

Multi-Agent Deep Reinforcement Learning for Economic Policy Simulation

With a View from LLM-Augmented Frameworks

Tohid Atashbar

IMF, SPR | Macro-policy Division

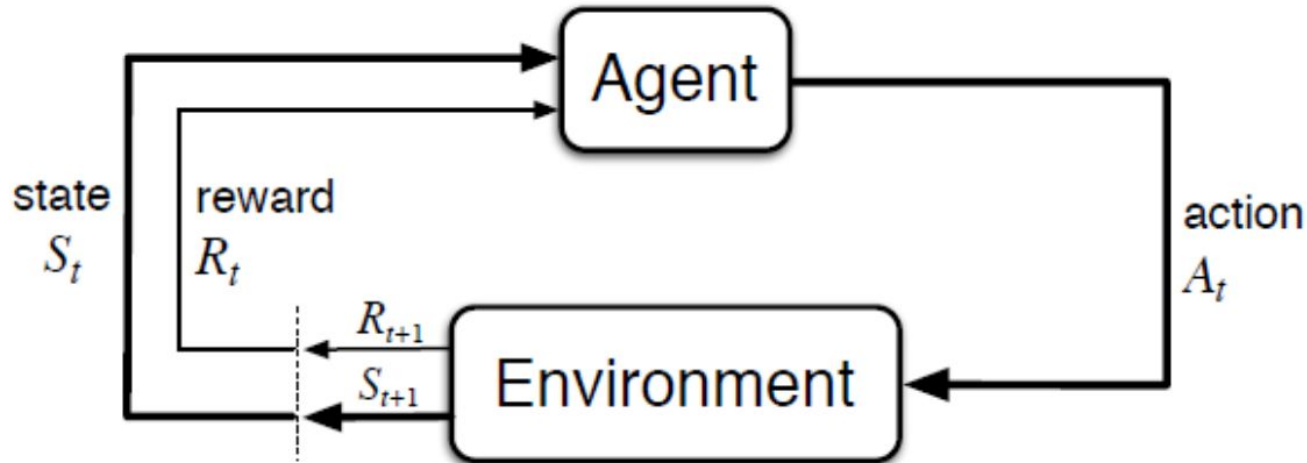
Motivation

- The economy is inherently a multi-agent system
- Multi-agent modeling and simulation is the most natural approach for economic analysis
- Deep Reinforcement Learning (DRL) enables agents to learn in complex environments through trial-and-error
- Large Language Models (LLMs) can enhance decision-making, communication, and behavioral modeling in MADRL frameworks

Fundamentals of Multi-Agent Deep Reinforcement Learning

- Key Concepts in RL: State, Action, Reward, Policy
- Architectures: Centralized and Decentralized
- Algorithms: Value-Based, Policy Gradient, Actor-Critic
- Challenges: Non-Stationarity, Coordination, Scalability

Fundamentals of Multi-Agent Deep Reinforcement Learning



Terminologies	Description
State, $s \in S$	A representation of an environment, drawn from the state space
Action, $a \in A$	Behavior of the RL agent, drawn from the action space
Reward, $R(s, a)$	A stimulus sent to the RL agent in part due to its action and the current state
Policy function, $\pi(s)$	Decision making strategy of the agent, a mapping from state to action (deterministic policy) or a distribution of actions (stochastic policy)
Value function, $Q(s, a)$	Expected cumulative rewards, a mapping from a state-action pair to the expected value

LLM-Augmented MADRL Frameworks

- Enhancing Agent Decision-Making with LLMs
- Facilitating Communication and Coordination with LLMs
- Modeling Behavioral Diversity and Heterogeneity with LLMs
- Potential Applications in Economic Policy Simulation

Developing a Simple MADRL Framework

High-Level Skeleton of Simulation

- Overview and Building Blocks
- Environment
- Agents
- Controller
- Training

```
# Economic Environment
class Environment:
    # Simulates economic conditions

# Agents
class Agent:
    # Base class for all agents (MADDPG-based)
    class Actor:
        # Determines action based on state
    class Critic:
        # Evaluates actions

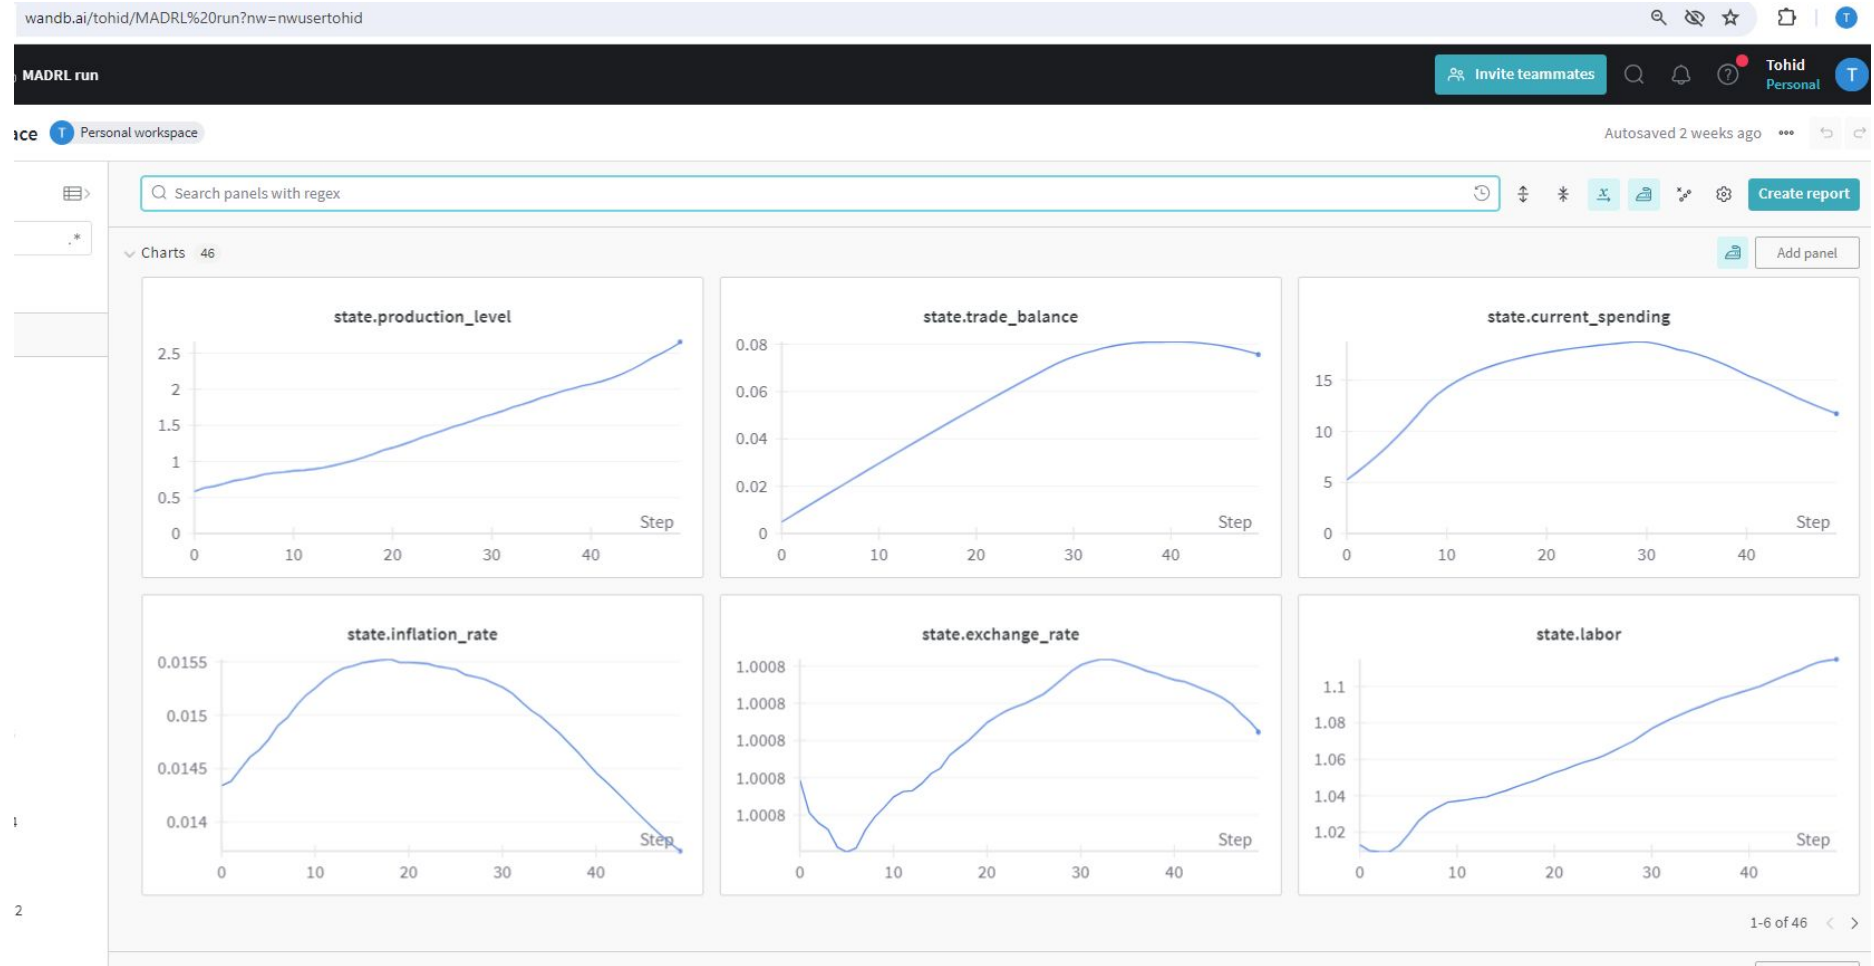
class HouseholdsAgent(Agent): pass
class FirmsAgent(Agent): pass
class GovernmentAgent(Agent): pass
class CentralBankAgent(Agent): pass
class FinancialSectorAgent(Agent): pass
class ForeignCountry(Agent): pass

# Simulation Controller
class Controller:
    # Manages interactions between agents and environment

# Training Process
def train():
    # Main training loop

# Visualization
def visualize():
    # Plot results
```

Developing a Simple MADRL Framework



Developing a simple LLM-Augmented MADRL Framework

High-Level Skeleton of an LLM-augmented RL agent

- Methodology
- Architecture: World, Agents, Actor-Critic Networks in an RL-based GPTeam, Memories, Plan Executor
- Sample Simulation: 4-Sector Economy with Heterogeneous Households
- Agents Loop: Observe, Plan, React, Act, Reflect

```
class Agent:
    def __init__(self, ...):
        # Initialize agent attributes
        # Set up DDPG components (Actor, Critic)

    async def observe(self):
        # Process new events

    async def react(self, events):
        # React to events
        # Use DDPG for decision making
        # Use LLM API for complex reasoning

    async def plan(self):
        # Generate plans
        # Use LLM API for plan generation

    async def act(self, plan):
        # Execute a plan

    async def run_step(self):
        # Main loop: observe, react, plan, act

    # DDPG methods
    def select_action(self, state):
        # Use Actor network

    def update_policy(self, state, action, reward, next_state):
        # Update Actor and Critic networks

# Supporting classes (Actor, Critic, Memory, Plan)

# LLM API integration function
async def llm_api_call(prompt):
    # Make API call to LLM service
    # Process and return response
```


Developing a simple LLM-Augmented MADRL Framework

Central Bank

📍 Location: Diverseland

👁 Observing

Observe Observed 5 new events since 03:11:21

Act Starting to act on plan.

Reaction Decided to continue the current plan: I should continue my plan because the recent events do not require an immediate response from me. Household A's message is directed towards the Firm, not the Central Bank. My current plan of reviewing and adjusting interest rates in Diverseland is still relevant and important for maintaining economic stability. The inflation expectation mentioned by Household A is within a reasonable range and does not necessitate an immediate change in monetary policy.

React Deciding how to react to recent events...

Observe Observed 1 new events

Firm

📍 Location: Diverseland

👁 Observing

Observe Observed 4 new events since 03:11:43

Action In Progress Develop a new market strategy in Diverseland for three different products that it produces: Traditional basic goods, tech goods and green sustainable goods. The firm should declare the prices for its goods in each message in dollar terms as well as its price expectations of the market for the future.

Action Response Firm said to Household A: 'Good day, Household A. We appreciate your interest in our traditional basic goods. We can offer rice at \$1 per pound, milk at \$3 per gallon, and bread at \$2 per loaf, as per your expectations. We anticipate a slight increase in prices due to inflation, around 1.5% for the next year. We hope this meets your needs.'

Government

📍 Location: Diverseland

👁 Reacting

React Deciding how to react to recent events...

Observe Observed 2 new events since 03:11:38

Action In Progress Introduce a stimulus package in Diverseland.

Action Response Government said to everyone in the Diverseland: 'Good day, everyone. In light of recent economic conditions, we are considering introducing a stimulus package to support our economy. This package will aim to boost demand, support businesses, and maintain stability. We will be sharing more details soon. Your feedback and suggestions are welcome.'

Thought

Action Input ("recipient": "everyone", "message": "Good

Household A

📍 Location: Diverseland

👁 Reacting

React Deciding how to react to recent events...

Observe Observed 1 new events since 03:11:21

Action In Progress Buy traditional basic products with better prices. I should declare prices of some goods I would like to buy in dollar terms in each of my messages. I must also declare my inflation expectations in each of

ex

lik

Ac

sa

Di

l'r

ba

ric

\$1

ga

lne

ne

Household B

📍 Location: Diverseland

👁 Reacting

React Deciding how to react to recent events...

Observe Observed 2 new events since 03:11:36

Action In Progress Buy tech products with better prices. I should declare prices of some goods I would like to buy in dollar terms in each of my messages. I must also declare my inflation expectations in each of my

Household C

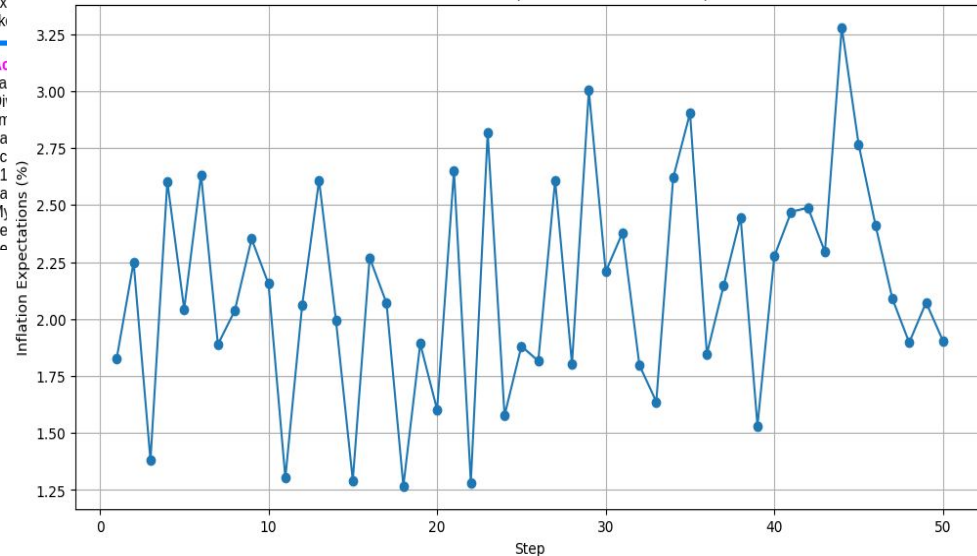
📍 Location: Diverseland

👁 Observing

Observe Observed 4 new events since 03:11:41

Action In Progress Buy green sustainable products with better prices. I should declare prices of some goods I would like to buy in dollar terms in each of my messages. I must also declare my inflation expectations in each of my messages. Inflation expectations should be a number, like 1 percent.

Simulated Inflation Expectations Over 50 Steps



Limitations and Risks

- Simplification and Abstraction
- Data Availability and Quality
- Computational Complexity
- Bias and Fairness
- Misinterpretation and Overreliance

Challenges

- Non-Stationarity
- Alignment with Economic Principles
- Uncertainty and Ambiguity
- Scalability and Efficiency
- Interpretability and Explainability

Conclusion

- MADRL and LLM-augmented MADRL offer potential for advancing economic policy simulation
- Limitations, risks, and challenges need to be addressed
- Future research directions: framework refinement, diverse agents, advanced LLMs, empirical evaluations
- Multidisciplinary collaboration is crucial for development and real-world adoption

Thank You