



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

MuJoCo Simulates Physics Without Planning Governance

by [Nick Clark](#) | Published March 28, 2026 | [PDF](#)

MuJoCo, now open-sourced by DeepMind, provides the physics simulation substrate that much of modern robotics and reinforcement learning research depends on. Its contact dynamics, articulated body modeling, and fast computation enable agents to explore physical interactions millions of times faster than real time. The simulation fidelity is genuine and the contribution to the field is substantial. But MuJoCo simulates the physical world. It does not govern the planning structures that agents use to reason about that world. An agent exploring MuJoCo trajectories has no containment boundary separating speculation from commitment, no branch classification governing which plans merit promotion, and no executive aggregation resolving conflicts between competing plans. The forecasting engine provides these governance structures.

What MuJoCo provides

MuJoCo's physics engine computes contact forces, joint dynamics, tendon routing, and actuator responses with speed and accuracy optimized for control applications. Reinforcement learning agents interact with MuJoCo environments to learn locomotion, manipulation, and dexterous control through trial and error. The simulator provides the physical ground truth against which agent policies are evaluated. Model-based planning approaches use MuJoCo as a differentiable dynamics model for trajectory optimization.

The simulation handles the physics. The agent's planning process, how it generates candidate actions, evaluates alternatives, and commits to execution, operates above the simulation layer. MuJoCo tells the agent what would happen if it took an action. It does not govern how the agent reasons about which actions to consider, how far to speculate, or when speculation should be promoted to execution.

The gap between simulation and planning governance

A reinforcement learning agent exploring trajectories in MuJoCo may generate thousands of candidate action sequences. Some are physically feasible and productive. Others are exploratory dead ends. Others are dangerous in ways the physics simulator faithfully models but the agent has no structural mechanism to quarantine. The agent learns through reward signals which trajectories to prefer. It does not have a governed planning structure that separates speculative exploration from committed execution.

When such agents are transferred to physical robots, the absence of planning governance becomes consequential. A policy that occasionally explores dangerous trajectories in simulation, learning to avoid them through negative reward, carries that exploration tendency into physical deployment where a single dangerous trajectory has real consequences. The simulation provided the physics. It did not provide the containment boundary that keeps speculative planning separated from executable plans.

The gap also limits the sophistication of planning in simulation. An agent without branch classification treats all candidate plans equivalently. An agent without executive aggregation cannot resolve conflicts between competing planning objectives. An agent without personality-modulated speculation cannot adjust its planning risk profile for different operational contexts.

What the forecasting engine provides

Planning graphs as first-class cognitive structures give agents in MuJoCo governed planning. The containment boundary ensures that speculative trajectories are evaluated within a bounded planning space before any can influence execution. Branch classification labels each candidate plan by type: exploratory, confirmatory, contingency, or committed. Only plans that pass through the executive aggregation process, where competing branches are resolved and validated, are promoted to execution.

For sim-to-real transfer, the forecasting engine ensures that planning governance learned in simulation transfers alongside the policy. The agent that crosses from MuJoCo to physical hardware brings its containment discipline with it. Speculative exploration remains bounded. Dangerous trajectories are quarantined by structure, not just by learned avoidance.

The structural requirement

MuJoCo provides the physical simulation substrate that robotics research needs. The structural gap is planning governance: the cognitive layer that controls how agents reason about the physical possibilities the simulator reveals. The forecasting engine provides containment, classification, and executive aggregation as first-class planning primitives. The agent that plans within a governed forecasting structure does not merely explore physics. It speculates within governed boundaries, classifies its own plans, and commits only what passes through structured validation.

[Forecasting Engine All 21 steps →](#)

Plan before you act. Contain speculation. Promote only what passes.

Primary Technical Disclosure

[Forecasting and Executive Graphs in Autonomous Cognitive Systems](#)

Secondary Technical

[Planning Graphs as First-Class Cognitive Structures](#)[Containment Layer and Delusion Boundary](#)[Branch Classification System](#)[Personality Field as Structural Modifier](#)[Executive Engine Multi-Agent Graph Aggregation](#)[Branch Dormancy and Deferred Promotion](#)[Proactive Speculative Maintenance \(Dream State\)](#)[Planning Graph Archival for Cognitive Forensics](#)[Cross-Agent Planning Graph Visibility](#)[Slope-Constrained Speculative Simulation](#)[Structural Separation From Verified Memory](#)[Forecasting Engine Architecture](#)[Forecasting Execution Cycle](#)[Emotional Modulation of Planning](#)[Executive Graph Conflict Resolution](#)[Planning Graph Delegation and Forking](#)[Temporal Anchoring and Lifecycle Management](#)[Forecasting as Coordination Primitive](#)[Forecasting-Modulated Discovery Traversal](#)[Forecasting as Confidence Input](#)[Integrity-Constrained Forecasting](#)[Forecasting for Training Curriculum](#)[Biological Signal to Forecasting Coupling](#)[Substrate-Agnostic Forecasting Deployment](#)

Applications (General)

[Surgical Robot Planning Through Governed Speculative Branches](#)[Defense Tactical Planning With Contained Speculation](#)[Forecasting Engine for Logistics Planning](#)[Forecasting Engine for Disaster Response Planning](#)[Forecasting Engine for Financial Portfolio Planning](#)[Forecasting Engine for Construction Project Planning](#)[Forecasting Engine for Epidemic Response Planning](#)[Forecasting Engine for Space Mission Planning](#)

Applications (Specific)

[da Vinci Plans Trajectories, Not Consequences](#)[Anduril's Lattice Plans Missions Without Speculative Containment](#)[Boston Dynamics Plans Motion, Not Missions](#)[Shield AI's Hivemind Cannot Contain Its Own Speculation](#)[MuJoCo Simulates Physics Without Planning Governance](#)[NVIDIA Isaac Sim Renders Worlds Without Governing Plans](#)[Unity ML-Agents Trains Without Governing Speculation](#)[Gazebo Simulates Robots Without Governing Their Plans](#)[Drake Optimizes Trajectories Without Governing Planning Structures](#)[robosuite Benchmarks Manipulation Without Governing Plans](#)
[Forecasting Engine overview →](#)

AQ

deterministic

autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



-
- nick@qu3ry.net
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie