Department of Applied Mathematics and Theoretical Physics | AI4ER CDT

FedRAIN-Lite: Federated Reinforcement Algorithms for Improving Idealised Numerical Weather and Climate Models

Pritthijit Nath[1] Sebastian Schemm[1] Henry Moss[1,2] Peter Haynes[1] Emily Shuckburgh[3] Mark Webb[4]

UNIVERSITY OF CAMBRIDGE · Met Office · UKRI Engineering and Physical Sciences Research Council · UKRI Natural Environment Research Council

1. Overview

1. **Physics-based model limitations:** Traditional climate models are computationally demanding and use simplified parametrisations for unresolved sub-grid processes, which introduce systematic biases that limit their fidelity under changing conditions.

2. **Computational overheads:** Current weather prediction systems rely on complex numerical methods and grid-based models that are expensive and inflexible. Offline parameter adjustment adds further cost and cannot adapt to evolving climate states.

3. **Improving Climate Models with RL:**
   a. **Adaptive design:** Unlike supervised ML, RL agents learn by direct interaction with the climate model, maximising cumulative rewards through feedback-driven decisions.
   b. **Dynamic alternative:** RL adjusts parameters as a function of the model state, guided by rewards (comparisons against observations or re-analysis), while adhering to physical constraints.
   c. **Sparse rewards:** RL learns effectively from sparse or delayed signals, aligning with the availability of observational data at discrete intervals.
   d. **Long-term optimisation:** RL balances exploration and exploitation (during training) to refine parameterisations over time, enabling stable and scalable improvements to model fidelity.

4. **FedRAIN-Lite:** We extend single-agent RAIN (Nath et al., 2024) with a federated RL framework that learns region-specific corrections while preserving global coherence in idealised Budyko-Sellers EBMs.

[1] Nath P, Moss H, Shuckburgh E, Webb M. RAIN: Reinforcement Algorithms for Improving Numerical Weather and Climate Models. arXiv; 2024. ArXiv:2408.16118 [cs]. Available from: https://arxiv.org/abs/2408.16118.

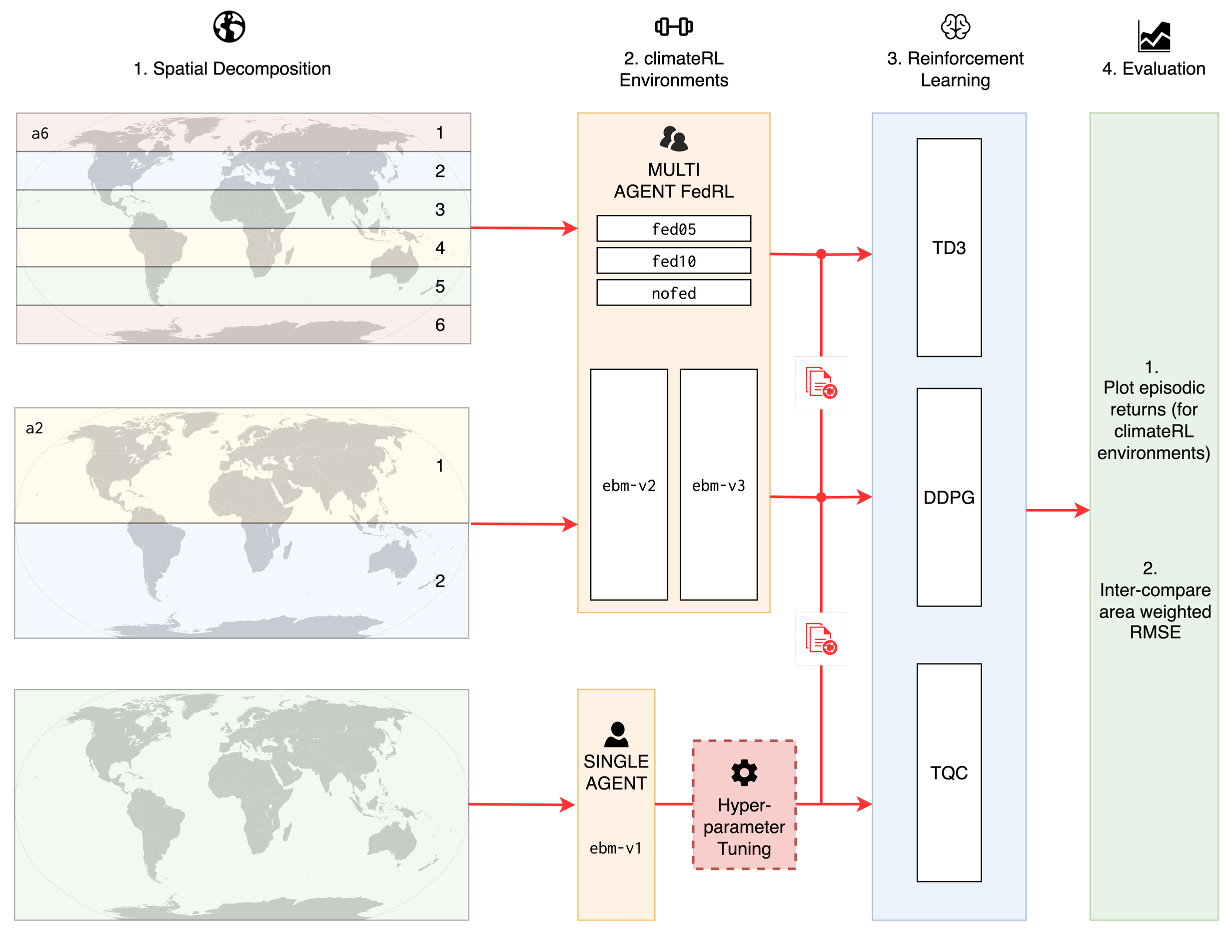

Experimental Outline: RL agents trained in single and federated settings across climateRL environments, evaluated using episodic returns and area-weighted RMSE (areaWRMSE).

3. Results

1. Federated setups (`ebm-v2`, `ebm-v3`) converge significantly faster than the single-agent baseline (`ebm-v1`), with stable training after 2.5k–5k steps compared to >10k in `ebm-v1`.

2. FedRL improves zonal skill, cutting area-weighted RMSE by over 50% in tropical and mid-latitude bands relative to static and non-federated baselines.

3. Frequent aggregation (`fed05`) yields the most consistent gains, while less frequent updates (`fed10`) show higher variance and less stable coordination.

4. DDPG consistently outperforms TD3 and TQC across all environments, with lower errors, faster convergence, and stronger generalisation across latitude bands

5. Robust performance across both coarse (`a2`) and fine (`a6`) decompositions demonstrates scalability of FedRAIN-Lite towards GCM-like architectures.

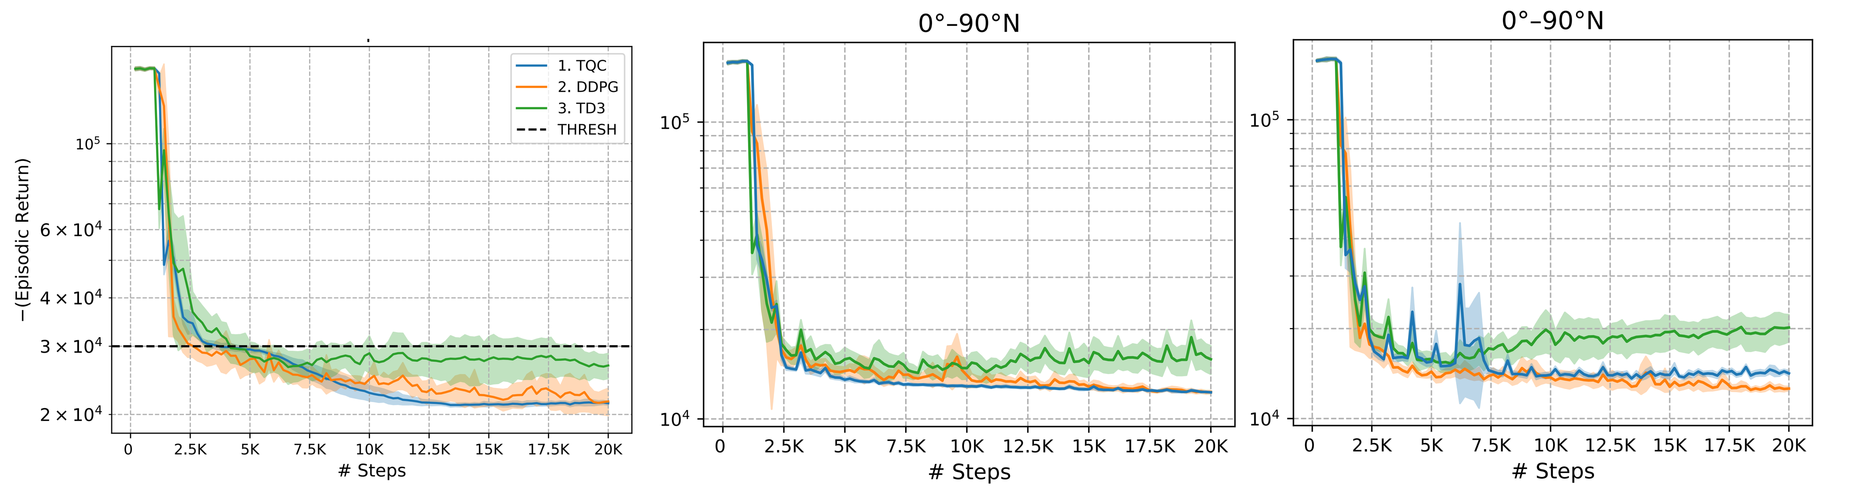

ebm-v1: single-agent, global input, global reward and latitude-specific parameters
ebm-v2: multi-agent FedRL setup with shared global profile input and local rewards
ebm-v3: multi-agent FedRL setup with sliced inputs and local rewards, mirroring GCM-like partitioning

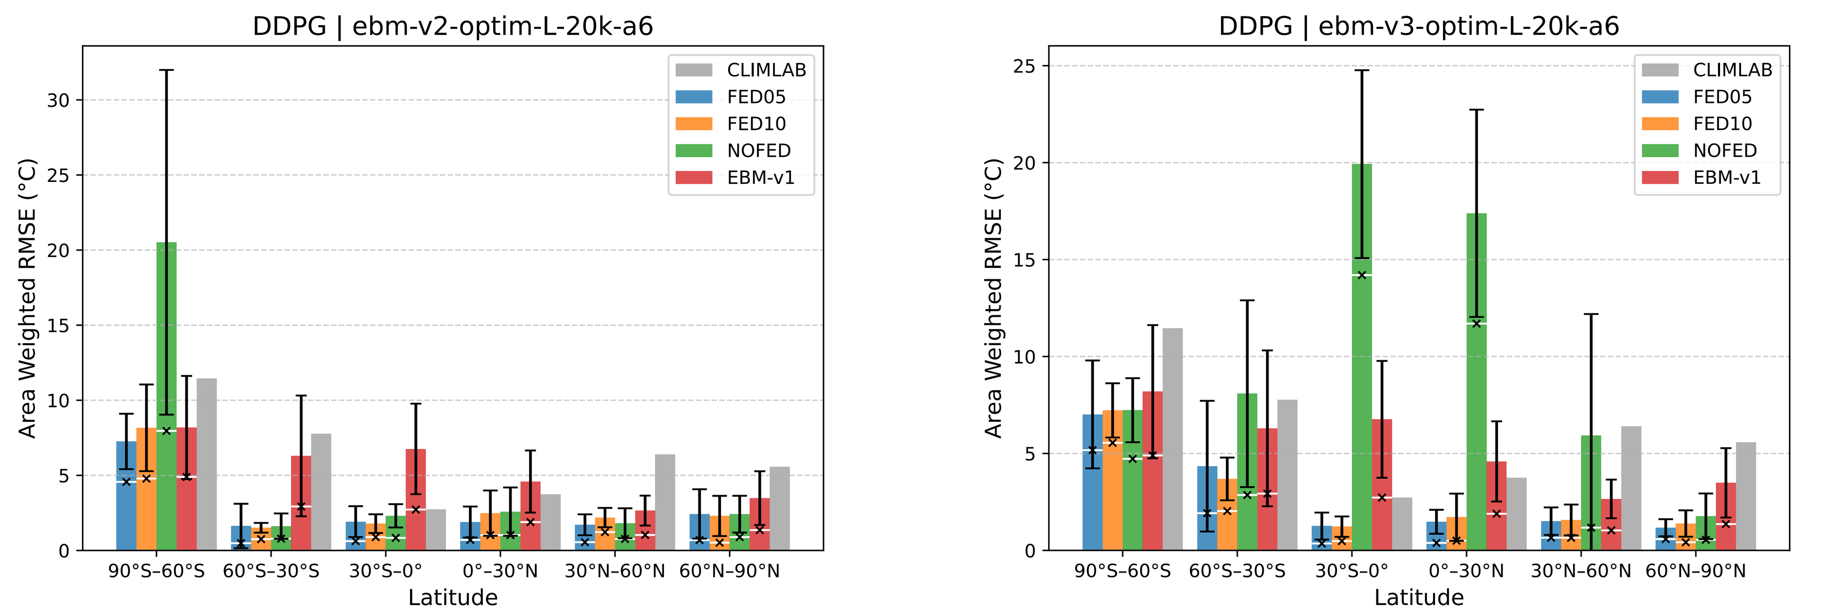

Zonal skill (areaWRMSE) of DDPG in `ebm-v2`

Zonal skill (areaWRMSE) of DDPG in `ebm-v3`

2. Federated Reinforcement Learning

1. Multiple RL agents are assigned to latitude bands, allowing each to specialise in learning region-specific correction policies that adapt to local climate dynamics.

2. In `ebm-v2`, agents optimise local rewards while observing the full zonal temperature profile, whereas in `ebm-v3`, agents receive only local temperature slices, creating a more decentralised and GCM-like setting.

3. Periodic global aggregation synchronises policy networks every *K* episodes (`fed05` or `fed10`), balancing local adaptation with global stability and enabling coherent coordination across latitude bands.

4. Decentralised design mirrors the modular, spatially decomposed structure of operational GCMs, reducing optimisation complexity and supporting geographically adaptive skill.

5. Federated coordination (FedRL) enables agents to share knowledge across regions, helping faster training and generalisation, while preserving the flexibility to adapt to regional variability.

`ebm-v1` single-agent setup: Global agent sees the full temperature profile and learns latitude-dependent OLR parameters.

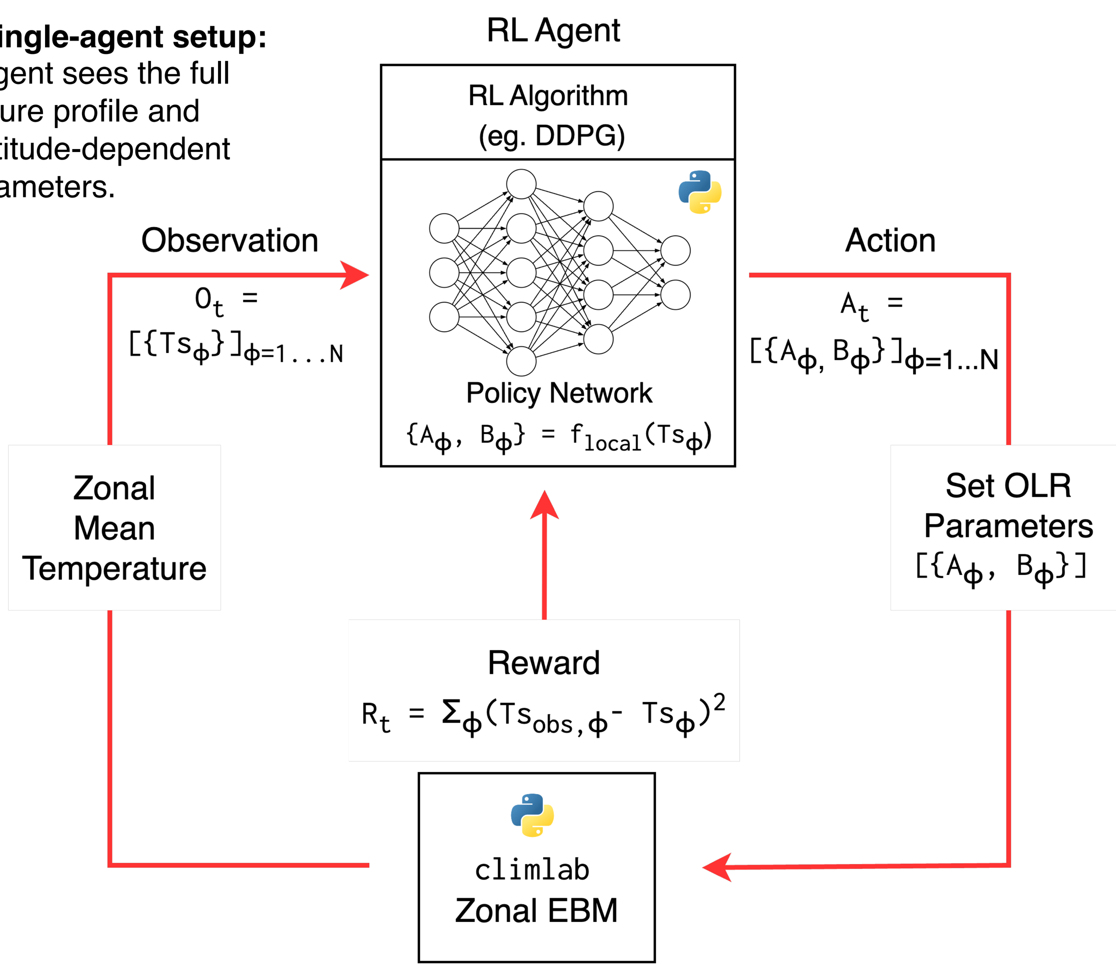

$$C(\phi)\frac{\partial T_s}{\partial t} = (1-\alpha(\phi))Q(\phi) - (A + BT_s) + \frac{D}{\cos\phi}\frac{\partial}{\partial \phi}\left(\cos\phi \frac{\partial T_s}{\partial \phi}\right)$$

`ebm-v3` multi-agent setup: Agents see local temperature slices, optimise region-specific rewards, and synchronise policies via FedRL aggregation every *K* episodes.

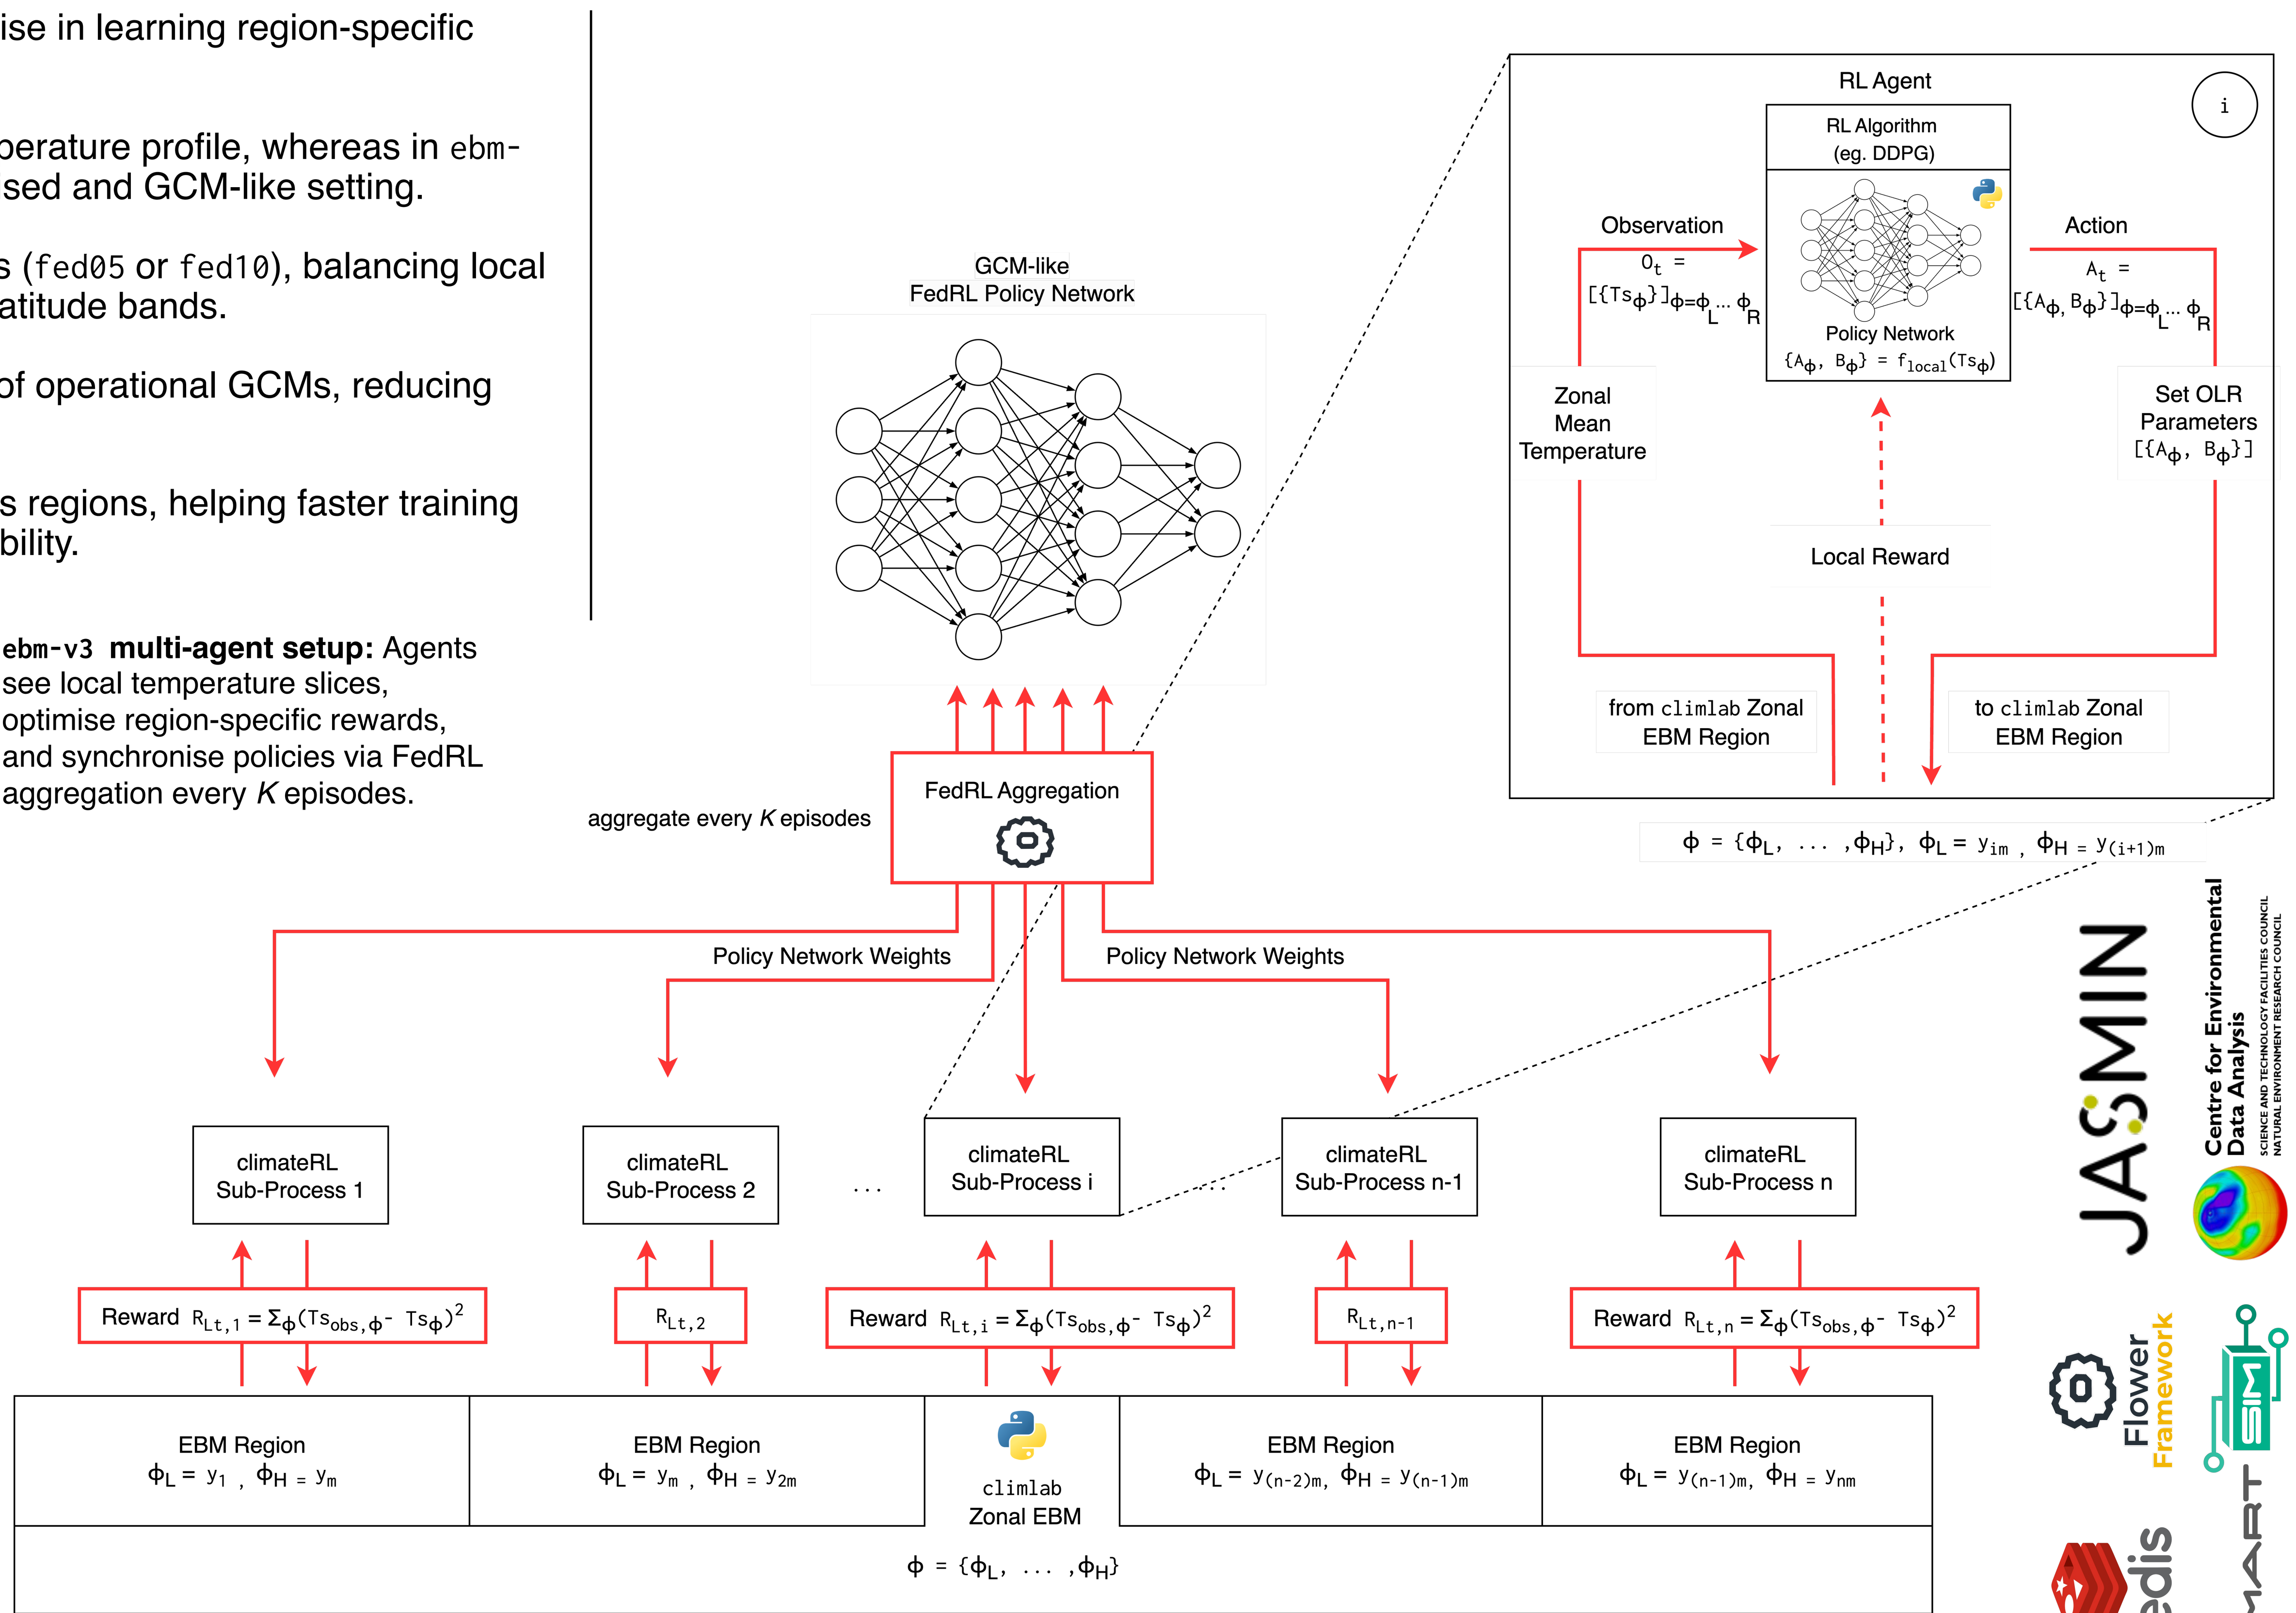

4. Conclusion

1. Combining RL with federated learning and spatial decomposition enables geographically adaptive, state-dependent parametrisations that evolve with model states.

2. FedRAIN-Lite aligns with the modular design of GCMs, allowing region-specific corrections while preserving global coordination.

3. DDPG emerges as a robust and efficient baseline, consistently achieving stable convergence and low zonal errors across all setups.

4. FedRAIN provides a scalable pathway from idealised EBMs to operational GCMs, supporting future data-driven and physically aligned climate modelling.

[1] Department of Applied Mathematics and Theoretical Physics, University of Cambridge, UK
[2] School of Mathematical Sciences, Lancaster University, UK
[3] Department of Computer Science and Technology, University of Cambridge, UK
[4] Met Office Hadley Centre, UK

P. Nath was supported by the UKRI Centre for Doctoral Training in Application of Artificial Intelligence to the study of Environmental Risks [EP/S022961/1]. Mark Webb was supported by the Met Office Hadley Centre Climate Programme funded by DSIT.

NEURAL INFORMATION PROCESSING SYSTEMS · AGU25 New Orleans, LA | 15–19 December 2025