

An anonymous repository with all code and per-dataset notebooks/markdowns is available at https://github.com/wbsnn2025/wbsnn_for_reviewer_iclr2026. We pin exact package versions and provide one-command scripts to regenerate all tables from a clean clone. Preprocessing uses train-only standardization and PCA; fixed splits, seeds, discovery budgets, and hyperparameters are documented in the repo README.