A Head-to-Head Battle of Policy Optimization: PPO vs. TRPO

Author: BenchChem Technical Support Team. **Date**: December 2025

In the landscape of reinforcement learning, Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO) stand out as two of the most influential algorithms for continuous control tasks. Both are designed to address the critical challenge of taking the largest possible improvement step on a policy without causing a catastrophic collapse in performance. While PPO emerged as a successor to TRPO, aiming for simpler implementation and better sample efficiency, the true performance differences are often nuanced and subject to specific implementation details. This guide provides an empirical comparison of their performance, supported by experimental data and detailed methodologies.

Core Concepts: A Tale of Two Optimization Strategies

At their core, both PPO and TRPO are policy gradient methods that aim to optimize a policy by taking iterative steps in the parameter space. The key difference lies in how they constrain the policy update to ensure stability.

Trust Region Policy Optimization (TRPO) formulates the problem as a constrained optimization. It seeks to maximize a surrogate objective function while ensuring that the KL-divergence between the old and new policies remains within a certain threshold, known as the trust region.[1][2] This approach guarantees monotonic policy improvement but involves complex second-order optimization, making it computationally expensive and difficult to implement.[1][3][4]

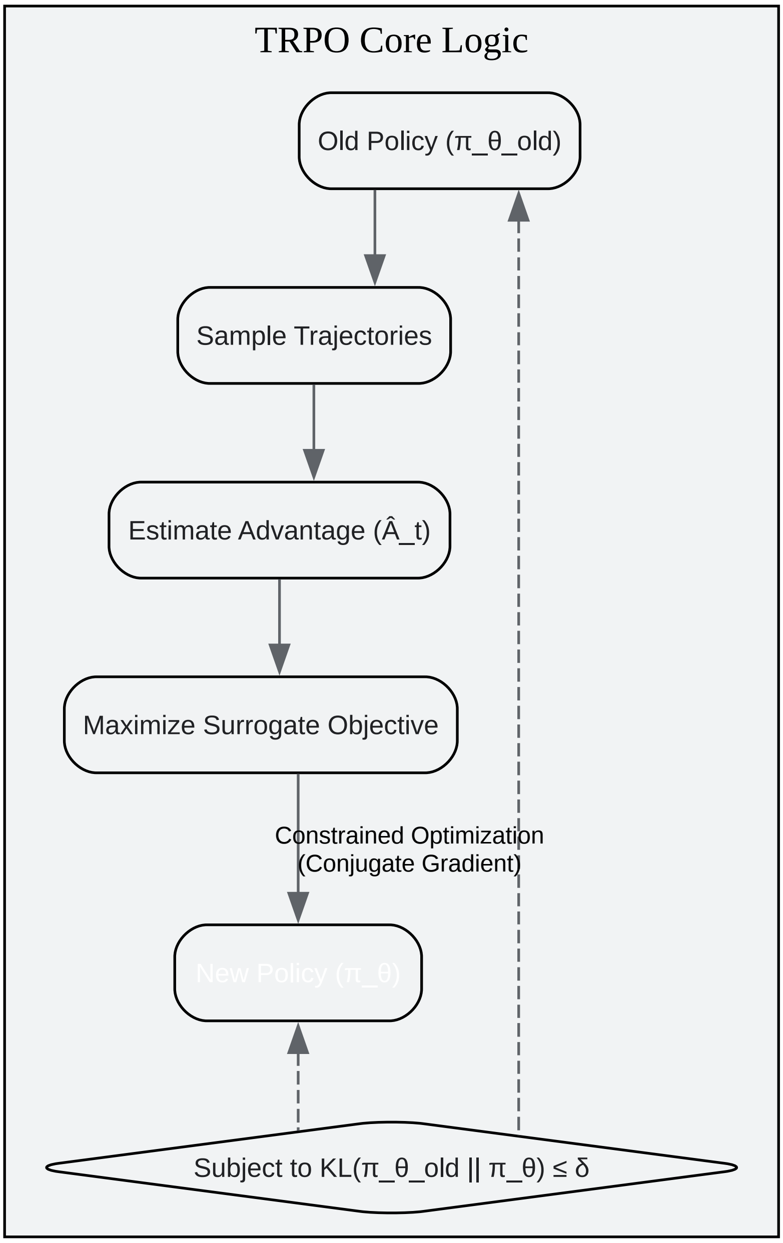

Click to download full resolution via product page

TRPO's constrained optimization workflow.

Proximal Policy Optimization (PPO) simplifies the process by using a clipped surrogate objective function. This clipping mechanism discourages large policy updates by limiting the change in the probability ratio between the new and old policies. This modification allows PPO to be optimized with first-order methods like stochastic gradient descent, making it significantly easier to implement and more computationally efficient.

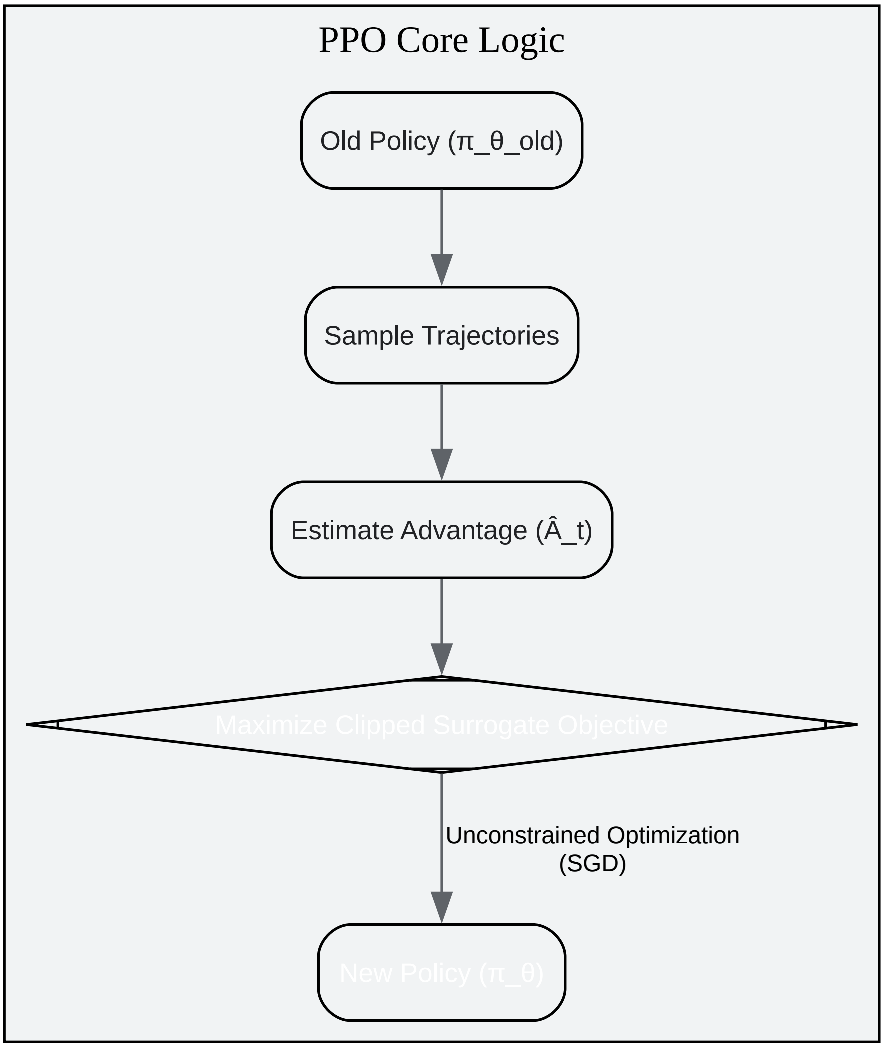

Click to download full resolution via product page

PPO's simplified clipped objective workflow.

Performance Showdown: MuJoCo Benchmarks

The MuJoCo continuous control benchmarks are a standard for evaluating the performance of reinforcement learning algorithms. The following table summarizes the performance of PPO and TRPO on several of these tasks. It is important to note that the performance of these algorithms can be significantly influenced by "code-level optimizations" which are often not part of the core algorithm description. These can include value function clipping, reward scaling, and observation normalization.

MuJoCo Task	PPO	TRPO	PPO (with code-level optimizations)	TRPO (with code-level optimizations)
Hopper-v2	~1816	~2009	~2175	~2245
Walker2d-v2	~2160	~2381	~2769	~3309
Swimmer-v2	~58	~31	~58	~94
Humanoid-v2	~558	~564	~939	~638

Note: The performance is measured as the average total reward. The values are indicative and sourced from various studies, including "Implementation Matters in Deep Policy Gradients: A Case Study on PPO and TRPO". The exact scores can vary based on hyperparameter tuning and random seeds.

The data reveals that while the base versions of TRPO sometimes outperform PPO, the inclusion of code-level optimizations significantly boosts the performance of both algorithms. Interestingly, a version of TRPO with these optimizations can outperform PPO on certain tasks. This highlights a critical finding: the implementation details can be more impactful on the final performance than the choice between the core PPO and TRPO algorithms.

Experimental Protocols

To ensure a fair and reproducible comparison between PPO and TRPO, a standardized experimental setup is crucial. The following outlines a typical protocol for benchmarking these algorithms on MuJoCo environments.

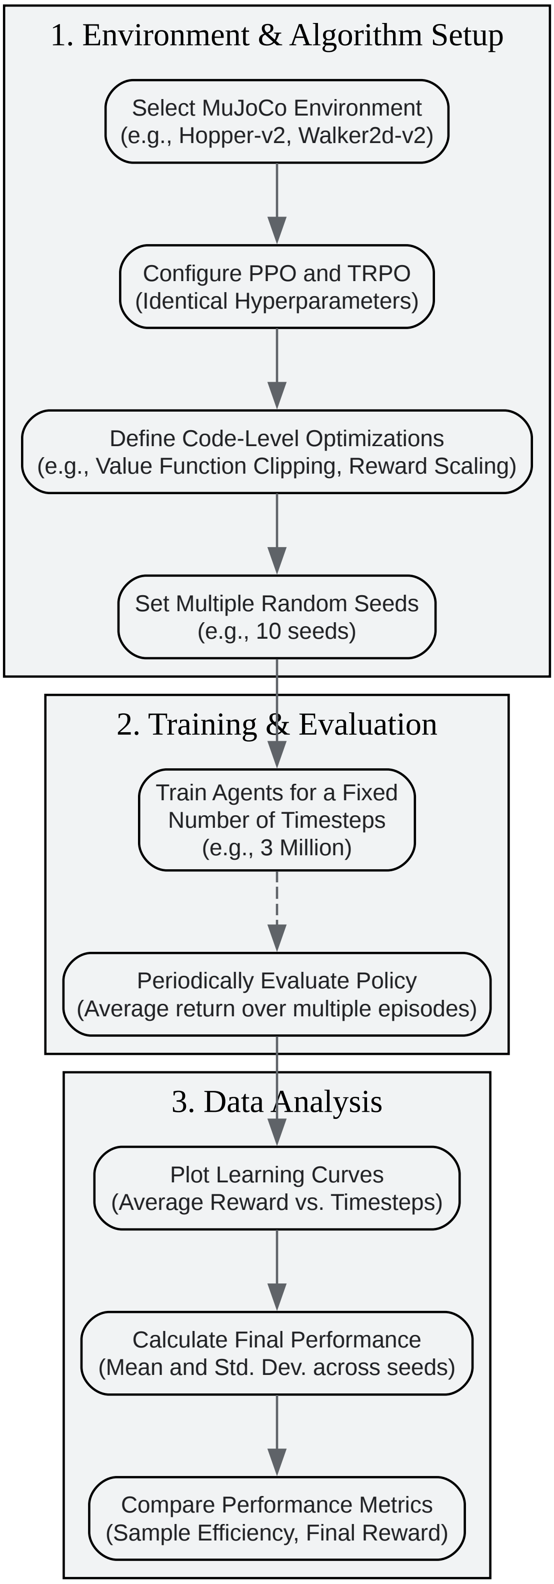
1. Environment & Algorithm Setup
Select MuJoCo Environment (e.g., Hopper-v2, Walker2d-v2)
Configure PPO and TRPO (Identical Hyperparameters)
Define Code-Level Optimizations (e.g., Value Function Clipping, Reward Scaling)
Set Multiple Random Seeds (e.g., 10 seeds)
2. Training & Evaluation
Train Agents for a Fixed Number of Timesteps (e.g., 3 Million)
Periodically Evaluate Policy (Average return over multiple episodes)
3. Data Analysis
Plot Learning Curves (Average Reward vs. Timesteps)
Calculate Final Performance (Mean and Std. Dev. across seeds)
Compare Performance Metrics (Sample Efficiency, Final Reward)

Click to download full resolution via product page

A standardized workflow for comparing PPO and TRPO.

1. Environment: The experiments are typically run on a suite of continuous control environments like those provided by OpenAI Gym's MuJoCo.

2. Hyperparameters: To isolate the effect of the core algorithm, it is essential to use the same set of hyperparameters for both PPO and TRPO where applicable. This includes:

- Neural Network Architecture: A common setup is a multi-layer perceptron (MLP) with two hidden layers of 64 units each, using tanh activation functions.
- Discount Factor (γ): Typically set to 0.99.
- GAE Parameter (λ): For Generalized Advantage Estimation, a value of 0.95 is common.
- Optimizer: Adam is frequently used for PPO, while TRPO uses the conjugate gradient method.
- Learning Rate: A common starting point is 3e-4.

3. Code-Level Optimizations: As their impact is significant, it is crucial to explicitly state which, if any, of these optimizations are used. These can include:

- Observation and reward normalization.
- Value function clipping.
- Gradient clipping.

4. Evaluation: The performance is measured by the average return over a number of episodes. This evaluation is performed periodically throughout the training process to generate learning curves. The final reported performance is typically the average of the returns over the last few training iterations across multiple random seeds to ensure statistical significance.

Conclusion: Simplicity and Performance in Practice

While TRPO offers theoretical guarantees of monotonic policy improvement, its complexity makes it challenging to implement and computationally demanding. PPO, on the other hand, provides a simpler, first-order optimization approach that is more accessible and often achieves comparable or even superior performance, especially when considering wall-clock time.

The empirical evidence suggests that while both algorithms are highly effective, the performance gap between them is often less significant than the impact of implementation-specific "code-level optimizations". For researchers and practitioners, PPO generally offers a better balance of ease of implementation, computational efficiency, and high performance, making it a popular and robust choice for a wide range of reinforcement learning problems. However, for applications where stability is paramount and computational resources are not a primary constraint, TRPO remains a viable and powerful alternative.

References

- 1. TRPO and PPO · Anna's Blog [gaoyuetianc.github.io]
- 2. medium.com [medium.com]
- 3. transferlab.ai [transferlab.ai]
- 4. proceedings.neurips.cc [proceedings.neurips.cc]
- To cite this document: BenchChem. [A Head-to-Head Battle of Policy Optimization: PPO vs. TRPO]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12371345#empirical-comparison-of-ppo-and-trpo-performance]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [Contact our Ph.D. Support Team for a compatibility check]

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com