

# LLMs for Enhanced Code Review & Optimizing Compute Cluster Efficiency with Reinforcement Learning

**Alexey Rybalchenko**

Software Development for Experiments (SDE) group, GSI Helmholtz Centre for Heavy Ion Research

**GSI/FAIR AI Workshop**

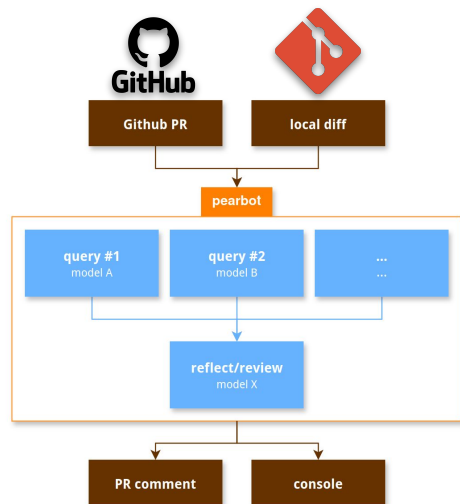
GSI, October 29, 2024





# LLMs for Enhanced Code Review

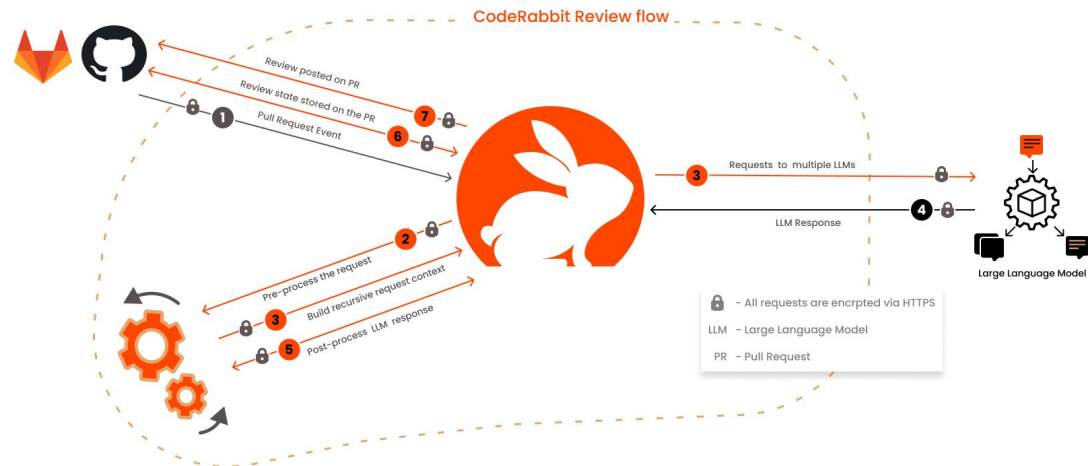
## Pearbot:



<https://github.com/GSI-HPC/pearbot>

- GitHub App for reviewing Pull Requests.
- Local execution mode for diffs or annotated commits.
- Agent ensemble approach for improved results.
- Customizable model(s) via the ollama setup.
- Execution on low-end hardware and/or without GPU.
- Customizable prompt(s).

## CodeRabbit:



<https://coderabbit.ai/>


- AI-based Pull Request reviewer via a Github/Gitlab App.
- Combination of various OpenAI models.
- Previously open source, now a closed project.
- Free to use for open source projects.

# CodeRabbit

## Interesting Results

- + Unbiased, 24/7, scalable, multi-lingual, customizable, with a broad knowledge base, can be kept up-to date with new data.
- + Good at **identifying logical errors** early, that other automated tools or even human review may overlook.
- + Can potentially reduce the time and resources needed for code reviews.
- May produce unnecessary output, when no actionable changes are necessary or such are not deduced by the model.
- Weaker models easily dive into hallucinations.
- Limited understanding of complex projects.
- Potential for false positives, flagging issues that aren't actually problematic.


```
examples/region/fairmq-start-ex-region-advanced.sh.in Outdated Hide resolved
36 + xterm -geometry 90x40+550+40 -hold -e @EX_BIN_DIR@/$PROCESSOR1 &
37 +
38 + PROCESSOR2="fairmq-ex-region-processor"
39 + PROCESSOR2+=" --id processor1"
```

 **coderrabbitai** (bot) on Nov 15, 2023 ...

There seems to be a copy-paste error. The `PROCESSOR2` should have a unique identifier, but it is currently set to `processor1`, which is the same as `PROCESSOR1`. This should be corrected to ensure that each processor has a unique ID.

```
- PROCESSOR2+=" --id processor1"
+ PROCESSOR2+=" --id processor2"
```

```
examples/region/fairmq-start-ex-region-advanced-external.sh.in Outdated Hide resolved
10 + fi
11 +
12 + if [[ $2 =~ ^[0-9]+$ ]]; then
13 +     msgSize=$1
```

 **coderrabbitai** (bot) on Nov 24, 2023 ...

The assignment of `msgSize` is incorrect; it should be assigned the value of `$2` instead of `$1`.

```
- msgSize=$1
+ msgSize=$2
```

# Pearbot

## Quality Improvements


### Self-Consistency

LLMs are probabilistic, so even with Chain-of-Thought, a single generation might produce incorrect results. Self-Consistency introduces enhanced accuracy by selecting the most frequent answer from multiple generations (at the cost of higher compute):

```
John found that the average of 15 numbers is 40.  
If 10 is added to each number then the mean of the numbers is?  
Report the answer surrounded by three backticks, for example: ```123```
```

Running the above several times and taking the most commonly returned value for the answer would make use of the self-consistency approach.

<https://www.llama.com/docs/how-to-guides/prompting/>



- Multi-Agent Initial Reviews:**
  - Multiple AI models generate initial code reviews.
- Reflection<sup>[1][2]</sup> by a “decider” Agent:**
  - A separate, potentially more advanced model analyzes the initial reviews.
  - This agent synthesizes and refines the feedback from multiple sources, rejects potentially less impactful comments.
  - It prioritizes the most important issues and suggestions.
  - The reflection step helps in producing a more comprehensive and coherent final review.
- Prompt improvements:**
  - Specific & useful code review examples.
  - Examples include Chain-of-Thought<sup>[3]</sup> type of reviews, that include some reasoning why the suggestions would be good.

[1] Madaan, Aman, et al. "Self-refine: Iterative refinement with self-feedback." *Advances in Neural Information Processing Systems* 36 (2024). <https://doi.org/10.48550/arXiv.2303.17651>  
[2] Shinn, Noah, et al. "Reflexion: Language agents with verbal reinforcement learning." *Advances in Neural Information Processing Systems* 36 (2024). <https://doi.org/10.48550/arXiv.2303.11366>  
[3] Wei, Jason, et al. "Chain-of-thought prompting elicits reasoning in large language models." *Advances in neural information processing systems* 35 (2022). <https://doi.org/10.48550/arXiv.2201.11903>

Li, Junyou, et al. "More agents is all you need." *arXiv preprint arXiv:2402.05120* (2024). <https://doi.org/10.48550/arXiv.2402.05120>

# Pearbot

Backend

**ollama**<sup>[3]</sup> (via python lib and HTTP request):  
open-source large language model server,  
written in Go, backed by **llama.cpp**<sup>[4]</sup> (C++):

- Efficient serving of large language models
- CPU/GPU/CPU+GPU hybrid inference to partially accelerate models larger than the total VRAM capacity
- Supports many model architectures:  
llama, gemma2, qwen2, ...
- Support for multitude of model quantization techniques and precisions for faster inference and reduced memory use
- Usage Metrics



**LLaMA**<sup>C++</sup>

```
-----  
Model: llama3.1  
  Family: llama, Format: gguf  
  Parameter Size: 8.0B, Quantization: Q4_0  
  Context Length: 131072  
Prompt tokens: 1825  
Tokens generated: 355  
Total tokens: 2180  
Speed: 97.79 tokens/second  
Generation time: 3.63 seconds  
Total duration: 4.25 seconds  
-----
```

# Cluster Job Scheduling with Reinforcement Learning (RL)

Idea

Schedule jobs when electricity prices are lower and/or electricity is “greener”.

The RL approach could be beneficial if:

- There are complex job dependencies/patterns.
- Electricity prices are volatile.
- Need to optimize for many variables simultaneously.
- Lots of historical data.



[https://energy-charts.info/charts/price\\_spot\\_market/chart.htm?!=en&c=DE](https://energy-charts.info/charts/price_spot_market/chart.htm?!=en&c=DE)

# Cluster Job Scheduling with Reinforcement Learning

## Concepts

Agent:

- Scheduling advisor

Environment:

- Cluster

Observations:


- Nodes & their states
- Job queue
- Predicted prices

Actions:

- Turn nodes on / off.

Reward:

- Processing jobs during favorable times.

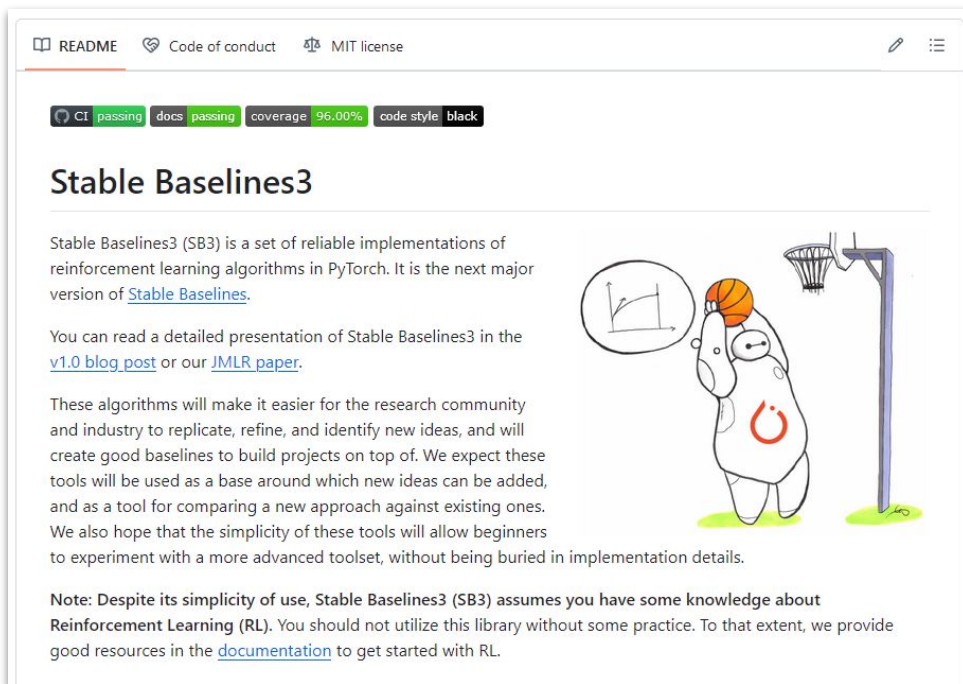


[https://gymnasium.farama.org/introduction/basic\\_usage/](https://gymnasium.farama.org/introduction/basic_usage/)



# Cluster Job Scheduling with Reinforcement Learning

## Tools




Stable Baselines3 (SB3) is a set of reliable implementations of reinforcement learning algorithms in PyTorch. It is the next major version of [Stable Baselines](#).

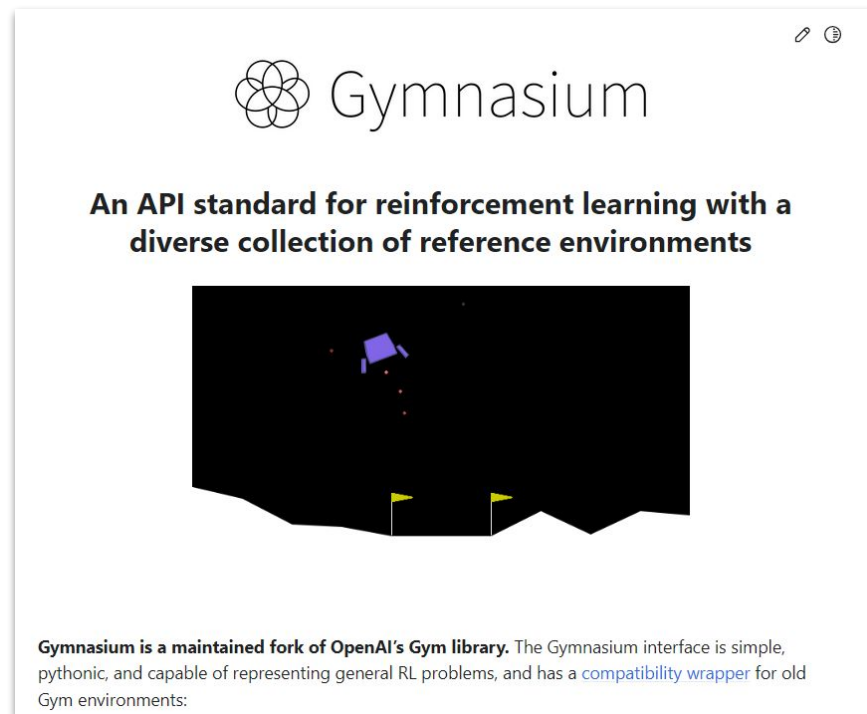
You can read a detailed presentation of Stable Baselines3 in the [v1.0 blog post](#) or our [JMLR paper](#).

These algorithms will make it easier for the research community and industry to replicate, refine, and identify new ideas, and will create good baselines to build projects on top of. We expect these tools will be used as a base around which new ideas can be added, and as a tool for comparing a new approach against existing ones. We also hope that the simplicity of these tools will allow beginners to experiment with a more advanced toolset, without being buried in implementation details.

**Note:** Despite its simplicity of use, Stable Baselines3 (SB3) assumes you have some knowledge about **Reinforcement Learning (RL)**. You should not utilize this library without some practice. To that extent, we provide good resources in the [documentation](#) to get started with RL.

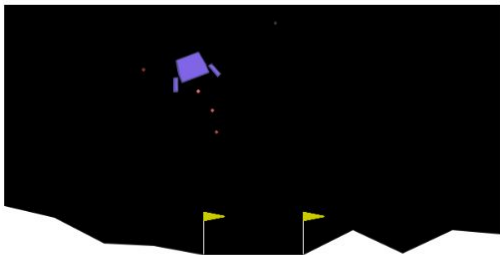


<https://github.com/DLR-RM/stable-baselines3>



## Gymnasium

An API standard for reinforcement learning with a diverse collection of reference environments



**Gymnasium is a maintained fork of OpenAI's Gym library.** The Gymnasium interface is simple, pythonic, and capable of representing general RL problems, and has a [compatibility wrapper](#) for old Gym environments:

<https://gymnasium.farama.org/>

# Cluster Job Scheduling with Reinforcement Learning

Some early results **WORK IN PROGRESS**

