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 Autonomous Driving
Category	Robotics and Automation

Objective Design and implement an RL agent for basic autonomous driving tasks such as lane-keeping, overtaking, or collision avoidance. The project aims to compare RL-based driving policies with simple rule-based baselines in lightweight simulators such as `Highway-env` [2].

Motivation Autonomous driving is one of the most exciting application areas for reinforcement learning [1]. Real-world driving requires decision-making under uncertainty, balancing safety, efficiency, and comfort. While training full-scale self-driving systems is computationally prohibitive, simplified environments like `Highway-env` [2] allow students to explore key challenges such as lane following, overtaking, and collision avoidance. This project provides hands-on experience with RL in automation and transportation, while keeping the scale feasible for course resources.

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
   - select a simple autonomous driving environment such as `Highway-env`,
   - implement an RL agent (e.g., DQN, PPO, SAC) to control vehicle behavior, and
   - implement a rule-based baseline (e.g., keep-lane, fixed-speed controller) for comparison.
2. Environment modification: you can use a pre-implemented driving environment. To give the implementation some level of novelty, you **must** modify the environment with **at least one** of the following modifications:
   - adding Gaussian noise to sensor observations (e.g., speed, distance),
   - introducing delays or perturbations in vehicle actions,
   - modifying traffic density or vehicle speed distributions,
   - penalizing unsafe behaviors such as aggressive lane changes.
3. Evaluation: the final project should report key evaluation of the implemented algorithms in the modified environment. In this respect, the results should

- compare safety metrics (collision rate) and efficiency metrics (average speed, time to goal),
- evaluate robustness under environment modifications, and
- analyze convergence speed and stability of training.

4. The results should be elaborated through
   - ablation experiments, e.g., removing penalties for unsafe behavior or modifying reward functions, and
   - providing discussions on trade-offs between safety, efficiency, and performance.

Milestones The following milestones are to be accomplished through semester.

1. Literature Review and Setup
   - Review RL for autonomous driving and familiarize with `Highway-env`.
   - Finalize driving tasks and environment modifications.
2. Implementation
   - Implement rule-based baseline.
   - Implement RL agent (DQN, PPO, or SAC) and validate on unmodified environment.
   - Train RL agent on modified environment.
3. Evaluation and Analysis
   - Collect and plot driving performance metrics (collisions, average speed).
   - Compare RL agent vs baseline under different conditions.
   - Perform ablation experiments.
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 using small networks, limited episodes, and short rollouts.
3. Students are encouraged to carefully balance safety and efficiency in their reward design.
4. Environment modifications should be realistic but lightweight to ensure experiments remain tractable.
5. Students are expected to balance training feasibility with achieving meaningful results.
6. 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] Alex Kendall, Jeffrey Hawke, David Janz, Przemyslaw Mazur, Daniele Reda, John-Mark Allen, Vinh-Dieu Lam, Alex Bewley, and Amar Shah. Learning to drive in a day. In *2019 international conference on robotics and automation (ICRA)*, pages 8248–8254. IEEE, 2019.

[2] Edouard Leurent. An environment for autonomous driving decision-making. `https://github.com/eleurent/highway-env`, 2018.