Course Project

Code of Honor. All external resources used in the project, including research papers, open-source repositories, datasets, and any content or code generated using AI tools, e.g., ChatGPT, GitHub Copilot, Claude, Gemini, must be *clearly cited* in the final submission. The final report must also include *a clear breakdown of individual group member contributions.* Any lack of transparency in the use of external resources or in reporting group contributions will be considered academic dishonesty and will significantly impact the final evaluation.

Topic	RL for Wireless Resource Allocation
Category	Resource Allocation

Objective Design and implement an RL agent for wireless resource allocation in a small multi-user system. The agent should allocate power or bandwidth among users to maximize throughput while ensuring fairness. The project aims to compare RL-based allocation policies with traditional heuristics such as round-robin or max-SINR scheduling.

Motivation Wireless resource allocation (power control, channel assignment, bandwidth sharing) is a key problem in communication networks. Classical heuristics such as round-robin and max-SINR are simple but cannot adapt to rapidly changing wireless conditions. RL provides a way to design adaptive resource allocation policies that can balance efficiency and fairness in dynamic environments [2, 4, 3, 1]. This project introduces students to RL applications in communication systems, highlighting the connection between algorithm design and real-world ECE challenges.

Requirements The final submission should address the following requirements while the details can be freely decided by the group members.

1. Implementation: in this respect, you should
   - design a simple wireless system with a few users (e.g., 2–5) and time-varying channel conditions,
   - implement an RL agent (e.g., DQN, PPO) that allocates power or bandwidth, and
   - implement baseline policies (e.g., round-robin, max-SINR) for comparison.
2. Environment modification: to ensure originality, you **must** extend any pre-implemented environment of your choice with **at least one** of the following modifications:
   - add constraints such as total power budget or fairness requirement,
   - add noise to channel state information, or
   - simulate heterogeneous users with different Quality-of-Service (QoS) demands.
3. Evaluation: the final project should report key evaluation metrics in the modified environment. In this respect, the results should
   - compare RL policies with baseline heuristics,

- analyze system throughput, fairness, delay, and
- report robustness under changing channel and traffic conditions.

4. The results should be elaborated through
   - ablation experiments, e.g., different reward designs (throughput vs fairness), and
   - providing discussions on trade-offs between complexity and performance.

MILESTONES The following milestones are to be accomplished through semester.

1. Literature Review and Setup
   - Review RL applications in wireless resource allocation, and select a study that is feasible to implement.
   - Design a simple wireless system simulation.
   - Finalize environment modifications.
2. Implementation
   - Implement baseline policies (round-robin, max-SINR).
   - Implement RL agent (DQN or PPO) and validate on a simple scenario.
   - Train agent on your modified environments.
3. Evaluation and Analysis
   - Collect and plot system performance metrics (throughput, fairness).
   - Compare RL agent vs baselines under different conditions.
   - Perform ablation experiments (reward variations, environment modifications).
4. Final Report and Presentation

SUBMISSION GUIDELINES The main body of work is submitted through Git. In addition, each group submits a final paper and gives a presentation. In this respect, please follow these steps.

- Each group must maintain a Git repository, e.g., GitHub or GitLab, for the project. By the time of final submission, the repository should have
  - Well-documented codebase
  - Clear `README.md` with setup and usage instructions
  - A `requirements.txt` file listing all required packages or an `environment.yaml` file with a reproducible environment setup
  - Demo script or notebook showing sample input-output
  - *If applicable,* a `/doc` folder with extended documentation
- A final report (maximum *5 pages*) must be submitted in a PDF format. The report should be written in the provided formal style, including an abstract, introduction, method, experiments, results, and conclusion.
  Important: Submissions that do not use template are considered *incomplete.*
- A 5-minute presentation (maximum *5 slides including the title slide*) is given on the internal seminar on Week 14, i.e., *Dec 1 to Dec 5,* by the group. For presentation, any template can be used.

Final Notes While planning for the milestones please consider the following points.

1. Creativity in modifying the environment and designing reward functions is encouraged.
2. Training should remain feasible by keeping the system small (few users, low-dimensional states).
3. Students are expected to balance research-oriented exploration with achievable implementation.
4. Teams are expected to manage their computing needs and are advised to perform early tests to estimate runtime and training feasibility. As graduate students, team members can use facilities provided by the university, e.g., ECE Facility. Teams are expected to inform themselves about the limitations of the available computing resources and design accordingly.

References

[1] Abdulmalik Alwarafy, Mohamed Abdallah, Bekir Sait Ciftler, Ala Al-Fuqaha, and Mounir Hamdi. Deep reinforcement learning for radio resource allocation and management in next generation heterogeneous wireless networks: A survey. *arXiv preprint arXiv:2106.00574*, 2021.

[2] Le Liang, Hao Ye, and Geoffrey Ye Li. Spectrum sharing in vehicular networks based on multi-agent reinforcement learning. *IEEE Journal on Selected Areas in Communications*, 37(10):2282–2292, 2019.

[3] Hao Ye, Geoffrey Ye Li, and Biing-Hwang Fred Juang. Deep reinforcement learning based resource allocation for v2v communications. *IEEE Transactions on Vehicular Technology*, 68(4):3163–3173, 2019.

[4] Nan Zhao, Ying-Chang Liang, Dusit Niyato, Yiyang Pei, Minghu Wu, and Yunhao Jiang. Deep reinforcement learning for user association and resource allocation in heterogeneous cellular networks. *IEEE Transactions on Wireless Communications*, 18(11):5141–5152, 2019.