CSC 36000: Modern Distributed Computing *NextGen with AI Agents*

Written Assignment #2: Project Review, Distributed AI and Distributed Architectures

Total Points: 100

Instructions

The goal of this assignment is to synthesize the work from your Project so far, and to create a clear plan for your Final Report and Presentation. This assignment asks you to reflect on that work and structure your final arguments.

Please answer the following questions to the best of your ability by **writing them manually on a paper**. Assignment submission is to be done by taking all photos of your results, making sure you are uploading all the photos and then upload them on Brightspace.

Your explanations should be clear and concise, demonstrating your understanding of your own project and the key course concepts. **This assignment must be completed with your assigned project group**, but each student must submit their own manually written copy.

Academic Integrity Statement (2 points)

I, **Write your Student Name here replacing this comment**, having Student ID **Write your Student ID here replacing this comment**, fully support this courses CSC36000's AI policy and the City College of New York's policy on Academic Honesty, which states the following in the Syllabus Document under GenAI policy and University policy

Note: You have to write the above academic integrity statement on a Paper manually on top of the manual writing components of this writing assignment. You will take photos of the writeup that you wrote manually and you will upload those photos showcasing your homework result. It is very important to manually write your homework as instructed for all questions.

Problem 1: Project Progress Review (25 points)

a) (5 points) State your project title and the full names of all group members.

b) (10 points) In 2-3 sentences, summarize your thoughts on the **key finding** or **main accomplishment** that your project seeks to complete. What part of your project is working well or complete?

c) (10 points) In 2-3 sentences, summarize the **biggest challenge** or **unexpected problem** you are currently facing). What key part of your project is *not* working as expected yet? How do you plan to fix it?

Problem 2: Project Evaluation & Metrics (25 points)

a) (10 points) In your Final Report, how will you *quantitatively* prove your project is successful? Share the **primary metrics** for distributed computing that you are using to evaluate your system (e.g., latency, F-measure, accuracy, energy saved, task completion time). Why did you select these metrics?

b) (15 points) Every good project needs a **baseline** (e.g., a "naive" algorithm, random choice, or a simpler method). Re-state your baseline. Then, **describe the single most important graph or table** you will create for your Final Report to compare your method against this baseline. Clearly state what the X-axis, Y-axis, and data series would be (for a graph) or what the columns and rows would be (for a table).

Problem 3: DQN Target Calculation (10 points)

An agent is training using a Deep Q-Network (DQN) with a discount factor $\gamma = 0.9$.

- At state s, the agent's **main network** predicts the Q-values:
  - $Q(s, \text{'left'}) = 15$
  - $Q(s, \text{'right'}) = 18$
- The agent decides to take the action **'left'**.
- After taking this action, it receives a **reward** $r = 5$ and lands in a new state s'.
- The **target network** (a separate, fixed copy) is used to evaluate the next state s' and predicts:
  - $Q_{\text{target}}(s', \text{'left'}) = 20$
  - $Q_{\text{target}}(s', \text{'right'}) = 10$

Calculate:

a) (5 points) The target Q-value y (the "ground truth" for this step).

$$y = r + \gamma \max_{a'} Q_{\text{target}}(s', a')$$

b) (5 points) The Mean Squared Error (MSE) loss for this single experience.

$$\text{Loss} = (y - Q(s, \text{'left'}))^2$$

Show your work for both calculations.

Problem 4: Distributed Deep Reinforcement Learning (25 points)

a) (7 points) Why is a tabular Q-learning approach (i.e., using a simple table) not feasible for a complex problem like an Atari game with an 84x84 screen? What is this specific problem called, and how do Deep Q-Networks (DQN) solve it?

b) (6 points) Deep Q-Learning (DQN) introduced two key innovations to stabilize training. Name and briefly describe **one** of them (i.e., Experience Replay or Fixed Q-Targets).

c) (6 points) What are the two main weaknesses of DQN? What class of methods (e.g., REINFORCE) was introduced to solve these limitations?

d) (6 points) Briefly explain the distributed architecture of **Asynchronous Advantage Actor-Critic (A3C)**. What is the role of the "Global Network" and what do the "Worker Agents" do?

Problem 5: Distributed Training Architectures (15 points)

In the lectures, we discussed two primary architectures for distributed training: the Parameter Server and Ring All-Reduce.

a) (8 points) Compare these two architectures in terms of their **Scalability** and **Reliability (Fault Tolerance)**. Which one is bandwidth-optimal and why?

b) (7 points) The Parameter Server architecture is said to have a "Single Point of Failure". What does this mean, and what component of the architecture does it refer to?

Bonus Problem: Distributed Training Architectures (Contd.) (10 points)

Read your answers in Problem 5 and Followup on this Bonus Question (optional) if you would want to answer a bonus question.

a) (3 points) What is a "straggler" in distributed training?

b) (7 points) Which of the two architectures in Problem 5 is generally more vulnerable to stragglers, and why?