REINFORCEMENT LEARNING: THEORY AND PRACTICE

Reproducibility

Profs. Amy Zhang and Peter Stone

TEXAS
The University of Texas at Austin

Logistics

Final project guidelines:

- Please be as detailed as possible regarding: the problem you are trying to solve, the RL algorithms/hyperparameters used, codebases used, and other relevant aspects of the project. Ideally you would include as much details so that somebody can reproduce your setup and experiments.
- We suggest having a figure/diagram for the specific problem/environment you are trying to solve. This can also be useful for the presentation video.
- Be clear on why the problem is a sequential decision-making one and formulate the problem as a Markov decision process with details about what the states, actions, and rewards are.
- There should be clear evaluation so we can interpret the results. For most projects, you will want to have some sort of baseline(s) you compare against, so make sure what those are is clear.

Logistics

Final project guidelines:

- For the literature review section, you should relate the works to what you actually end up doing.
- The report should have clearly separated sections and be formatted like a conference paper. Examples of sections you may want to include: abstract, introduction/motivation, related work, methodology, experiments/results, conclusion, etc.
- **All writing should be your own -- all quotes must be clearly attributed.**
- Include at least 10 citations with full bibliographic references to acknowledge where your ideas came from.
- Be very clear about what code you've used from other sources, if any. Clear citations are essential. Failure to credit ideas and code from external sources is cheating.
- Make sure you evaluate both the good and bad points of your approach.

Deadlines

- Due date: Monday April 29th 11:59 pm and can be submitted late until Saturday May 4th 11:59 pm (1% of project grade lost per day).

Discussion Exercise

Laith Altarabishi:
Given the high variance of results as a consequence of hyper parameter tuning, to what degree can we really reason about the innovation of new algorithms/methods in the field? What if there are potentially great breakthroughs that are only restricted by a poor choice of hyper-parameters? Or vice versa, great experimental results that seem to have hit the perfect sweet spot of hyper-parameters making a method/algorithm seem like a great improvement compared to methods that didn't tune hyper-parameters with the same level of accuracy?

Discussion Exercise

Laith Altarabishi:
Given the high variance of results as a consequence of hyper parameter tuning, to what degree can we really reason about the innovation of new algorithms/methods in the field? What if there are potentially great breakthroughs that are only restricted by a poor choice of hyper-parameters? Or vice versa, great experimental results that seem to have hit the perfect sweet spot of hyper-parameters making a method/algorithm seem like a great improvement compared to methods that didn't tune hyper-parameters with the same level of accuracy?

How do we evaluate results?
Brittleness of hyperparameters = bad, sweeps showing robustness to changes in hyperparameters = good
Description of how hyper parameters were tuned, using the same approach across all methods = good

Reading Responses

William Avery

Coming from the computer vision research domain, I think many of the same problems exist. Researchers use different hardware, experimental hyperparameters, augmentation setups, so on and so forth. This creates a problem where it is difficult to compare new approaches apples to apples, and the time and resources required to replicate and verify others' work is often untenable, as it eats into personal research as well as funding in the form of TACC credits. What do you suggest as a solution to this type of problem, as I highly doubt hardware setups and experimental practices will unify anytime soon?

Reading Responses

William Avery

Coming from the computer vision research domain, I think many of the same problems exist. Researchers use different hardware, experimental hyperparameters, augmentation setups, so on and so forth. This creates a problem where it is difficult to compare new approaches apples to apples, and the time and resources required to replicate and verify others' work is often untenable, as it eats into personal research as well as funding in the form of TACC credits. What do you suggest as a solution to this type of problem, as I highly doubt hardware setups and experimental practices will unify anytime soon?

Opensourcing code, virtual environments, compute

Experimental setups *should* unify!

All setup information should be reported for an apples-to-apples comparison.

Reading Responses

May Liu

This paper was published in 2019 it seems like? I wonder what are some progress that have been made towards making RL research more reproducible over the years that follow.

MLRC2023 Home Submission Form Past Iterations

ML Reproducibility Challenge 2023

Welcome to the ML Reproducibility Challenge 2023 (**MLRC 2023**). This is the seventh edition of the event (v1, v2, v3, v4, v5, v6). The primary goal of this event is to encourage the publishing and sharing of scientific results that are reliable and reproducible. In support of this, the objective of this challenge is to investigate reproducibility of papers accepted for publication at top conferences by inviting members of the community at large to select a paper, and verify the empirical results and claims in the paper by reproducing the computational experiments, either via a new implementation or using code/data or other information provided by the authors.

Call For Papers

Kaggle Awards for the ML Reproducibility Challenge 2022

Hi Kagglers!

We're very happy to announce that we will be sponsoring $500k in awards for the ML Reproducibility Challenge 2022. The Reproducibility Challenge is an annual challenge where participants work to reproduce one of the papers from Top ML Conferences & Journals in 2022 and submit a report for review & acceptance. You can find all official participation requirements and documentation like key dates, templates, and review criteria for the challenge on their website at: https://paperswithcode.com/rc2022.

Deep Reinforcement Learning at the Edge of the Statistical Precipice

Rishabh Agarwal* Google Research, Brain Team MILA, Université de Montréal

Max Schwarzer MILA, Université de Montréal

Pablo Samuel Castro Google Research, Brain Team

Aaron Courville MILA, Université de Montréal

Marc G. Bellemare Google Research, Brain Team

Product Solutions Open Source Pricing

google-research / rliable Public

Code Issues 3 Pull requests 1 Actions

master 1 B

agarwl Merge pull

NeurIPS 2021 outstanding paper award