 — addressing reviewer feedback on RL baselines, multi-seed validation, mathematical formulations, and cross-embodiment transfer

Marc Hesse

Independent Researcher, Potsdam, Germany
info@mhflocke.com

Abstract

We present MH-FLOCKE, an embodied AI platform in which simulated quadruped creatures learn locomotion through a biologically grounded cognitive architecture. Unlike end-to-end reinforcement learning (RL) approaches that treat the body as an optimization target, MH-FLOCKE implements a 15-step closed-loop processing cycle that integrates proprioception, embodied emotions, episodic memory, motivational drives, a Global Workspace for attentional competition, metacognitive self-assessment, and reward-modulated spike-timing-dependent plasticity (R-STDP) in a spiking neural network (SNN). In this revision, we address reviewer feedback by providing: (1) mathematical formulations of all core learning rules (Izhikevich dynamics, R-STDP, cerebellar forward model, competence gate), (2) a PPO baseline comparison on the same morphology showing MH-FLOCKE's neural learning core achieves 3.5x the walking distance (45.15 ± 0.67 m vs. 12.83 ± 7.78 m at 50k steps), (3) multi-seed statistical validation across 10 seeds and 80 runs with mean \pm std for all metrics, and (4) cross-embodiment transfer to the Unitree Go2 quadruped robot without architectural changes. Systematic ablation across 60+ runs on the original morphology and 80 runs on Go2 isolates the contribution of each component: vestibular reflexes eliminate all falls, motor babbling increases flat-terrain distance by 763%, the cerebellar forward model produces measurable prediction errors, and an olfactory sensory environment enables stimulus-driven behavior switching. The SNN+Cerebellum core (B1) achieves the lowest variance of any condition ($\sigma = 0.67$ m across 10 seeds), demonstrating highly consistent learning. We report an interaction effect where the full cognitive architecture reduces locomotion distance compared to the neural learning core alone, and present this finding with analysis as a documented result. A companion video is available at <https://www.youtube.com/watch?v=Jo7UM6pEFMg>. The architecture and all ablation data are publicly available at <https://github.com/MarcHesse/mhflocke>.

1. Introduction

Quadruped locomotion is a benchmark problem in embodied AI, yet the dominant approach remains end-to-end reinforcement learning (RL), which discovers motor policies through reward maximization without explicit modeling of the biological subsystems that produce locomotion in animals. While these approaches achieve impressive results in simulation and sim-to-real transfer [1, 2], they conflate learning mechanisms that operate on fundamentally different timescales: spinal reflexes (milliseconds), cerebellar adaptation (seconds), behavior selection (minutes), and memory consolidation (hours).

Biological quadruped locomotion emerges from the interaction of at least six distinct neural subsystems: spinal Central Pattern Generators (CPGs) that produce rhythmic motor patterns [3], vestibular reflexes that maintain postural stability [4], cerebellar forward models that predict sensory consequences of motor commands [5], brainstem motor babbling that calibrates the sensorimotor map during development [6], motivational drives that select behaviors based on internal states [7], and sensory systems that provide external stimuli for goal-directed behavior [8]. In the intact animal, these systems operate in parallel, with higher layers modulating but not replacing lower ones.

MH-FLOCKE (named after the author's late dog) implements this biological hierarchy as a computational architecture. The central contribution is a 15-step closed-loop processing cycle (Section 3.8) that runs at every simulation timestep, integrating all subsystems into a single coherent perception-action-learning cycle. Unlike modular robotics architectures that communicate through message-passing interfaces, MH-FLOCKE's cognitive loop is deeply interleaved: emotions modulate learning rates (Step 4), memory retrieval influences behavior

selection (Step 5), and metacognitive self-assessment gates exploration (Step 8). The creature does not have a body and a brain that communicate through an interface; the brain is embodied.

We make the following contributions. While the individual components (CPGs, SNNs, cerebellar models, R-STDP) are well-established in their respective literatures, their integration into a single closed-loop architecture that operates at every simulation timestep is, to our knowledge, novel. Specifically, we contribute: (1) a systematic ablation study across 60+ runs (original morphology) and 80 runs (Unitree Go2) that quantifies the contribution of each biological subsystem with multi-seed statistical validation, (2) a PPO baseline comparison on the same morphology demonstrating 3.5x distance advantage with 11.6x lower variance at 50k steps across 10 seeds, (3) cross-embodiment transfer to a production quadruped robot (Unitree Go2) without architectural changes, (4) mathematical formulations of all core learning rules for reproducibility, and (5) an honest report of an interaction effect where sensory steering interferes with cerebellar learning. The MH-FLOCKE framework builds on Integrity-OS [9], an earlier system that achieved 99.1% hallucination reduction through neural-symbolic hybrid architectures. The complete source code is publicly available at <https://github.com/MarcHesse/mhflocke> under the Apache 2.0 license.

2. Related Work

2.1 RL-Based Locomotion

Deep RL has achieved robust quadruped locomotion in simulation [1] and sim-to-real transfer [2, 10]. These approaches learn end-to-end policies mapping observations to joint torques, achieving high performance on locomotion benchmarks. However, the learned policies are opaque: it is unclear which biological principles (if any) the network has discovered, and the policies do not generalize to non-locomotion behaviors without retraining. In Section 5.6, we provide a direct comparison with PPO [28] on the same morphology, demonstrating that while PPO achieves stable but limited locomotion, MH-FLOCKE produces substantially greater distance with richer behavioral repertoires.

2.2 Central Pattern Generators

CPG models for robot locomotion have a long history [3, 11]. Oscillator networks produce stable gaits without sensory feedback, and can be modulated by descending signals for speed and gait transitions. MH-FLOCKE uses CPGs as the foundation layer, with a competence-gated transition from CPG-dominated to actor-dominated control as the creature learns (Section 3.2).

2.3 Cerebellar Forward Models

The cerebellum is hypothesized to implement forward models that predict sensory consequences of motor commands [5, 12]. Computational models typically use Purkinje cell supervised learning with climbing fiber error signals. Our implementation (Section 3.4) uses a parallel fiber to Purkinje cell pathway with prediction-error-driven climbing fiber pulses. The mathematical formulation is provided in Section 3.9.

2.4 Embodied Cognition and Consciousness

Global Workspace Theory (GWT) [13, 14] proposes that consciousness arises from competition among specialized modules for access to a shared broadcast. MH-FLOCKE implements GWT as Step 7 of the cognitive loop, with sensory, motor, predictive, error, and memory modules competing for broadcast, modulated by emotional valence and motivational drives. The Perturbational Complexity Index (PCI) [15] provides a quantitative complexity metric. While we make no claims about machine consciousness, the PCI value provides a measurable correlate of architectural complexity that distinguishes the full system from ablated versions.

3. Architecture

MH-FLOCKE consists of a simulated quadruped body (MuJoCo physics engine, 12 actuators in 4 legs) and a cognitive architecture organized in 9 packages: brain, body, behavior, bridge, integrity, llm, viz, self_improvement, and utils. The core computational substrate is a spiking neural network (SNN) with 5000+ neurons, Izhikevich dynamics [29], and four neuromodulators (DA, 5-HT, NE, ACh). The architecture is organized as nested closed

loops across timescales.

3.1 Spinal Reflexes

The fastest layer operates at every simulation timestep (200 Hz). Muscle tone maintains joint stiffness against gravity. Stretch reflexes resist perturbations proportional to joint displacement. Golgi tendon organ simulation limits excessive force to prevent self-damage. Crossed extension reflexes coordinate contralateral limbs during stumbling [16]. Reflexes are additive: they modulate motor output from higher layers but never replace it.

3.2 Central Pattern Generator

A phase-coupled oscillator network produces rhythmic gait patterns for walk and trot [3, 11]. The CPG provides a stable locomotion baseline that requires no learning. A competence gate (formalized in Eq. 4) blends CPG output with learned actor output: at training start, CPG weight is 90%; as the actor's velocity exceeds 0.03 m/s, the gate smoothly transitions to 40% CPG / 60% actor. This ensures the creature can walk from the first step while the actor learns to modulate and improve the innate pattern.

3.3 Motor Babbling

During the first 7,000 steps (babbling phase), the creature performs exploratory motor movements with reduced amplitude (70% of full range) and 25% additive noise [6]. This calibrates the sensorimotor map: the SNN learns the relationship between motor commands and sensory consequences. Motor babbling is the single most impactful component on flat terrain, increasing distance by 763% in ablation (from 1.50m to 12.95m at 50k steps on the original morphology). On hilly terrain, the CPG alone provides sufficient sensory variety, making babbling less critical.

3.4 Cerebellar Forward Model

A cerebellar learning module implements the Marr-Albus-Ito framework [5, 12]: parallel fiber inputs carry motor efference copies, Purkinje cells learn to predict sensory consequences, and climbing fiber error signals (triggered when prediction error exceeds a threshold) drive supervised learning. The deep cerebellar nuclei (DCN) output provides motor corrections blended with CPG and SNN output. The prediction error and learning rule are formalized in Eq. 3. A critical implementation detail: the forward model must be updated before the early-return gate in the training loop (Issue #71), otherwise the FM never receives training data during the CPG-dominated phase.

3.5 Vestibular System

Quaternion-derived upright estimation provides gravitational reference [4]. The vestibular signal gates cerebellar corrections: when the creature is falling (upright < 0.3), the correction magnitude is reduced to prevent the cerebellum from learning during transient states. This single mechanism eliminated all falls across every ablation configuration (27 falls in ablation #4 to 0 falls from ablation #7 onward).

3.6 Behavior Planning and Sensory Environment

A BehaviorPlanner selects from 8 behaviors (walk, trot, sniff, alert, rest, look_around, mark, motor_babbling) based on motivational drive state and sensory stimuli [7, 8]. The SensoryEnvironment module (v0.4.0) provides olfactory gradient targets and periodic acoustic events. Scent sources spawn 2-4m ahead of the creature and respawn when distance exceeds 5m (modeled as turbulent plume intermittency [17]). Olfactory steering modulates hip abduction asymmetrically to create gentle turning arcs toward scent sources [18, 19]. Sounds arrive 50% from scent direction (informative) and 50% from random directions (environmental noise). This replaces random behavior switching with stimulus-driven, biologically motivated behavior transitions.

3.7 The 15-Step Cognitive Loop

The central orchestrator is CognitiveBrain, which executes a 15-step closed loop at every simulation timestep. This loop integrates all subsystems into a coherent processing cycle from raw sensation through emotion, cognition, learning, and back to motor output. Each step corresponds to a distinct neural function:

Step 1 — SENSE: Raw proprioceptive and exteroceptive sensor parsing (height, upright, velocity, joint angles).

Step 2 — BODY SCHEMA: Efference copy comparison detects anomalies between expected and actual sensor states [20].

Step 3 — WORLD MODEL: Learned state predictor (MLP) generates prediction errors that drive curiosity and learning signals.

Step 4 — EMOTIONS: Embodied valence-arousal model derives emotional state from body signals. Somatic markers modulate neuromodulators [21].

Step 5 — MEMORY: Sensorimotor episodic memory records and retrieves experience sequences for prediction [22].

Step 6 — DRIVES: Motivational drives (survival, exploration, comfort, social) steer behavior selection and reward modulation.

Step 7 — GLOBAL WORKSPACE: Module competition for global broadcast, modulated by emotion and drives. Winner broadcasts to SNN hidden neurons [13, 14].

Step 8 — METACOGNITION: Self-assessment of confidence, consciousness level (PCI), learning progress, and module activity [15].

Step 9 — CONSISTENCY: Integrity checking inspired by anterior cingulate conflict monitoring. Detects prediction-body-memory conflicts [23].

Step 10 — COMBINED REWARD: Merges extrinsic reward with curiosity, empowerment, drive modulation, and emotional valence.

Step 11 — R-STDP LEARNING: Reward-modulated spike-timing-dependent plasticity (Eq. 2). Cerebellar populations are protected from R-STDP.

Step 12 — SYNAPTOGENESIS: SNN spike patterns form concept nodes in a knowledge graph. Astrocyte calcium gating regulates local plasticity [24, 25].

Step 13 — HEBBIAN LEARNING: Unsupervised coincidence learning captures correlations that reward signals alone cannot detect.

Step 14 — DREAM MODE: Periodic offline replay consolidating episodic memory into procedural knowledge via R-STDP [26].

Step 15 — NEUROMODULATION: DA (reward sensitivity), 5-HT (mood stability), NE (exploration noise), ACh (attention) adjusted from somatic markers.

This 15-step loop runs at every simulation timestep (200 Hz), creating a tight perception-action-learning cycle. Cognition is not a separate layer on top of motor control — it is deeply interleaved.

3.8 Nested Timescales

The architecture operates as nested closed loops across timescales. At the fastest timescale (every step): spinal reflexes maintain posture. At medium timescale (50-100 steps): the cerebellum corrects motor patterns via climbing fiber pulses. At slow timescale (1000+ steps): the BehaviorPlanner switches behaviors based on drives and stimuli. At the slowest timescale (dream intervals): episodic memory consolidates into procedural and conceptual knowledge via synaptogenesis. Each layer can override layers below it (metacognitive confidence gates exploration; consistency resets neuromodulators; drives override behavior) but the lower layers provide the foundation that makes higher cognition possible.

3.9 Mathematical Formulations

This section was added in revision to address reviewer request for mathematical detail crucial for reproducibility.

3.9.1 Izhikevich Neuron Model.

The SNN uses the Izhikevich neuron model [29] for computational efficiency with biologically realistic firing patterns. The membrane dynamics are governed by:

$$\begin{aligned}dv/dt &= 0.04v^2 + 5v + 140 - u + I \\du/dt &= a(bv - u) \\ \text{if } v \geq 30 \text{ mV, then } v &:= c, u := u + d\end{aligned}$$

where v is the membrane potential (mV), u is a recovery variable, I is the total synaptic input current, and (a, b, c, d) are parameters that determine the firing pattern. We use regular spiking (RS) parameters ($a=0.02, b=0.2, c=-65, d=8$) for excitatory neurons and fast spiking (FS) parameters ($a=0.1, b=0.2, c=-65, d=2$) for inhibitory neurons [29]. The network contains 5000+ neurons with an 80/20 excitatory/inhibitory ratio.

3.9.2 Reward-Modulated STDP (R-STDP).

Synaptic plasticity follows a three-factor learning rule combining spike-timing correlations with a global reward signal [30]. The eligibility trace for synapse (i, j) is:

$$de_{ij}/dt = -e_{ij}/\tau_e + \text{STDP}(\Delta t_{ij})$$

where $\tau_e = 1000$ ms is the eligibility trace time constant and the STDP kernel is:

$$\text{STDP}(\Delta t) = A^+ \exp(-\Delta t / \tau^+) \text{ if } \Delta t > 0 \text{ (pre before post)}$$

$$\text{STDP}(\Delta t) = -A^- \exp(\Delta t / \tau^-) \text{ if } \Delta t < 0 \text{ (post before pre)}$$

with $A^+ = 0.01, A^- = 0.012$ (slight LTD bias), $\tau^+ = \tau^- = 20$ ms. The weight update is modulated by the global reward signal $R(t)$:

$$\Delta w_{ij} = \eta \cdot R(t) \cdot e_{ij}$$

where $\eta = 0.001$ is the base learning rate. The reward signal $R(t)$ combines extrinsic reward (forward velocity, upright stability) with intrinsic components (curiosity from prediction error, empowerment, drive satisfaction, emotional valence) as computed in Step 10 of the cognitive loop. Cerebellar Purkinje cell populations are explicitly excluded from R-STDP to prevent interference with supervised forward model learning.

3.9.3 Cerebellar Forward Model.

The cerebellar module predicts the next sensory state $s_{\blacksquare}(t+1)$ from the current motor command $m(t)$ and proprioceptive state $s(t)$:

$$s_{\blacksquare}(t+1) = W_{PF} \cdot [m(t); s(t)]$$

$$e_{FM}(t) = || s(t+1) - s_{\blacksquare}(t+1) ||_2$$

When $e_{FM}(t)$ exceeds a threshold $\theta_{CF} = 0.01$, a climbing fiber pulse triggers weight updates:

$$\Delta W_{PF} = -\alpha_{CB} \cdot (s(t+1) - s_{\blacksquare}(t+1)) \cdot [m(t); s(t)]^T$$

with cerebellar learning rate $\alpha_{CB} = 0.005$. The deep cerebellar nuclei (DCN) output produces motor corrections $\Delta m(t)$ that are blended with CPG and SNN output. Vestibular gating reduces corrections when upright < 0.3 to prevent learning during falls.

3.9.4 Competence Gate.

The competence gate smoothly transitions motor control from CPG-dominated to actor-dominated:

$$w_{CPG}(t) = w_{max} - (w_{max} - w_{min}) \cdot \sigma(k \cdot (v_{actor}(t) - v_{thresh}))$$

where $w_{max} = 0.9, w_{min} = 0.4, k = 50$ is the sigmoid steepness, $v_{actor}(t)$ is the actor's achieved forward velocity, and $v_{thresh} = 0.03$ m/s is the competence threshold. $\sigma(x) = 1/(1+\exp(-x))$ is the logistic sigmoid. The final motor output is:

$$m_{out}(t) = w_{CPG}(t) \cdot m_{CPG}(t) + (1 - w_{CPG}(t)) \cdot m_{actor}(t) + \Delta m_{CB}(t)$$

where m_{CPG} is the CPG output, m_{actor} is the SNN actor output, and Δm_{CB} is the cerebellar correction. On the Go2 platform, dynamic PD scaling adjusts gains from 0.4 (standing) to 1.5 (fallen) to account for the heavier morphology.

4. Experimental Setup

4.1 Simulation Environment

All experiments use MuJoCo physics with a timestep of 5ms (200 Hz). Two morphologies are tested: (1) the original Bommel quadruped with 4 legs, 3 joints per leg (hip, knee, abduction), and 12 actuators, and (2) the Unitree Go2 quadruped robot model with identical joint topology but different mass distribution, link lengths, and actuator limits. Morphology is specified in MuJoCo XML and validated for mass distribution and joint limits. Two

terrain conditions are tested: flat (difficulty 0.0) and hilly (difficulty 0.3, procedurally generated). Training runs are 50,000 steps (4.2 minutes of simulated time at 200 Hz). Go2 runs use --auto-reset 500 because the Go2 cannot self-right after falls.

4.2 Ablation Design

We use a 3x2 ablation design crossing system complexity with terrain difficulty:

A (CPG only): Spinal CPG + reflexes + vestibular. No SNN, no cerebellum, no drives, no sensory environment. This is the biological floor — an anencephalic preparation.

B (CPG + SNN + Cerebellum): Adds SNN with R-STDP, actor-critic, drives, behavior planner, sensory environment, and cerebellar forward model. Tests the neural learning core.

C (Full system): All components including the complete 15-step cognitive loop with GWT, metacognition, dream mode, and synaptogenesis.

Each configuration is tested on flat (subscript 1) and hilly (subscript 2) terrain, yielding 6 conditions. On the original morphology (Bommel), ablation iterations #4, #7, #9, #10 capture the progression as components were added. On Go2, all 6 conditions plus a PPO baseline are run with **10 random seeds each** (60 biological + 10 PPO flat + 10 PPO hilly = **80 total runs**), reporting mean \pm standard deviation for all metrics. Seeds used: 7, 42, 55, 123, 256, 314, 808, 999, 1337, 2025.

4.3 PPO Baseline

Added in revision to address reviewer request for RL baseline comparison.

A Proximal Policy Optimization (PPO) [28] baseline is trained on the identical Go2 MuJoCo model using Stable Baselines3 [31]. The observation space includes joint angles, velocities, body orientation, and height (identical to the sensory input available to MH-FLOCKE's SNN). The action space maps to the same 12 actuators. PPO is trained for 50,000 environment steps with default hyperparameters (learning rate $3e-4$, clip range 0.2, 64 minibatch size, 2048 steps per update). Ten seeds are run for statistical comparison.

4.4 Metrics

Maximum distance (meters): furthest point reached from origin, measured as Euclidean distance. Falls: number of fall transitions (upright < 0.3 after being upright > 0.7). Scents found (sf): number of olfactory source targets reached (radius 3.0m). Actor competence: 0.0 (fully CPG) to 1.0 (actor trained, CPG at minimum 40%). All metrics are logged in FLOG, a custom binary format with msgpack-encoded frames. Statistical significance is assessed via mean and standard deviation across 10 seeds.

5. Results

5.1 Ablation Progression (Original Morphology)

Table 1 shows the progression of the full system (C1, C2) across ablation iterations on the original Bommel morphology as components were added. The most significant improvement came from vestibular reflexes and motor babbling (#4 to #7): falls dropped from 27 to 0 on hilly terrain, and flat-terrain distance increased by 763% (from 1.50m to 12.95m). The forward model fix (#9) added 3.9% on flat terrain with measurable prediction errors for the first time. The sensory environment (#10) introduced behavior switching at the cost of 10% forward distance but with 11 scent sources found.

Table 1: Full system progression across ablation iterations (original morphology, single seed)

Run	#4	#7	#9	#10	Δ #4 \rightarrow #10	Key change
C1 flat	1.50m/1F	12.95m/0F	13.45m/0F	12.14m/0F	+710%	Babbling+vest.
C2 hilly	4.47m/27F	12.18m/0F	12.40m/0F	6.45m/1F	+44%	Steer interact.

5.2 Component Isolation (Ablation #10, Original Morphology)

Table 2 shows the full ablation #10 results across all 6 configurations with the sensory environment active on the original Bommel morphology.

Table 2: Ablation #10 results (50k steps, Sensory v0.4.0 + FM fix + Steering 0.05, original morphology)

Config	#7	#9	#10	Falls	SF	vs #9	Note
A1 CPG flat	15.96m	16.44m	16.44m	0	0	±0%	No change
A2 CPG hilly	17.20m	16.92m	16.92m	0	0	±0%	No change
B1 SNN flat	10.07m	3.43m	12.91m	0	19	+277%	Sensory rescues
B2 SNN hilly	9.89m	12.45m	14.71m	0	13	+18%	Drives help
C1 Full flat	12.95m	13.45m	12.14m	0	11	-10%	Steer→CB
C2 Full hilly	12.18m	12.40m	6.45m	1	8	-48%	Regression

5.3 Sensory Environment Validation

The sensory environment underwent four iterations to achieve functional olfactory navigation. The initial implementation placed scent sources at fixed positions with quadratic distance decay, yielding zero scent sources found because the creature drifted laterally during motor babbling. The final version (v0.4.0) uses distance-based respawning (> 5m triggers new source ahead), linear distance decay (smell = $1/(1+dist)$), and olfactory steering via asymmetric hip abduction (gain = 0.05). Validation at 100k steps: 22.02m distance, 19 scents found, smell strength stable at 0.18-0.25. Without the sensory environment, the same configuration reached 20.85m but with zero scents and stagnating distance after 100k steps.

5.4 Forward Model Breakthrough

Prior to ablation #9, the cerebellar forward model showed zero prediction error across all runs, indicating it was not receiving training data. The root cause (Issue #71) was a timing error: the forward model update occurred after an early-return gate that bypasses processing during the CPG-dominated phase. Moving the FM update before this gate produced measurable prediction errors ($e_{FM} = 0.005-0.006$) and doubled correction magnitudes (0.031 to 0.059). C1 improved by 3.9% (12.95m to 13.45m), confirming that cerebellar adaptation provides a measurable contribution on flat terrain.

5.5 200k Long Run Analysis

A 200k-step run of the full system (C1, pre-sensory) reached 31.03m with zero falls. Distance progression showed clear stagnation: 2.8m/10k steps in the first 50k, decreasing to 0.4m/10k in the final 50k. The CPG-only configuration would extrapolate to approximately 65m at 200k steps, making the full system 48% of the CPG baseline at this timescale. This growing gap is not a failure: the full system spends time on non-locomotion behaviors (alert 15%, look_around 6%) driven by its cognitive architecture.

5.6 Cross-Embodiment Transfer: Unitree Go2

Updated with 10-seed validation.

To test whether the architecture generalizes beyond a single morphology, we transferred MH-FLOCKE to the Unitree Go2 quadruped robot model without any architectural changes. The Go2 has the same joint topology (4 legs, 3 joints each, 12 actuators) but significantly different physical properties: heavier mass, different link lengths, different actuator torque limits, and a different center of mass. The only adaptation required was dynamic PD gain scaling (0.4 standing, 1.5 fallen) and an auto-reset at 500 steps because the Go2 cannot self-right after falls.

Table 3 shows the Go2 ablation results across 10 seeds with mean \pm standard deviation. The key finding is that MH-FLOCKE's neural learning core (B1, SNN+Cerebellum) achieves **45.15 \pm 0.67 m**, compared to **12.83 \pm 7.78 m** for PPO on the same morphology — a **3.5x advantage** with dramatically lower variance. The architecture transferred without modification, demonstrating that the biological learning mechanisms generalize across body plans.

Table 3: Go2 ablation results (50k steps, 10 seeds, mean \pm std). PPO trained with Stable Baselines3 [31].

Config	Distance (m)	Falls	Upright Str.	Actor C.	PCI	Note
A1 CPG flat	40.73 \pm 6.14	0.2 \pm 0.4	49627	0.90	0.786	Moderate var.
A2 CPG hilly	26.96 \pm 9.68	1.0 \pm 1.2	47483	0.87	0.795	Terrain -34%
B1 SNN+CB flat	45.15 \pm 0.67	0.0 \pm 0.0	50000	1.00	0.231	Best, 0 falls

Config	Distance (m)	Falls	Upright Str.	Actor C.	PCI	Note
B2 SNN+CB hilly	23.55 ± 3.95	0.0 ± 0.0	50000	0.96	0.258	Terrain -48%
C1 Full flat	22.35 ± 11.50	3.3 ± 2.8	39677	0.69	0.256	High var., falls
C2 Full hilly	20.21 ± 2.54	0.7 ± 0.8	46691	0.77	0.258	Stable
PPO flat	12.83 ± 7.78	0.0 ± 0.0	50000	—	—	3.5x worse
PPO hilly	9.49 ± 5.29	0.0 ± 0.0	50000	—	—	Terrain -26%

Several notable findings emerge from the 10-seed validation: (1) B1 achieves remarkable consistency: $\sigma = 0.67$ m across 10 seeds is the lowest variance of any biological condition, indicating that the SNN+Cerebellum learning core converges to a highly reliable locomotion strategy. All 10 seeds achieve zero falls and full actor competence. (2) A1 outperforms C1 on pure distance (40.73 m vs. 22.35 m), consistent with the original morphology finding that the full cognitive architecture trades locomotion efficiency for behavioral diversity. (3) C1 shows the highest variance ($\sigma = 11.50$ m) and is the only configuration with substantial falls (3.3 ± 2.8), suggesting that the additional cognitive layers introduce instability at the 50k-step training horizon. (4) Terrain effect is asymmetric: B1 loses 48% on hilly terrain ($45.15 \rightarrow 23.55$ m), while C1 loses only 10% ($22.35 \rightarrow 20.21$ m), suggesting the full system's behavioral diversity provides some terrain robustness. (5) PPO shows high variance ($\sigma = 7.78$ m flat) compared to B1 ($\sigma = 0.67$ m), indicating that biological priors not only improve mean performance but also stabilize learning across random seeds.

Notably, the Go2 B1 results show substantially higher distances than the original Bommel morphology (45 m vs. 13 m at 50k steps). This is likely due to the Go2's optimized mass distribution and actuator configuration, which is designed for locomotion, whereas Bommel is a procedurally generated quadruped.

5.7 CPG Fall at Step 49.5k: Evidence for SNN Necessity

Documenting a key finding from the Go2 ablation.

A critical event in the A1 (CPG-only) configuration occurred at step 49,500 in one of the ten seeds: as the SNN's learned velocity contribution dropped (due to lack of sustained reward signal in the CPG-only config), the Competence Gate (Eq. 4) increased CPG weight to 88%. The resulting rigid, high-amplitude CPG gait caused the Go2 to fall. This demonstrates a key architectural principle: the CPG alone is insufficient for sustained locomotion because it lacks the adaptive modulation that the SNN provides. The B and C configurations, which include the SNN, showed zero falls across all seeds (B) or recoverable falls (C), confirming that learned motor adaptation is essential for stability.

5.8 Behavioral Time Breakdown

New section providing quantitative behavioral analysis.

To quantify the behavioral diversity introduced by the cognitive architecture, we analyzed the time spent in each behavior state across all 10 seeds per condition. Behavior state is logged every 1000 steps in FLOG training frames.

Table 4: Behavioral time breakdown (50k steps, 10 seeds, mean ± std %). A configurations show only walk (CPG has no behavior planner).

Behavior	A1 CPG flat	B1 SNN+CB flat	C1 Full flat	C2 Full hilly
walk	100.0 ± 0.0	63.1 ± 24.4	71.2 ± 17.9	68.2 ± 21.0
alert	—	12.4 ± 10.7	12.3 ± 9.1	11.4 ± 8.5
trot	—	8.0 ± 8.3	3.5 ± 4.4	4.7 ± 4.9
motor_babbling	—	6.5 ± 8.0	6.5 ± 8.0	6.5 ± 8.0
look_around	—	5.3 ± 9.7	4.5 ± 4.0	5.1 ± 7.0
sniff	—	4.7 ± 14.1	2.0 ± 6.1	4.1 ± 12.2

The behavioral breakdown reveals that CPG-only configurations (A) spend 100% of time in walk, confirming zero behavioral diversity. B and C configurations spend 29-37% of time in non-locomotion behaviors (alert, look_around, sniff, motor_babbling). This directly explains the distance gap between A1 (40.73 m) and B1/C1: the creature is not failing to walk — it is choosing to do other things. Alert behavior (12-17%) is the dominant non-locomotion behavior, indicating frequent pauses for sensory attention. Motor babbling occupies 6.5%

uniformly, reflecting the fixed 7000-step babbling phase. B1 shows more trot (8.0%) than C1 (3.5%), indicating that the full cognitive architecture's additional neuromodulatory dynamics may suppress the walk-to-trot transition. High variance across seeds (walk std = 17-24%) indicates that individual seeds develop different behavioral profiles, reflecting the stochastic nature of the drive and sensory systems.

6. Discussion

6.1 The Distance Gap as Behavioral Richness

A persistent finding is that the full system (C) walks less far than the CPG baseline (A). At 50k steps on the Go2, A1 reaches 40.73m while C1 reaches 22.35m (55%). The behavioral time breakdown (Section 5.8) provides quantitative evidence: C1 spends only 71% of time in walk, with the remainder allocated to alert (12%), motor_babbling (7%), look_around (5%), trot (4%), and sniff (2%). We argue this is a feature, not a limitation: the purpose of the cognitive architecture is not to maximize distance but to produce a creature that exhibits animal-like behavioral diversity.

6.2 Steering-Cerebellum Interaction

The most surprising result on the original morphology is the negative interaction between olfactory steering and cerebellar learning. In the B configuration (no cerebellum), the sensory environment dramatically improved performance (B1: +277%). In the C configuration, it reduced performance (C1: -10%, C2: -48%). The mechanism: olfactory steering modifies motor output after the cerebellar correction, creating an unpredictable perturbation from the forward model's perspective. The solution would be to integrate steering before the cerebellar pathway. We report this as a genuine negative result and architectural lesson.

6.3 B1 Recovery: Drives Compensate for Missing Cerebellum

B1 (CPG + SNN, flat terrain) collapsed from 10.07m to 3.43m between ablations #7 and #9 on the original morphology due to the forward model timing fix disrupting SNN training dynamics. In ablation #10, with the sensory environment, B1 recovered to 12.91m (+277%), demonstrating that motivational drives and olfactory stimulation can substitute for cerebellar motor corrections.

6.4 Limitations

MH-FLOCKE has several limitations. While cross-embodiment transfer to Go2 is demonstrated, only two morphologies have been tested; generalization across diverse body plans (hexapods, bipeds) remains open. The SNN uses Izhikevich neurons rather than more biologically detailed compartmental models. Sim-to-real transfer has not been attempted. The 200 Hz cognitive loop is computationally expensive, limiting real-time applications. The C2 regression on the original morphology indicates that the steering-cerebellum interaction is not yet resolved. While we implement GWT and measure PCI, we make no claims about machine consciousness. The expanded 10-seed validation reveals that the full cognitive architecture (C) introduces significant instability compared to the neural learning core (B), with falls and high variance on the Go2 platform. This indicates that the cognitive layers require further engineering before they provide a net benefit to locomotion performance. The current results demonstrate that the SNN+Cerebellum core is the primary contributor to locomotion learning, while the cognitive architecture's value lies in behavioral diversity rather than motor performance.

6.5 Comparison with PPO Baseline

Updated with 10-seed validation.

The PPO baseline achieves 12.83 ± 7.78 m on the Go2 (flat), compared to 45.15 ± 0.67 m for MH-FLOCKE's neural learning core (B1). This **3.5x advantage** is striking and is accompanied by a dramatic difference in consistency: PPO's standard deviation is 11.6x larger than B1's (7.78 m vs. 0.67 m). Across 10 seeds, B1 never falls and always reaches full actor competence, while PPO's performance ranges widely.

PPO was trained with default hyperparameters for 50k environment steps — this is a relatively short training horizon for RL, which typically benefits from millions of steps [1]. However, 50k steps is the same budget given to MH-FLOCKE, making the comparison fair in terms of **sample efficiency**.

The advantage of MH-FLOCKE likely stems from the CPG providing a strong locomotion prior: while PPO must discover walking from scratch through reward optimization, MH-FLOCKE begins with an innate gait pattern and uses learning to refine it. This mirrors biological development, where neonatal stepping reflexes provide a foundation for learned locomotion [6]. The trade-off is interpretability and biological plausibility versus asymptotic performance: given sufficient training time (millions of steps), PPO would likely achieve higher pure forward distance, but without the behavioral diversity, sensory navigation, or emotional repertoire that MH-FLOCKE produces.

On hilly terrain, the pattern holds: PPO achieves 9.49 ± 5.29 m compared to B2's 23.55 ± 3.95 m (2.5x advantage). The terrain penalty is proportionally similar for both systems (~26% for PPO, ~48% for B2), suggesting that terrain difficulty affects learning efficiency independently of the architecture.

6.6 Cognitive Architecture Overhead on Go2

Updated: replaces previous B=C section with 10-seed findings.

The 3-seed validation in the previous revision showed identical B and C results, suggesting the cognitive layers had no effect. The expanded 10-seed validation reveals a dramatically different picture: B1 achieves 45.15 ± 0.67 m while C1 achieves only 22.35 ± 11.50 m — the full cognitive architecture **halves locomotion distance** and introduces significant instability.

This divergence has three identified causes:

- 1. Computational overhead.** The 15-step cognitive loop (GWT competition, metacognition, consistency checking, dream mode, synaptogenesis) introduces additional sources of stochastic variation through the Global Workspace competition and neuromodulatory dynamics.
- 2. Behavioral interference.** The full system's drives and behavior planner trigger non-locomotion behaviors (sniff, alert, rest, look_around) that interrupt forward progress. In B, drives are present but the simpler architecture allows more consistent locomotion-dominant behavior.
- 3. Fall cascades.** C1 shows 3.3 ± 2.8 falls — the only configuration with substantial falls. When the cognitive layers modulate motor output through neuromodulatory pathways during a critical phase of gait, the resulting perturbation can trigger a fall. With auto-reset at 500 steps, each fall costs learning progress. In seeds where falls occur early, the compounding effect produces the high variance ($\sigma = 11.50$ m).

This is an honest result that reveals a **maturity gap** in the architecture: the cognitive layers are implemented but not yet properly coupled to motor learning. Specifically, dopaminergic neuromodulation does not yet influence R-STDP learning rates, and metacognitive confidence does not gate motor output during unstable phases. The B configuration represents the current **performance ceiling**, while C represents the current **architectural ambition** — the gap between them is the concrete engineering work remaining.

Interestingly, on hilly terrain C2 (20.21 ± 2.54 m) shows much lower variance than C1 ($\sigma = 11.50$ m), with fewer falls (0.7 vs. 3.3). This suggests that terrain variability may actually stabilize the cognitive architecture by providing richer sensory input that keeps the Global Workspace productively engaged rather than introducing spurious competition.

7. Future Work

Several directions are planned. Muscle synergies [27] would reduce the actor's search space from 12 independent joints to 4-5 synergy dimensions. Recovery learning would address the current failure mode where CPG overrides righting reflexes. The steering-cerebellum interaction (Section 6.2) will be addressed by integrating olfactory steering as an SNN input rather than a post-hoc motor modification. Neuromodulator-to-R-STDP coupling will make drives directly influence learning dynamics, resolving the cognitive overhead issue (Section 6.6). A navigation ablation using scent-finding efficiency as the primary metric will provide a fairer comparison of cognitive layer contributions. Multi-agent interaction would leverage the Theory of Mind module already implemented in CognitiveBrain. Sim-to-real transfer on a physical Unitree Go2 is a longer-term goal. The MH-FLOCKE framework is being prepared for open-source release, building on the Integrity-OS codebase previously published on Zenodo [9].

8. Conclusion

MH-FLOCKE demonstrates that biologically grounded, modular cognitive architectures can produce quadruped locomotion learning with rich behavioral repertoires and strong cross-embodiment generalization. The 15-step closed-loop architecture integrates six timescales of processing, from spinal reflexes to dream-based memory consolidation, in a single coherent computation that runs at every simulation timestep. Systematic ablation quantifies the contribution of each component: vestibular reflexes eliminate falls, motor babbling enables flat-terrain learning, the cerebellar forward model provides measurable motor corrections, and the sensory environment replaces random behavior switching with stimulus-driven goal pursuit. Cross-embodiment transfer to the Unitree Go2 achieves 3.5x the distance of a PPO baseline at identical sample budgets (45.15 ± 0.67 m vs. 12.83 ± 7.78 m across 10 seeds), with the SNN+Cerebellum core showing 11.6x lower variance than PPO. Behavioral time analysis reveals that B and C configurations spend 29-37% of time in non-locomotion behaviors, quantitatively explaining the distance gap as behavioral richness. The honest reporting of cognitive architecture overhead and the B≠C finding provides concrete architectural lessons for future embodied AI systems.

A companion video demonstrating the Go2 learning to walk is available at <https://www.youtube.com/watch?v=Jo7UM6pEFMg>.

The complete source code, including all training scripts, the FLOG analysis dashboard, and the Unitree Go2 configuration, is available at <https://github.com/MarcHesse/mhflocke> under the Apache 2.0 license. Comprehensive documentation with architecture details, mathematical formulations, and API references for all modules is available at <https://mhflocke.com/docs/>. The Go2 model files are redistributed under BSD-3-Clause from the MuJoCo Menagerie project [31].

Disclosure

AI writing tools (Claude, Anthropic) were used to assist with manuscript preparation, code development, and language editing. All research design, architecture implementation, experimentation, data collection, analysis, and conclusions are the sole work of the author.

References

- [1] Rudin, N. et al. (2022). Learning to walk in minutes using massively parallel deep RL. CoRL.
- [2] Miki, T. et al. (2022). Learning robust perceptive locomotion for quadrupedal robots. *Science Robotics*, 7(62).
- [3] Ijspeert, A.J. (2008). Central pattern generators for locomotion control. *Neural Networks*, 21(4), 642-653.
- [4] Angelaki, D.E. & Cullen, K.E. (2008). Vestibular system: the many facets of a multimodal sense. *Ann. Rev. Neurosci.*, 31, 125-150.
- [5] Wolpert, D.M. et al. (1998). Internal models in the cerebellum. *Trends Cogn. Sci.*, 2(9), 338-347.
- [6] Prechtl, H.F.R. (1997). The importance of fetal movements. In *Textbook of Perinatal Medicine*.
- [7] Tinbergen, N. (1951). *The Study of Instinct*. Clarendon Press.
- [8] Sokolov, E.N. (1963). *Perception and the Conditioned Reflex*. Pergamon Press.
- [9] Hesse, M. (2025). Integrity-OS: Neural-Symbolic Hybrid Architecture. Zenodo. DOI: 10.5281/zenodo.18450340.
- [10] Hwangbo, J. et al. (2019). Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26).
- [11] Righetti, L. & Ijspeert, A.J. (2008). Pattern generators with sensory feedback for quadruped locomotion. ICRA.
- [12] Miall, R.C. & Wolpert, D.M. (1996). Forward models for physiological motor control. *Neural Networks*, 9(8), 1265-1279.
- [13] Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- [14] Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness. *Cognition*, 79, 1-37.
- [15] Casali, A.G. et al. (2013). A theoretically based index of consciousness. *Sci. Transl. Med.*, 5(198), 198ra105.
- [16] Tresch, M.C. et al. (1999). The construction of movement by the spinal cord. *Nature Neurosci.*, 2(2), 162-167.
- [17] Murlis, J. et al. (1992). Odor plumes and how insects use them. *Ann. Rev. Entomol.*, 37(1), 505-532.
- [18] Catania, K.C. (2006). Olfaction: underwater sniffing by semi-aquatic mammals. *Nature*, 444(7122), 1024-1025.
- [19] Porter, J. et al. (2007). Mechanisms of scent-tracking in humans. *Nature Neurosci.*, 10(1), 27-29.
- [20] Makin, T.R. et al. (2008). On the other hand: dummy hands and peripersonal space. *Behav. Brain Res.*, 191(1), 1-10.
- [21] Damasio, A.R. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. Putnam.
- [22] Tulving, E. (1972). Episodic and semantic memory. In *Organization of Memory* (pp. 381-403). Academic Press.
- [23] Botvinick, M.M. et al. (2001). Conflict monitoring and cognitive control. *Psychol. Rev.*, 108(3), 624-652.
- [24] Huttenlocher, P.R. (1979). Synaptic density in human frontal cortex. *Brain Res.*, 163(2), 195-205.
- [25] Araque, A. et al. (1999). Tripartite synapses: glia, the unacknowledged partner. *Trends Neurosci.*, 22(5), 208-215.
- [26] Walker, M.P. & Stickgold, R. (2004). Sleep-dependent learning and memory consolidation. *Neuron*, 44(1), 121-133.
- [27] Bizzi, E. & Cheung, V.C.K. (2013). The neural origin of muscle synergies. *Front. Comput. Neurosci.*, 7, 51.

- [28] Schulman, J. et al. (2017). Proximal policy optimization algorithms. arXiv:1707.06347.
- [29] Izhikevich, E.M. (2003). Simple model of spiking neurons. IEEE Trans. Neural Networks, 14(6), 1569-1572.
- [30] Fremaux, N. & Gerstner, W. (2016). Neuromodulated STDP and three-factor learning. Front. Neural Circuits, 9, 85.
- [31] Raffin, A. et al. (2021). Stable-Baselines3: Reliable RL implementations. JMLR, 22(268), 1-8.