CASE STUDY · ENERGY | BP

BP: Autonomous Valve Manipulation at 1/90,000th the Compute Cost

ThoughtForge Active Inference benchmarked against Deep Reinforcement Learning across three phases.

10.6 seconds	**$0.0055**	**99.7%**
TO TRAIN A WORKING MODEL	CLOUD COST PER MODEL	TURN ACCURACY UNDER VARIATION

1 SUMMARY

BP wanted a robotic arm that could reliably manipulate a variety of valves with no prior knowledge of each valve: no CAD model, no measured geometry, no knowledge of its physical properties. BP was also evaluating Deep Reinforcement Learning (DRL), so ThoughtForge implemented the solution both ways, once with its patented Active Inference algorithm and once with DRL, and benchmarked the two side-by-side.

ThoughtForge trained a working model in 10.6 seconds for $0.0055 in compute cost, compared to 11.1 days and $300-$700 for DRL models. ThoughtForge achieved 99.7% accuracy for the full reach, grasp, and turn task. The best DRL model achieved only 40% and, in the hardest conditions, turned the valve the wrong way.

2 PROBLEM

Conventional industrial manipulator control solves robotic control against a fixed frame of reference and known, stored positions.t depends on the world matching the model it was given: parts in known places, with known geometry, presented the same way every cycle. Shift a component slightly, and the system must stop and be updated, or the environment corrected. Real-world environments are not predictable: valves vary in shape, stiffness, and slipperiness, edges are worn off, valves corrode, sensors are noisy, and change is inevitable.

DRL is a popular approach to training AI, but faces two fundamental problems:

1. Sample efficiency: it is impossible to know in advance how much data a task will need due to edge cases, and training time is unpredictable.
2. Brittle outside the training set: when a DRL model meets conditions it never saw in training, there is no telling what it will do.

These problems translate into schedule risk, cloud spend, and infrastructure investment before knowing whether a task can be automated.

3 REQUIREMENTS

Use a robotic arm to turn valves to arbitrary degrees under the following conditions:

- Three valve geometries: square, circular and triangular
- Manipulate valves when the shape is unknown
- Use off the shelf sensors rather than precision instrumentation
- Able to manipulate valves under varying rotational resistance and surface slipperiness /friction

4 APPROACH

Both approaches were evaluated on a common testbed: a custom OpenAI Gym environment wrapped a MuJoCo physics simulation of a 7-DOF Fetch manipulator controlled through direct joint torques. The simulation stepped at 10 ms with automated metrics collected for every run. Two task definitions were used: 1) The Full Task gave the arm 12 seconds to reach out, grasp, and turn the valve. 1) The Simplified Task started with the valve already grasped and ran for 5 seconds.

In Phase 1, DRL was implemented using Stable Baselines, covering DDPG, PPO, and SAC and SAC + HER (Hindsight Experience Replay), across two months of observation space reduction, reward tuning, larger topologies, learning rate changes, action noise, and activation function changes, with input from industry practitioners. Phase 2 added Domain Randomization to pursue robustness.

ThoughtForge applied its patented Active Inference framework, which continuously updates a generative control model through environmental interaction, with no backpropagation.

Phase 1 established the performance baseline against DRL.
Phase 2 tested robustness to varied physical properties, motor torque, and contact friction.
Phase 3 pushed ThoughtForge to its limits against moving valves, adjusting sensor noise, and adding faulty sensors to the Full Task.

5 RESULTS

Phase 1: Simplified Task. Best DRL result (SAC) against ThoughtForge:

MEASURE	SAC (BEST DRL)	THOUGHTFORGE
Solved the Full Task	No	**Yes**
Training iterations	15,000,000	**1,500**
Time to train one model	11.1 days	**10.6 seconds**
Cloud cost per model	$300	**$0.0055**
Compute (TFLOPs)	1,434.9	**0.016**
Model size, compressed	3.4 MB	**44 KB**
Turn error, circular valve	4.10 degrees	**1.74 degrees**

ThoughtForge was 10,000x more sample-efficient and 90,000x cheaper to train, on a model 80x smaller on disk. DRL results also varied sharply from run to run, meaning several models would need to be trained and then the best one selected.

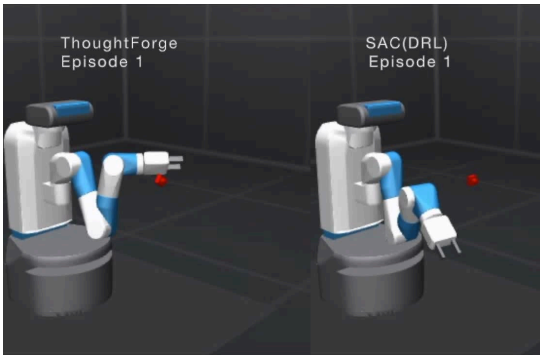

VMuJoCo simulation of ThoughtForge training of the Full Task.

Phase 2: Robustness. Turn accuracy as the test environment diverges from the training environment, PPO (the strongest DRL model, trained with Domain Randomization) against ThoughtForge. Ranges span the three valve shapes:

TEST CONDITION	PPO ACCURACY	THOUGHTFORGE
Default properties	72.0% to 96.7%	**99.7% to 99.8%**
75% variation, physical properties	33.4% to 84.2%	**99.7% to 99.8%**
75% variation, motor torque	49.9% to 76.0%	**99.2% to 99.4%**
90% decrease, contact friction	below 0% to 21.8%	**99.7% to 99.8%**

In the low-friction condition, the PPO model turned the triangular and circular valves away from the target angle. PPO was able to generalize across all valve shapes, but no DRL model learned to turn a valve to the desired angle within the project timeframe, and fewer than 25% of DRL training runs survived to completion. ThoughtForge generalised to all three shapes and to the desired turn angle without the need for domain randomization or adaptation.

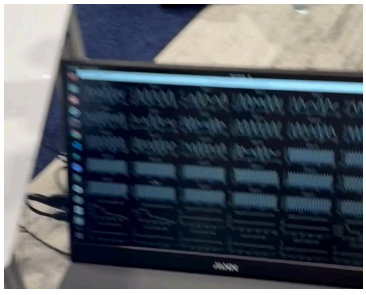

ThoughtForge Valve Manipulation Model and real-time learning demonstrated live at a conference.

Phase 3: ThoughtForge alone on the Full Task. Further DRL comparison was not required. Maintaining above 95% accuracy and robustness, ThoughtForge's Active Inference model tolerated:

- Valve location drifting 0.2 m per second under a random walk, and 5 cm circular motion at 0.1 Hz
- Sensor noise of 0.5 cm on location and 2 degrees on orientation unfiltered, rising to 10 cm and 20 degrees with Kalman and low pass filtering
- Faulty sensors reporting stale data while the valve drifted 2 cm per second in position and 40 degrees per second in orientation
- Valve orientation drifting 20 degrees per second under a random walk
- 50 degrees of noise on the turn angle sensor, with or without filtering

WHAT THIS MEANS FOR THE INDUSTRY AND OUR CUSTOMERS

Cost and schedule risk largely disappear.
A model that trains in seconds for fractions of a cent removes the need for cloud training infrastructure and the multi week retraining cycles that make DRL programs slip whenever reality turns out to have been under modelled.

Robustness holds where it matters.
A valve seized with corrosion, slick with oil, or reached from a vessel that is moving are the normal conditions of a BP site, not edge cases. ThoughtForge held above 99% accuracy across exactly these variations; the DRL models did not.

Cheap sensors become viable.
Tolerance to noisy, drifting, and stale sensor data means commodity instrumentation rather than precision-calibrated hardware, and a 44KB model that runs on resource-constrained devices at the edge.

ABOUT THOUGHTFORGE

ThoughtForge is a robot control system for both simple and complex tasks.
Using our patented AI, ThoughtForge-powered robots easily adapts in real time to new parts, dynamic changes, and varying environments, planning motion and paths on-the-fly with unmatched accuracy and robustness.