

Plastic recurrent excitation reaches the old PING regime

exp073 · 20 July 2026 · Draft

Abstract

This exploratory experiment tested the next biologically plausible path after exp071/exp072: keep the cumulative-potential readout, keep the historical one-seed validation-only SHD split, and make recurrent $E \rightarrow E$ excitation both nonzero and plastic while preserving Dale signs. The original matched COBA/PING design still failed at the required local smoke gate: PING remained finite and active, but the matched COBA cell repeatedly produced skipped/non-finite updates even after reducing the requested W_{EE} scale from 0.3 0.01 through 0.03 0.01, 0.003 0.001, and finally 0.0003 0.0001. After that gate failure, an explicitly authorized PING-only continuation ran the surviving cell on Modal. The first eight-epoch short stage reached 58.33% validation accuracy; after reviewing that trajectory, the user authorized the forty-epoch continuation. The final40 run selected 71.32% held-out validation accuracy at epoch 39 and finished at 68.38%. W_{EE} remained nonnegative and trained away from its tiny initialization. The result is a modest PING-only recovery to the old 70% regime, not a matched COBA/PING result and not a clear jump to a new high-accuracy regime.

Goal prompt

```
/goal Design and execute a new exploratory SHD experiment on a new branch/PR,
testing whether nonzero Dale-constrained trainable recurrent excitation moves
the matched COBA/PING models beyond the current ~70% regime.
```

```
Use current main as the base. Build exp073 around the best surviving SHD
recipe from exp071/exp072, especially the cumulative-potential readout. Do
not rerun the old baseline unless needed for debugging; use exp071/exp072 as
the historical baseline.
```

```
Experiment:
```

- run matched COBA and PING cells with nonzero W_{EE} initialization and --trainable-w-ee;
- keep W_{in} and W_{out} trainable as in the current training path;
- keep Dale signs enforced; do not use free signed recurrence;
- start with a modest nonzero W_{EE} initialization, e.g. --w-ee 0.03 0.01, unless local inspection of current parameter scales suggests a safer nearby value;
- keep input preprocessing, dataset split, readout, optimizer, batch size, and evaluation protocol matched across COBA and PING;
- keep the comparison exploratory and single-seed unless I explicitly ask for

replication.

Protocol:

- implement through experiments/exp073.py and existing tools/snn CLI capabilities if possible;
- do not edit tools/snn unless the current CLI cannot express nonzero trainable W_{EE} cleanly; if a tools/snn change is required, stop and explain exactly why before editing;
- run a local smoke/scout stage first with few epochs/samples to confirm finite loss, active spiking, and W_{EE} actually changes during training;
- if both cells are finite and active, run a short RunPod exploratory stage, fewer than 40 epochs for iteration speed;
- if accuracy looks meaningfully better than the exp071/exp072 historical ~70% regime without pathological firing or NaNs, run a 40-epoch confirmation for the promising cell(s);
- COBA and PING may run in parallel on separate RunPod pods when useful;
- RunPod spending is authorized up to \$40 total; reap all pods when finished.

Deliverables:

- artifacts/data/exp073 with provenance, numbers.json, reproducer, training curves, firing-rate diagnostics, W_{EE} diagnostics, and matched rasters where relevant;
- writings/exp073.typ as a cold-readable Demolab experiment report comparing exp073 to the historical exp071/exp072 baseline;
- update the Spiking Heidelberg Digits collection so exp073 is the next canonical experiment;
- build successfully;
- commit focused implementation/results changes, push the branch, and open a new PR;
- do not merge unless I explicitly ask.

End at the review gate with:

- best/selected COBA and PING accuracies;
- whether trainable nonzero W_{EE} appears to move us beyond the ~70% regime;
- W_{EE} diagnostics showing it was nonzero and trained;
- firing-rate/pathology summary;
- exact compute spend;
- validation performed;
- PR link and rendered exp073 link.

Methods

The intended comparison inherited the exp071/exp072 SHD recipe: seed 42, 7340 development-training utterances, 816 held-out validation utterances, batch size 32, Adam learning rate 0.0004, 1 ms simulation steps, 1000 ms windows, matched input preprocessing, and the cumulative-potential readout with signed abstract readout weights and a trainable bias. The official SHD test set remained sealed and inaccessible to the runner.

The only intended new scientific ingredient was `--trainable-w-ee` with nonzero `--w-ee`. Dale's law stayed enabled; W_{EI} , W_{IE} , and W_{II} stayed frozen. COBA kept the no-inhibitory-

loop recipe and no voltage-gradient dampening; PING kept the registered inhibitory loop and dampening 1000.

During implementation, the local CLI accepted `--w-ee` but training did not pass it into model construction. Commit 247e186 fixed that plumbing and added a focused regression test before this experiment ran. A later transport fix moved exp073's Modal worker function to module scope and aligned its Python image with the host, without changing any scientific parameter.

Local gate


The final registered local scout used the same nonzero W_{EE} initialization now recorded in the PING-only configuration: `--w-ee #r.config.w_ee_init.at(0) #r.config.w_ee_init.at(1)`.



The gate required both matched cells to be finite, active, non-saturated, and free of skipped or non-finite updates before any paid matched run. COBA failed that finite-update criterion while PING passed, so the locked matched COBA/PING result could not be promoted. The cloud result below is therefore the later authorized PING-only continuation, not a matched-network claim.

Authorized PING-only continuation

After the matched local gate failed, the user authorized a design pivot: continue with the surviving PING cell only, without recasting it as a matched COBA/PING result. This continuation uses the same one seed, development-training/validation split, input preprocessing, cumulative potential readout, nonzero trainable W_{EE} , and PING dampening 1000. The eight-epoch short run reached 58.33%; after trajectory review, the user authorized the forty-epoch continuation recorded here.

Cell	Selected acc.	Final acc.	Epoch	Final E rate	Final I rate
PING	71.32%	68.38%	39	19.52 Hz	55.54 Hz





The PING-only W_{EE} diagnostics show selected mean 0.01463611 and mean absolute movement from initialization 0.0146355.

The selected checkpoint's W_{EE} norm grew without violating Dale's nonnegativity constraint. The run was dispatched on modal with GPU A10G. The timestamped compute ledger records 2026-07-23T12:30:20.490Z to 2026-07-23T15:22:25.613Z UTC, estimated billable GPU time 10291.6 s, and estimated spend 3.149 USD. This is a timestamp estimate pending provider-billing reconciliation, not an exact invoice.

Timestamped result checkpoint

At 2026-07-23T15:22:25.613Z UTC, the user-authorized final40 result selected 71.32% validation accuracy at epoch 39 and finished at 68.38%. The run was finite and clean in the runner's summary, but the activity curves and live log show high activity with saturation warnings during training, so the result should be read as accuracy recovery with a dynamical caveat rather than a free improvement.

Conclusion

The locked matched experiment still failed at the local gate: COBA could not be honestly promoted. The later PING-only run is therefore exploratory continuation evidence for the stable PING recipe, not matched COBA/PING evidence. Its 71.32% selected validation accuracy reaches the exp071/exp072 historical 70% PING regime and slightly edges over it on this one seed, but the final epoch falls back to 68.38% and the high-activity/saturation diagnostics make this a cautious exploratory signal, not a defensible high-accuracy claim.

Paid compute spend: 3.915 USD. Provider billing exact: **false**. Active RunPod pods after this checkpoint: 0.