Deep reinforcement learning for magnetic control on WEST

Samy. Kerboua-Benlarbi[1,2] , R. Nouailletas[1], B. Faugeras[2], P. Moreau[1]
1 Institut de Recherche sur la Fusion par Confinement Magnétique, CEA Cadarache, 13115 St-Paul-lez-Durance, France
2 Laboratoire Jean-Alexandre Dieudonné, Université Côte D'Azur, CNRS, Parc Valrose 06108 NICE CEDEX 2, France

A learning-based approach

The limits of classical controllers

One of the many problems to be addressed in tokamak control is the magnetic control of the plasma, i.e. **the control of the voltages applied to the poloidal coils** located around the machine.

Magnetic controllers on WEST

- **Linear**
- **Specific tuning of control parameters**
- **Separate control loops**
- **Not very robust with respect to the coupled behaviour of plasma dynamics**
- **Need for reconstruction codes**

A call in reinforcements

In reinforcement learning (RL), **an agent maximizes its cumulative reward over time** $G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$, with $\gamma \in [0,1]$, **exploring possible outcomes to optimize its behavioral policy** $\pi : S \rightarrow A$.

Deep RL

- **Non-linear, black-box neural networks (NN)**
- **Single reward function**
- **Single control loop**
- **Handles the coupled behaviour of plasma dynamics**
- **Direct use of raw data**

Objective: Tracking of plasma's shape, position and current using deep reinforcement learning in the WEST use-case.

A general distributed framework for WEST

A **general framework based on [1,2,6]** has been developed**:**

- Logical components connected through **socket protocols (TCP, UDS)** for **multi-language interaction (C++, Python)**,
- **Multi-GPU, multi-threaded** to run multiple environments in parallel**.**

Training loop → Fast, asynchronous and reliable

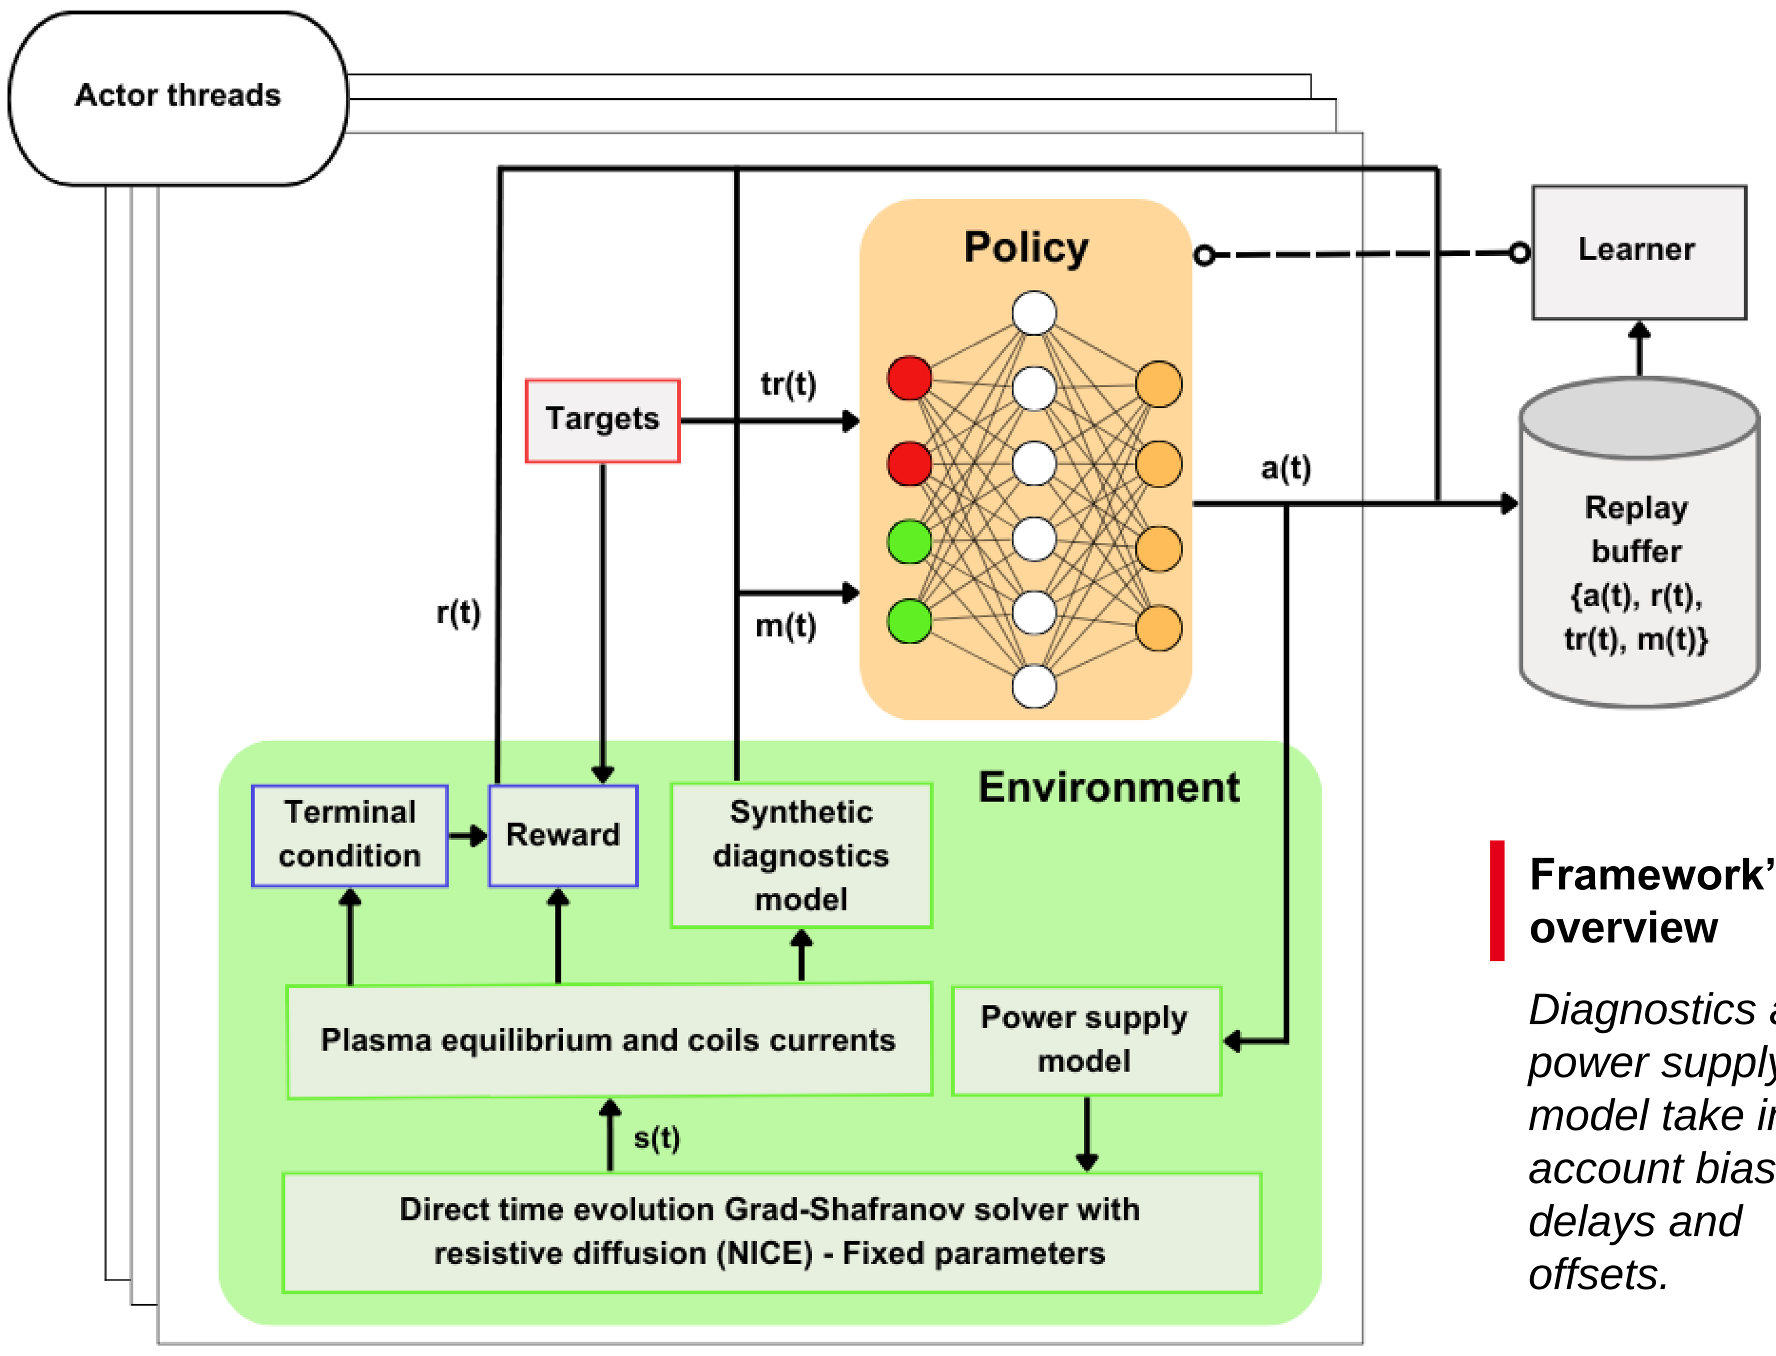

Framework's overview

Diagnostics and power supply model take into account bias, delays and offsets.

Training configuration

- *Reward*: **Follow defined plasma trajectory, and maintain initial Ip.**
- *Environment*: **NICE, a C++ free-boundary equilibrium solver** with **resistive diffusion** (η, J_{ni} fixed) [4,5]**.**

$$s(t) = \{\Psi(t),\ I_a(t)\} \rightarrow r(t)$$

- *Agent*: **MPO** [1,3], **95 MLP actors, recurrent learner,** in **Python**.

$$\underline{\textbf{magnetics}}\ m(t) \rightarrow a(t) = V_a(t)$$

- *Termination*: **Thresholds on currents, safety factor, simulated time.**

Study on baseline scenario

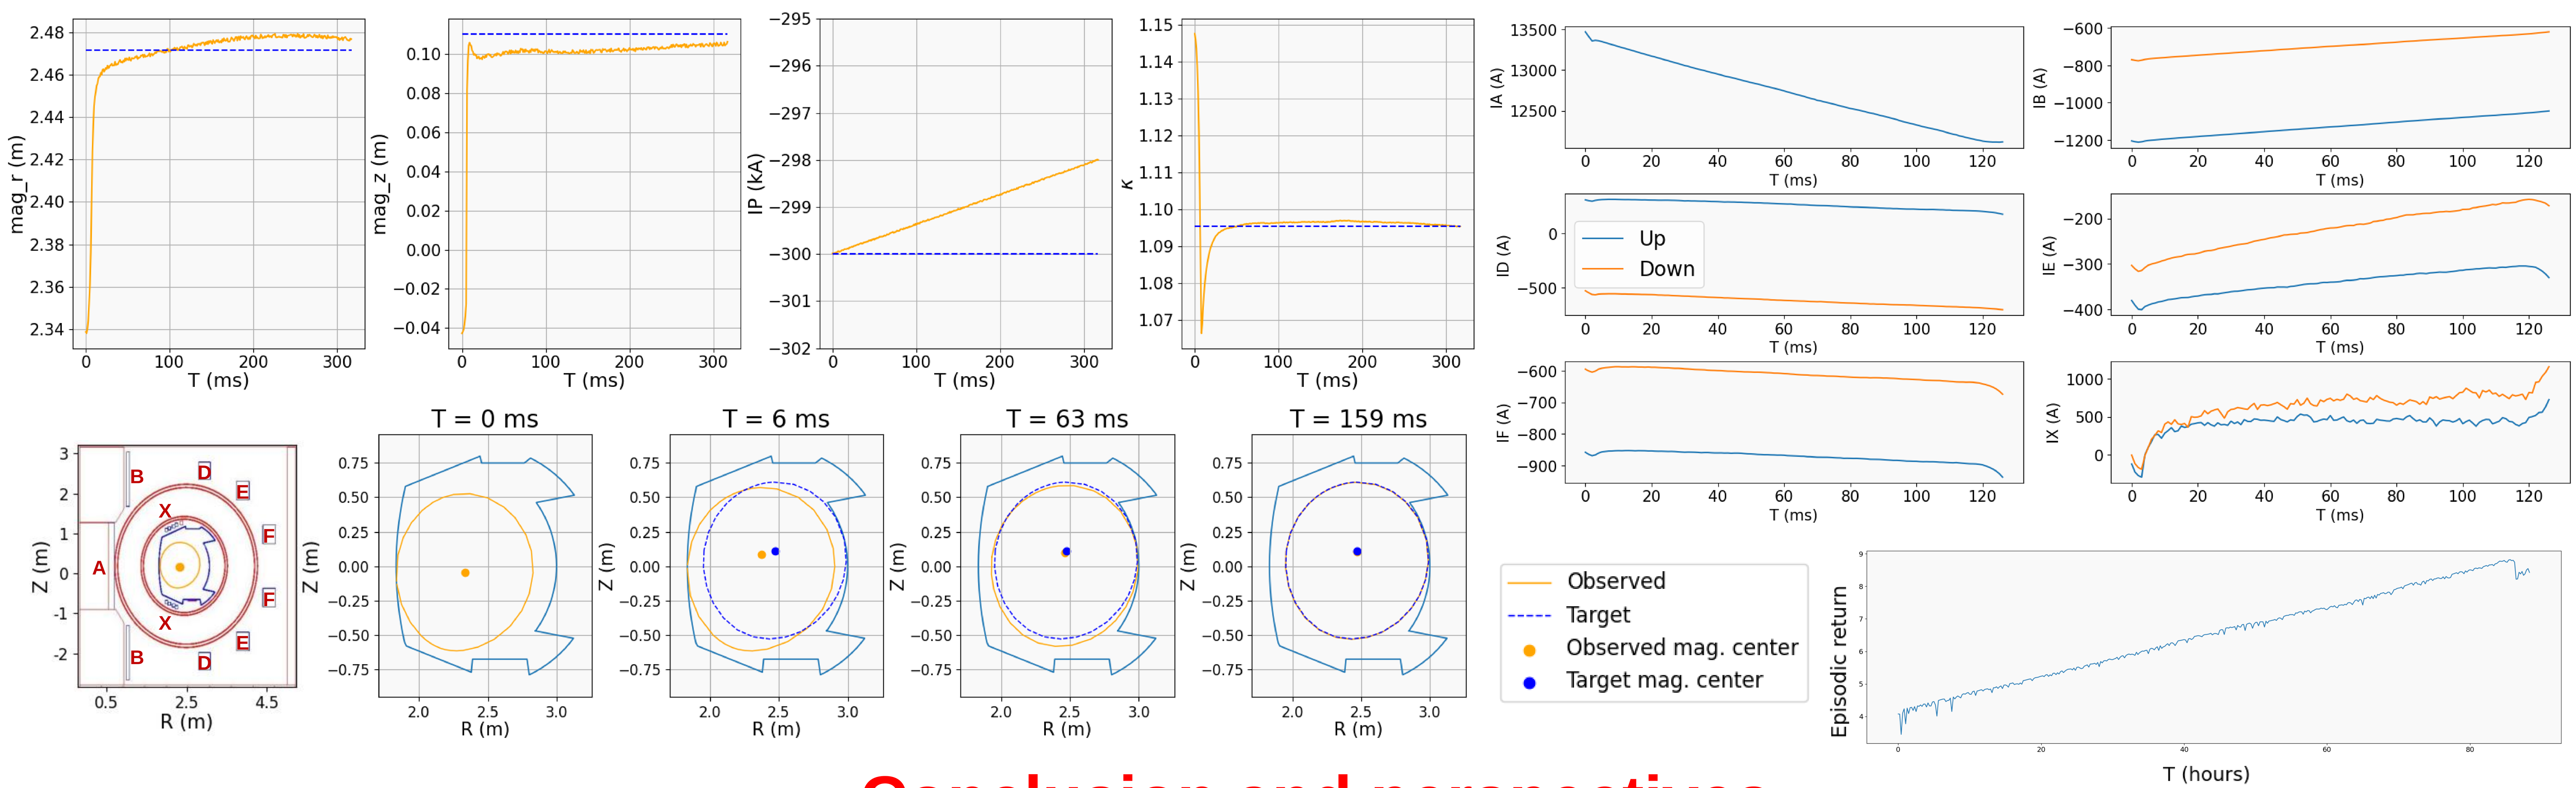

Conclusion and perspectives

- **Fast**
- **Reliable**
- **Maintainable**

Multi-language and multi-GPU framework

- **Integral control**
- **Robust reward engineering**
- **Curriculum learning**
- **Physics-informed NN**

Control of WEST plasma's shape and position

- **Non-linear tracking**
- **No need for reconstruction codes**

[1] Degrave, J., Felici, F., Buchli, J. et al., Nature 602, 414–419, 2022.
[2] W. Hoffman and B., Shahriari, et al., ArXiv abs/2006.00979 , 2020.
[3] A., Abdolmaleki, et al., ArXiv abs/1806.06920, ICLR , 2018.
[4] B., Faugeras. Fusion Engineering and Design, Volume 160, 2020
[5] H., Heumann, Journal of Computational Physics, Volume 442, 2021
[6] B. D., Tracey, et al., *arXiv preprint arXiv:2307.11546* (2023).

The authors would like to thank Federico Felici and the Fusion team at DeepMind, for their valuable advice and expertise regarding the use of reinforcement learning for magnetic control. This work has been supported by Capgemini Engineering, and the Provence-Alpes-Côte d'Azur region.