AUTHORS

Responsible Professor: Tom Viering
Supervisor: Gido van de Ven
Student: Oskar Hage

Mind the Gap: Layerwise Proximal Replay for Stable Continual Learning

TU Delft

01. Introduction

- **Continual learning** trains models on a sequence of tasks.
- A major challenge is catastrophic forgetting, where learning new tasks causes a drop in performance on earlier ones.
- A sharper, short-term version of this is the **stability gap**, which appears immediately after switching tasks (shown in Figure 1 below).
- This occurs even with methods like Experience Replay (ER).
- This study explores whether **Layerwise Proximal Replay** can help reduce the stability gap.

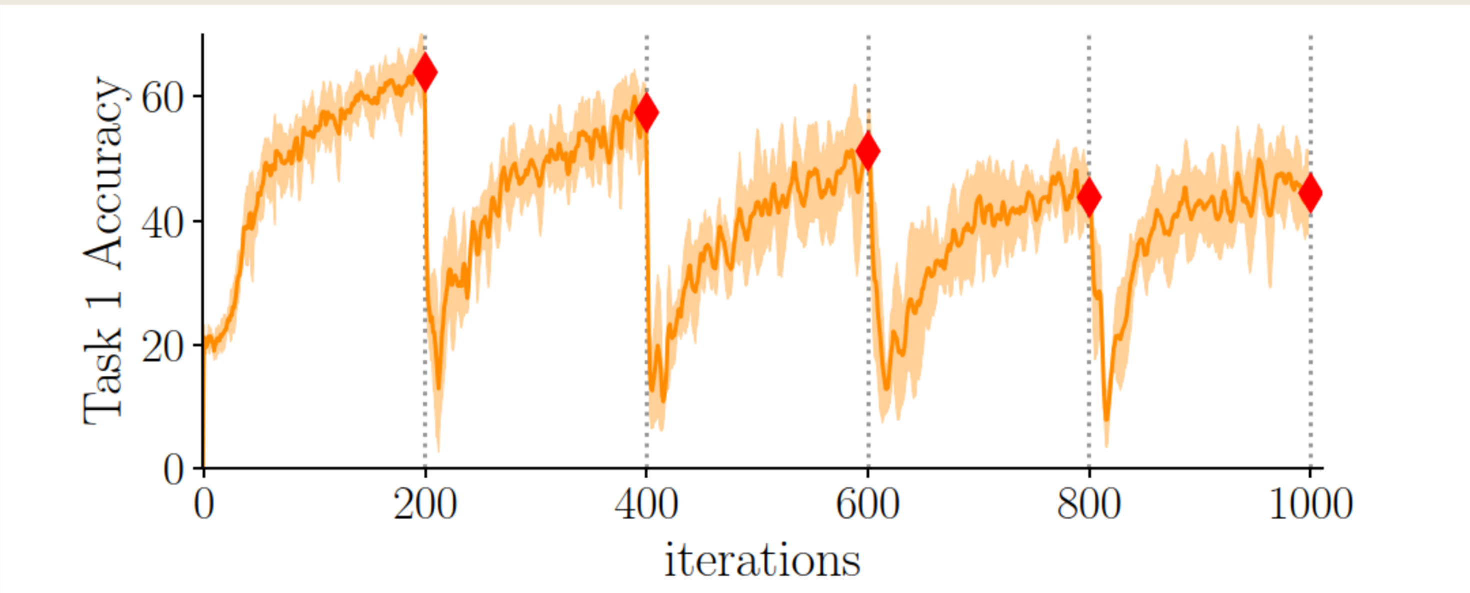

Figure 1: The Stability Gap as shown in [1].

02. Layerwise Proximal Replay (LPR)

- Introduced by Yoo et al. (2024) as a stability-focused continual learning method [2].
- Uses replayed data to estimate how each layer responds to inputs.
- Computes a per-layer pre-conditioner matrix to guide gradient updates.
- Encourages weight changes that preserve useful internal representations.
- Operates entirely at the layer level, without modifying the overall architecture.
- Parameter ω dictates strength of the regularisation.

03. Research Questions

- Does LPR aid in mitigating the stability gap, and how does changing ω affect the results?
- What is the relation between learning rate and ω?

04. Methodology

- We use Rotated MNIST as a continual learning benchmark with three tasks, each defined by a fixed image rotation (0°, 80°, 160°).
- A multi-layer perceptron is trained incrementally on all tasks using full replay.
- We then implement LPR on the same architecture.
- Test accuracy on all tasks is tracked across training to visualise and quantify the stability gap in both settings.
- We test on learning rates = [0.01, 0.1, 0.5, 1] and ω = [0.01, 0.1, 0.5, 1.0, 10.0]

05. Results

- Results show that LPR stabilises the drops test accuracy after after task switches. (Figure 2)
- Compared to the baseline, LPR improves the most at higher learning rates, up to 45 percentage points. (Table 1)
- Higher values of ω limit the plasticity of the model, inhibiting learning on future tasks. (Table 2)
- LPR achieves lower accuracy but also lower variance in final task accuracy compared to the baseline. (Table 2)

Table 1: Best LPR configuration per learning rate. Values reflect mean ± SEM over 5 runs. The last column reports the absolute reduction in stability gap relative to the baseline. Lower is better

Learning Rate	Best ω	Mean Gap (%)	SEM (%)	Reduction (p.p.)
0.01	1.0	0.09	0.20	**0.13**
0.10	0.01	1.38	1.64	**13.33**
0.50	0.01	0.65	0.55	**47.47**
1.00	0.01	0.63	0.61	**45.74**

LR	Baseline	$\omega = 0.01$	$\omega = 0.1$	$\omega = 0.5$	$\omega = 1.0$	$\omega = 10.0$
0.01	**71.66 ± 0.16**	60.05 ± 0.19	50.31 ± 0.15	47.64 ± 0.12	46.90 ± 0.27	44.44 ± 0.26
0.10	**94.74 ± 0.03**	90.57 ± 0.05	86.76 ± 0.07	83.27 ± 0.07	81.29 ± 0.12	75.09 ± 0.30
0.50	96.51 ± 0.05	**96.10 ± 0.03**	94.66 ± 0.02	93.13 ± 0.07	92.69 ± 0.04	89.79 ± 0.13
1.00	78.86 ± 15.46	**96.26 ± 0.15**	94.58 ± 0.49	93.77 ± 0.47	93.10 ± 0.39	89.36 ± 3.43

Table 2: Final test accuracy (%) across learning rates for the baseline and LPR with different ω values. Values are reported as mean ± SEM over 5 runs. Best result per learning rate is bolded.

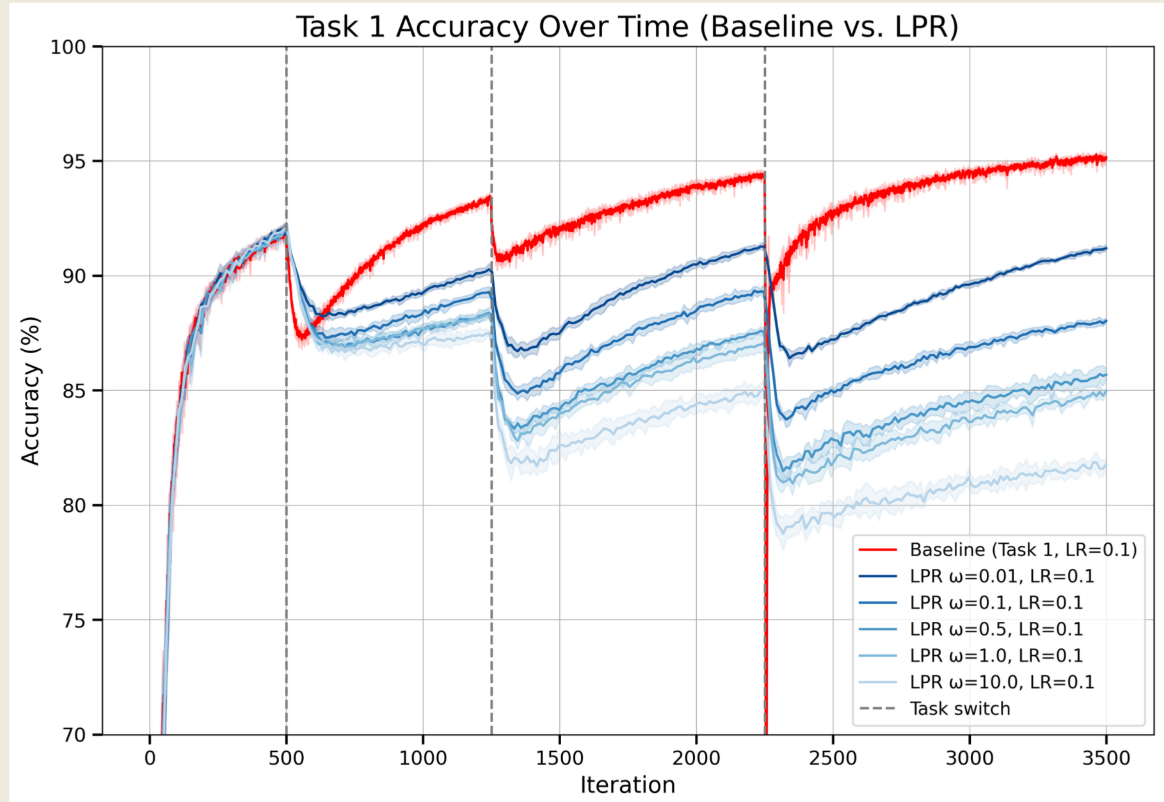

Figure 2: Task 1 accuracy over time for the baseline (in red) and LPR variants (in shades of blue) at learning rate η = 0.1. Lighter lines correspond to higher values of ω. Dashed vertical lines indicate task transitions. Note the y-axis starting at 70% accuracy and Task 1 dropping to below 70 after Task switch 3 (49.60%).

06. Conclusion

- LPR consistently reduces the stability gap after task switches.
- Benefits are most pronounced at higher learning rates, where baseline training is unstable.
- Low-to-moderate ω values (0.01–0.1) achieve the best trade-off between stability and adaptability.
- High ω reduces plasticity, leading to underfitting on new tasks and lower overall performance.
- LPR enables stable continual learning without requiring a lower learning rate.

07. Future Work / Limitations

- Extend experiments to deeper architectures and more complex continual learning benchmarks (e.g., class-incremental scenarios).
- Explore adaptive or task-aware tuning of ω, and optimise the efficiency of preconditioner updates for online settings.
- **Limitation:** Simplified LPR settings

Bibliography

[1] De Lange, M., van de Ven, G. M., & Tuytelaars, T. Continual Evaluation for Lifelong Learning: Identifying the Stability Gap. KU Leuven.
[2] Yoo, J., Liu, Y., Wood, F., & Pleiss, G. Layerwise Proximal Replay: A Proximal Point Method for Online Continual Learning. Proceedings of the 41st International Conference on Machine Learning,