


LIFELONG REPRESENTATIONS: A SURVEY ON CONTINUAL SELF-SUPERVISED LEARNING FOR VISION MODELS



Sergi Masip, Alicja Dobrzeniecka, Jonathan Swinnen, Joachim Collin, Bartłomiej Twardowski, Szymon Łukasik, Tinne Tuytelaars


MOTIVATION

Access to manual labeled data in real-world deployments might **not be possible** or **too costly**.



WHY SSL RESISTS FORGETTING

Task-Agnostic Representations




Supervised

vs.



Self-Supervised

Flatter loss landscapes



⚠ BEWARE

01

Compute & data hunger

02

Plasticity struggles







03

Risk of overfitting

04


Cross-modal drift

TAXONOMY OF METHODS




OPEN CHALLENGES


Evaluation Standardization




Rethinking Stability-Plasticity



Generalization Across SSL Objectives



Scaling to Foundation Models



CONCLUSIONS

SSL objectives resist forgetting via task-agnostic representations and flatter loss landscapes, but mitigation strategies are still needed. The field must also consider plasticity and move beyond small-scale benchmarks toward foundation model scale.

This paper has received funding from the Flemish Government under the Methusalem Funding Scheme (grant agreement n° METH/24/009).

ACKNOWLEDGEMENTS