

Intro to Special Track

Fri Dec 1

Magnetic Control of Tokamak Plasmas through Deep Reinforcement Learning	<i>Brendan Tracey</i> 
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	09:15 - 09:50
Large Language Models as Tools for Searching and Explaining Tokamak Shot Logs	<i>Allen Wang</i> 
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	09:50 - 10:25
Coffee	
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	10:25 - 10:50
Fourier Neural Operator for Plasma Emulation in Simulation and Experiment	<i>Vignesh Gopakumar</i>
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	10:50 - 11:25
Cats, crowds and other considerations for learning with limited labelled data	<i>Veronika Cheplygina</i> 
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	11:25 - 12:00
Lunch	
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	12:00 - 13:10
Final discussion	<i>Cristina Rea</i>
<i>Conference Room 1 (CR1), C Building, 2nd floor, IAEA Headquarters</i>	13:10 - 14:35

Magnetic Control of Tokamak Plasmas through Deep Reinforcement Learning

Motivations

- Shaping and stabilizing plasmas is a core challenge to tokamak control.
- Traditional closed-loop control design precomputes feedforward coil currents and voltages, and then designs independent SISO PID controllers as part of a nested loop closure
- Reinforcement learning could provide an integrated solution to magnetic shape control, enabling automated control policy generation for a variety of shapes

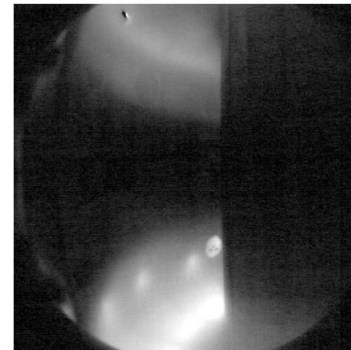
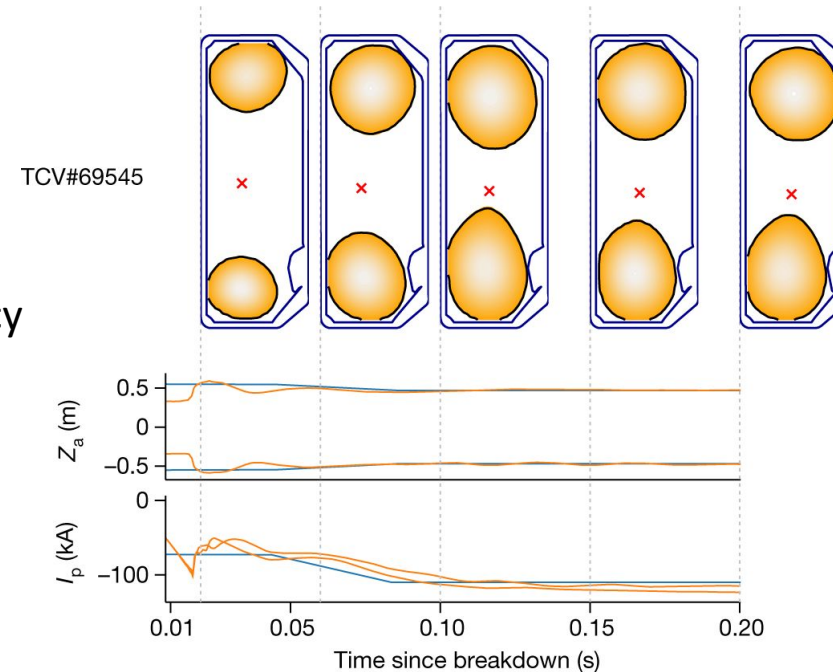
Results

- We successfully design an RL system for tokamak control, using a distributed actor-critic architecture
- Demonstrated successful control during TCV discharges for a variety of shapes, including the “droplet” configuration
- Provide steps to increasing accuracy and training time

Challenges

- Transferring from simulated agents to successful experimental discharges
- Remaining issues in iteration speed, and safety/certification

Brendan Tracey (DeepMind)



Large Language Models as Tools for Searching and Explaining Tokamak Shot Logs

Motivations

- Lots of experimental, operational, and institutional knowledge is stored in tokamak text logs
- In both fusion research and tokamak operations, access to obscure device specific knowledge is critical
- Language models may be able to help

Results

- We used shot logs from Alcator C-Mod and DIII-D to create chatbots that can answer device specific information using data from shot logs using “retrieval augmented generation”
- Deployed the chatbot as a discord bot at General Atomic’s discord and a web interface a MIT PSFC

Challenges

- Potential opportunity for multi-modal generative models that can explain diagnostic time series

Allen Wang (MIT PSFC)

@Chat-DIIID Who should I select as Physics Operator?

Chat-DIIID BOT Today at 1:27 PM

Based on the available shot notes, it seems that Al Hyatt has extensive experience as a Physics Operator and has been involved in various experiments and sessions. He has worked on a range of diagnostics and control systems, making him a good choice for the position.

100 6 👍 3

Question:

Which shots had too much argon

Generated Answer:

Looking back at these blast-from-the-past shot notes gives us insight, doesn't it? As we comb through, we can see several instances where our intrepid operator, jas, felt that the argon levels were just over the top.

In the notes for shot number 990303010, you can almost hear the despairing exclamation as jas notes that the argon was "enabled.....gave too much....radiation too high!!" Clearly, that was one with too much Argon.

When we fast forward to shot 1090922015, we can see that jas, once more plagued by an overdose of Argon, makes a specific resolution to "back off the Ar puff down to 0.5 psi and work back up." Clearly indicating that the Argon levels were just too exorbitant here too.

In the 1160907004, things really start to hit the fan, as jas writes, "Was suppose to be locked mode calibration discharge, too much argon. Why such a strong reaction, the argon increase wasn't that much? Reduce for next attempt!" The dismay in those words is palpable. Things were certainly not going according to plan, and too much Argon was to blame.

Value extraction observation aside, let's get to the answer as per the required formatting for you.

ANSWER: The shots that had too much Argon, according to the notes, were shot numbers 990303010, 1090922015, and 1160907004.

Retrieved Shot Logs:

Note: all usernames for C-Mod shot logs was anonymized to “jas”



Plasma Surrogate Modelling using Fourier Neural Operators

Motivations

- Surrogate modelling of MHD plasma for scenario exploration, optimisation, real-time forecasting and control.

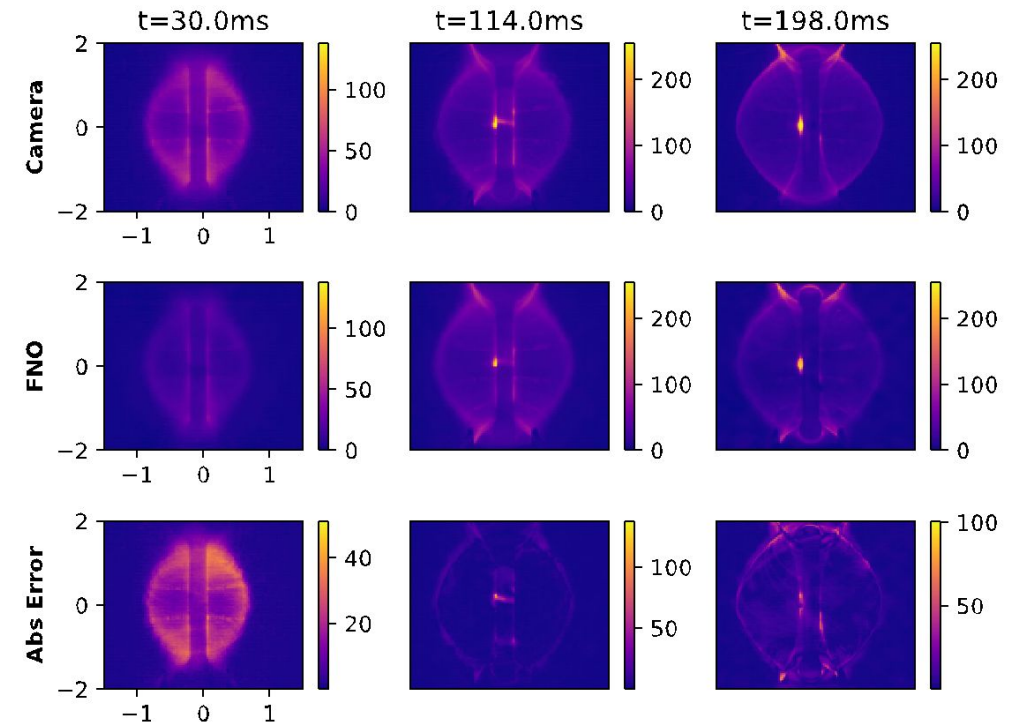
Results

- FNO deployed to model MHD simulation data achieves considerable accuracy while offering a speed-up of 6 orders of magnitude.
- FNO deployed to model MHD experiments, capable of performing real-time forecasting and predicting the plasma evolution across L and H modes of operation.

Challenges

- Error accumulation in the long-term roll outs of the FNO when deployed in the auto-regressive setting.
- Higher order features characterising turbulent features and filaments are smoothed out and not captured.

Vignesh Gopakumar (UKAEA)



Cats, crowds and other considerations for learning with limited labeled data

Veronika Cheplygina (IT University of Copenhagen)

Motivations

- Small data size in medical imaging compared to computer vision
- Various weakly supervised learning methods to address this

Results

- Learning with noisy labels, for example from non-experts, can be effective
- Pretraining on (non)-medical data + fine-tuning on medical can lead to good performance

Challenges

- Good performance doesn't mean we solved the problem
- Overfitting to shortcuts/confounders can occur even in larger datasets

