Department of Computer Science

University of Reading

From accurate neural closures to sustainable hybrid modelling

A scientifically constrained, performance-portable use case for MONC

Sambit Kumar Panda[1] | Todd R. Jones[1] | Muhammad Shahzad[1] | Bryan N. Lawrence[1] | Anna-Louise Ellis[2]

[1] University of Reading · [2] Met Office, UK

Met Office

AFESP — Advancing the Frontiers of Earth System Prediction

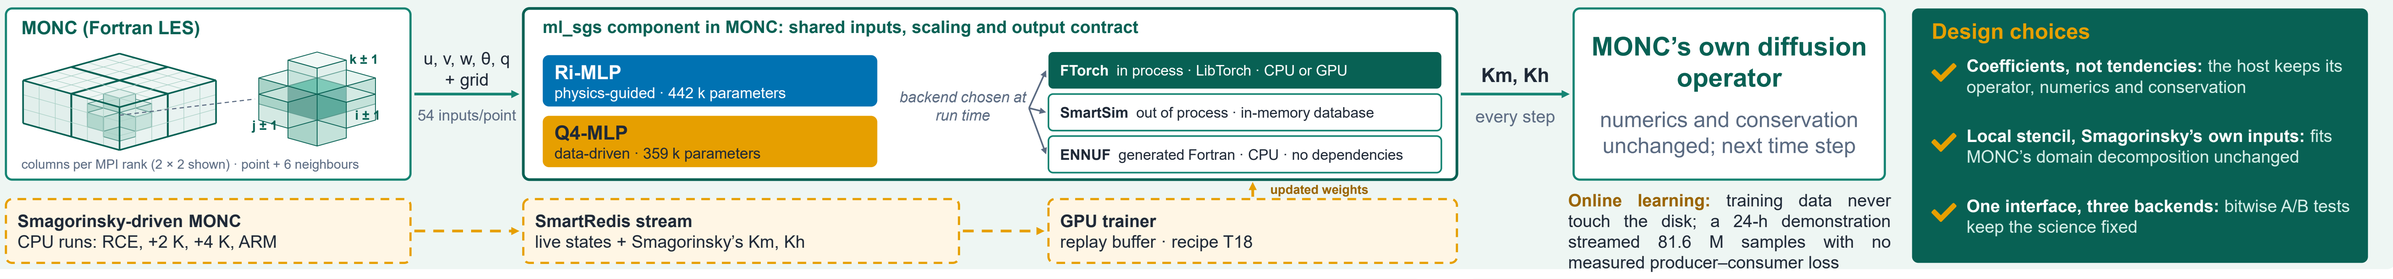
ECMWF Workshop on software strategies for sustainable physical modelling · Reading, 23–27 November 2026

When is a neural parameterisation a sustainable model component?
One worked example: MONC's subgrid closure, trained offline, coupled through three backends, run for a month and retrained while the model runs.

- **Stable when coupled** MET — 30 of 30 members to day 30, and every climate and grid tried; departures checked against RCEMIP CRMs, dynamic Smagorinsky next
- **Portable** MET — three backends, two clusters, bit-identical A/B; 63 of 63 cross-cluster comparisons overlap
- **Cost follows physics** IN PROGRESS — MONC's adaptive time step depends on each closure's flow; network compression is next
- **Evolvable** MET — retrained, distilled and bounded in place at unchanged cost; the same loop can take new regimes, grids and closures

One scientific component, three backends, one learning loop

Workshop themes: hybrid modelling · programming models · composition and sharing



Online learning: training data never touch the disk; a 24-h demonstration streamed 81.6 M samples with no measured producer–consumer loss

Design choices

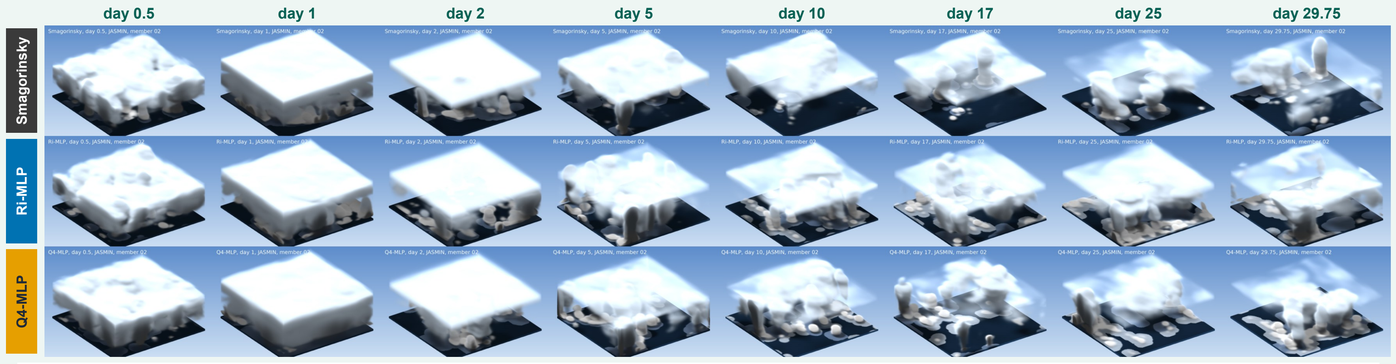
- **Coefficients, not tendencies:** the host keeps its operator, numerics and conservation
- **Local stencil, Smagorinsky's own inputs:** fits MONC's domain decomposition unchanged
- **One interface, three backends:** bitwise A/B tests keep the science fixed

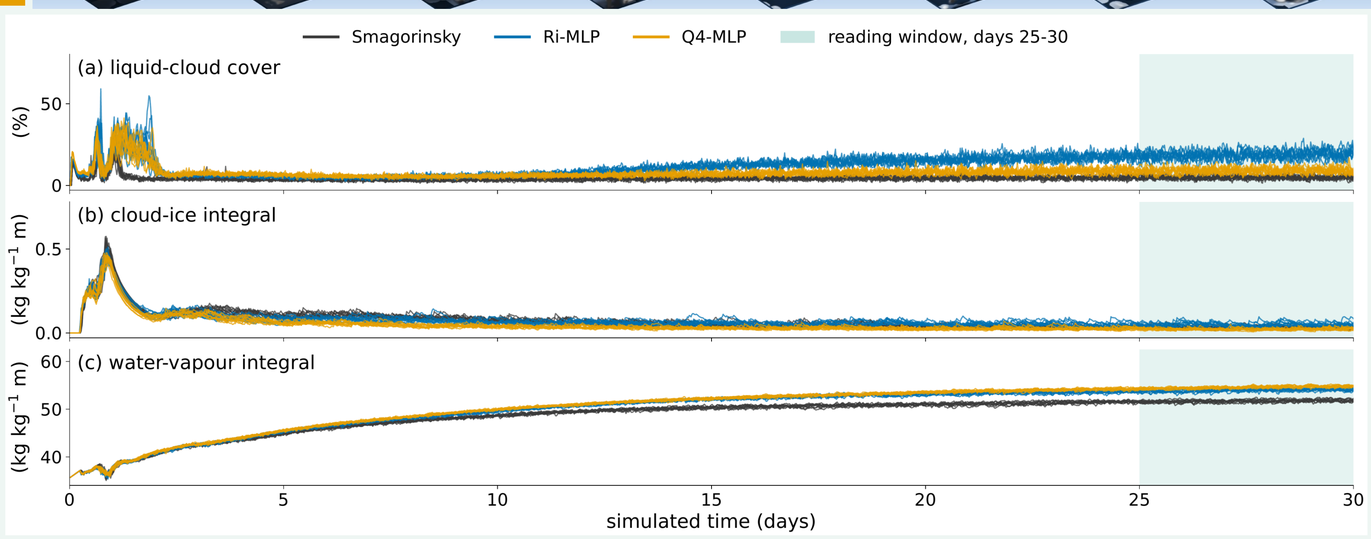
Stable when coupled

Workshop theme: hybrid modelling

trained offline, it runs for weeks in a coupled multi-physics model outside its training data; departures are judged against independent references

1 Trained offline, stable for a month in the coupled model





30 / 30 stable — members to day 30, no safeguard

+286 / +86% — liquid cloud, Ri / Q4

+0.23 / +0.35 K — θ below 2 km; vapour +4–6%

−22 / −24% — resolved KE; tops 0.8 / 2.2 km lower

- **Departures from Smagorinsky, a reference, not the truth:** liquid cloud, temperature and vapour sit beyond its member spread; ice moves in opposite directions (**+44%** Ri, **−33%** Q4); RACC2 repeats all six.
- **Independent check, 12–14 RCEMIP CRMs (300 K):** 5 km temperature inside their range, vapour just above it (48.3–48.4 vs 48.1 mm), OLR 6–11 W m⁻² below; dynamic Smagorinsky is next.
- **Spread (Jones et al. 2019):** the emulators' members agree more closely; formal tests match the range rule in **29 of 30** cells.

RCE 1 km, CASIM, JASMIN. Top: member 02 in 3-D (colour marks phase). Middle: 10 members per closure; numbers: member means over days 25–30 (shaded) against Smagorinsky.

Portable and benchmarked

programming models · operations

one interface on CPU or GPU; cost in simulated days per day

2 One interface, three backends, two clusters

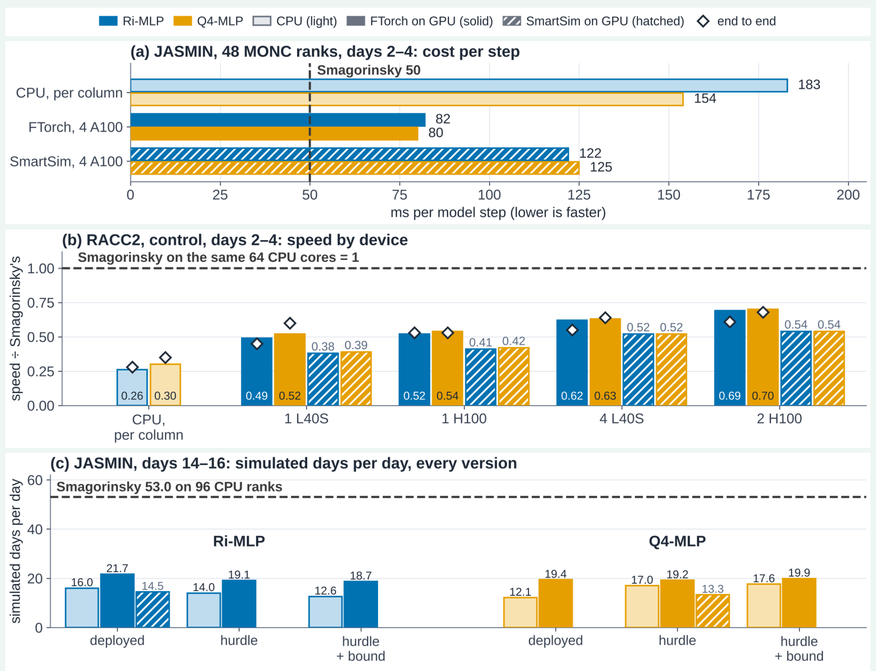
FTorch [3]
+ in process, fastest
− LibTorch build

SmartSim [4]
+ CPU + GPU apart, streams
− Redis, Python; +40–45 ms

ENNUF [5]
+ pure Fortran, Met Office HPC
− CPU only, ~5× slower

RACC2 → JASMIN: **five faults, all in the stack**, none in the network. Checks: SmartSim = FTorch bit for bit; bit-continuous restarts; an input contract. **63 of 63** cross-cluster comparisons overlap.

3 Cost follows the physics, not only the network



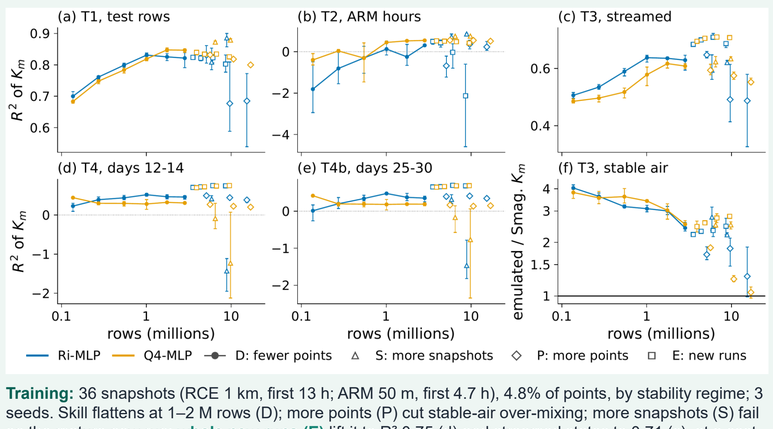
MONC limits its time step by both the flow and the mixing (advective and viscous Courant numbers), and the emulators feed their own K into that test. Each closure therefore runs at its own step (**0.49 / 0.65** of Smagorinsky's at 0.5 km), so per-step speed (a, b) is not simulated days per day (c); retraining left (c) unchanged, as in 30-day production (0.27–0.38×). **Next: distillation and symbolic regression on CPU and GPU, aiming for dynamic-Smagorinsky accuracy (4.4–4.6× Smagorinsky's cost) at a fraction of that cost.**

Evolvable, reusable, and run at scale

Workshop themes: hybrid modelling · composition and sharing · workflow orchestration

retrained, compressed and constrained while the model runs, and kept running for weeks on two clusters; each part below is a recipe others can reuse

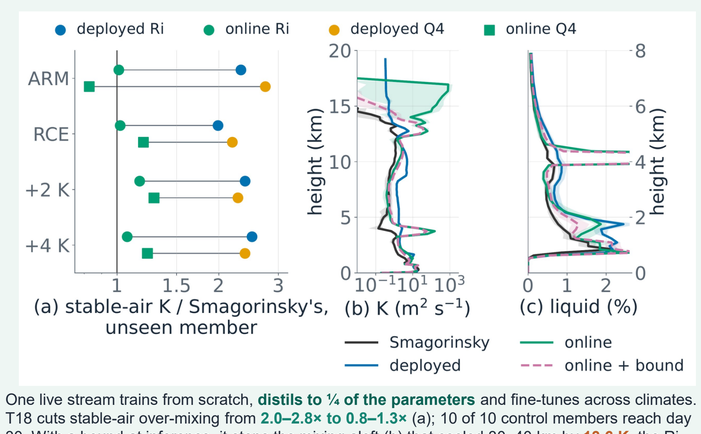
4 Sparse, mixed-resolution, mixed-regime training



Training: 36 snapshots (RCE 1 km, first 13 h; ARM 50 m, first 4.7 h), 4.8% of points, by stability regime; 3 seeds. Skill flattens at 1–2 M rows (D); more points (P) cut stable-air over-mixing; more snapshots (S) fail on the mature new run; **whole new runs (E)** lift it to R² 0.75 (d) and streamed states to 0.71 (c), at a cost on the ARM hours (b).

For the community: a few hours of sparse, mixed-resolution, mixed-regime LES, learned with multi-task heads, give a closure stable for a month when coupled; for mature states add **whole new runs** (R² 0.75 vs 0.46 / 0.31).

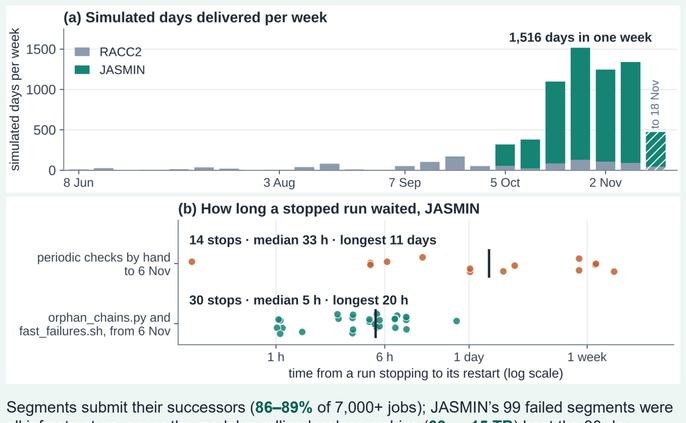
5 Learn, fine-tune and compress while it runs



One live stream trains from scratch, **distils to ¼ of the parameters** and fine-tunes across climates. T18 cuts stable-air over-mixing from **2.0–2.8× to 0.8–1.3×** (a); 10 of 10 control members reach day 30. With a bound at inference, it stops the mixing aloft (b) that cooled 30–40 km by **13.6 K**; the Ri-MLP's liquid excess (c) and a low RCEMIP OLR remain.

For the community: any running model becomes a training-data source at **no storage cost**; the same stream can fine-tune to new grids, regimes or closures and distil students for CPU-only HPC.

6 300 long runs: a workflow problem

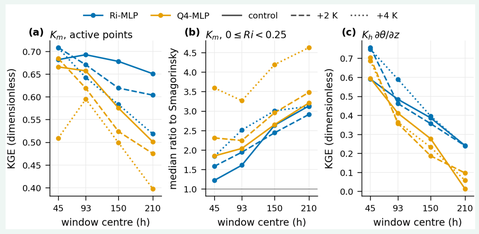


Segments submit their successors (**86–89%** of 7,000+ jobs); JASMIN's 99 failed segments were all infrastructure, never the model; a rolling lossless archive (**39 → 15 TB**) beat the 30-day scratch purge.

For the community: self-chaining segments, two small monitors and a rolling archive are a pattern any SLURM campaign can copy. **Watch for the silent stop, not the crash.**

7 Physics guidance holds up under warming

Workshop theme: hybrid modelling



Smagorinsky drives 10 days; both emulators scored on the same state.

- The **Ri-MLP** is closer in **all 12** climate × window cells (a), most at +4 K; heat-flux KGE 0.41–0.51 vs 0.30–0.31 (c).
- Both drift as the state leaves the training hours: stable-air Km 2.1–2.7× (Ri), 2.4–3.9× (Q4) (b).

Same state, Km KGE, 15 set-ups: accuracy and stability are separate

1 km, spin-up	1 km, days 25–30	unseen 2 km	unseen 0.5 km	coupled
0.56–0.67*	**Ri 0.58 / Q4 0.33**	**0.14–0.45**	**Ri < 0 / Q4 ≈ 0.1**	**all 15 stable**

* Q4-MLP at +4 K on the two smaller domains: 0.34, 0.43.

8 AI-assisted development, made checkable

Workshop theme: agentic coding

Researcher sets questions, approves, communicates → **Pre-registration** 52 dated plans, decisions, readings → **Agents** Claude Code on both clusters; Codex on RACC2 → **Runs and analysis** port, build, submit, monitor; figures, drafts → **Shared record** 579-commit log, 585-commit repo, 41 rules → **Cross-cluster review** the second cluster checks the results

next question

Errors the record made findable:

- a rebuild without microphysics, silent for a month
- a scaler that hid cloud liquid from 18 emulators
- a testcase droplet number in the wrong unit
- confident claims later found wrong

Agents speed up the work and make confident errors at the same pace; pre-registration, the shared record and a second cluster's review make those errors findable. Guardrails: the agent cannot cancel jobs on JASMIN; each cluster writes only its own folder.

Lessons for sustainable hybrid modelling

1. **Coefficients, not tendencies** — the host keeps its operator, numerics and conservation
2. **One component, many backends** — chosen at run time; bitwise A/B tests keep the science fixed
3. **Portability lives in the stack** — every fault we met was outside the network
4. **Test the contract, run long** — input contracts, borrowed set-ups, weeks of coupling
5. **Cost follows the physics** — the adaptive time step depends on each closure's flow: count simulated days per day
6. **Operate with a record** — monitors, a purge-safe archive, logs that make agents' errors findable

Next: registered days 25–30 readings of every month-long family (rechecks at days 45–50) · completing the bounded runs · month-long runs of the new-run training sets · distillation and symbolic regression · Lagrangian dynamic Smagorinsky as a second reference

Outlook: the same mixed-resolution, mixed-regime multi-task set-up is a baseline for learning costlier closures (dynamic Smagorinsky, the boundary-layer and turbulence schemes of weather and climate GCMs) and, with more regimes, a foundation model for subgrid closure.

References

[1] Panda et al., Physics-guided multi-task learning for subgrid scale turbulence parameterization, IJCNN 2026, IEEE. doi:10.1109/IJCNN69854.2026.11728747

[2] Panda et al., Physics-aware multi-task learning for atmospheric turbulence parameterization, PASC '26, ACM. doi:10.1145/3815572.3815750

Ensemble-spread analysis after Jones et al. (2019). In preparation: JAMES (month-long coupling, online evaluation, data scaling, fine-tuning); a cost and portability paper.

Software

[3] Atkinson et al. (2025), FTorch: a library for coupling PyTorch models to Fortran, J. Open Source Softw. 10(107), 7602. doi:10.21105/joss.07602

[4] Partee et al. (2022), Using machine learning at scale in numerical simulations with SmartSim: an application to ocean climate modeling, J. Comput. Sci. 62, 101707. doi:10.1016/j.jocs.2022.101707

[5] Reid, H. and Met Office (2025), ENNUF: Easy Neural Networks for Use in Fortran, v1.1.0. github.com/MetOffice/ennuf

Acknowledgements

Met Office (CASE studentship, HPC); AFESP, University of Reading; JASMIN and its ORCHID GPU cluster; RACC2; the FTorch (ICCS, Cambridge), SmartSim and SmartRedis, and ENNUF developers; Claude (Anthropic), used as a coding and analysis agent.

Contact

s.panda@pgr.reading.ac.uk

Papers, code and data:

IJCNN paper · IJCNN data · IJCNN code · **PASC paper** · PASC data · PASC code