DIPARTIMENTO DI
INGEGNERIA DELL'INFORMAZIONE
ELETTRONICA E TELECOMUNICAZIONI

SAPIENZA
UNIVERSITÀ DI ROMA

Review of PhD thesis "Continual Learning Without Access to Past Data" (Grzegorz Rypeść, Warsaw University of Technology)

The thesis proposes novel continual learning (CL) methods for the scenarios of incremental CL with no past memory (*exemplar-free*) and possibly unsupervised new classes (*continual category discovery*). These are challenging setups, making the three contributions in the thesis valid and important to the field.

This is also shown by the track record of publications, including three first-author publications in top-tier AI and computer vision conferences (NeurIPS, ICLR, ECCV). Overall, this immediately places the author among the most active and visible early-career contributors to continual learning research in the last few years.

The thesis itself consists of an introduction, state-of-the-art, and conclusion chapter, with the three papers "sandwiched" in the middle. While I imagine this format is formally acceptable (since the thesis was sent under review), it makes for a somewhat fragmented reading experience. Chapters change templates, fonts, and formatting conventions, there are multiple bibliographies throughout the document, and cross-references between the introductory chapters and in-between the included papers are relatively limited. I would have preferred a more homogeneous manuscript, but I do not consider this aspect relevant for the scientific evaluation.

Overall, the three contributions focus on exploiting statistical properties of latent representations in order to perform continual learning without storing past samples.

- The first contribution (SEED) proposes a selective ensemble-learning strategy in which only one expert is updated for each incoming task. Expert selection is driven by the geometry of latent representations modeled through Gaussian distributions. This allows the method to maintain plasticity while reducing interference between tasks and avoiding the parameter growth typically associated with expert-based continual learning approaches.

- The second contribution (CAMP) addresses continual category discovery, where only part of the incoming data is labeled and novel classes must be discovered autonomously. The key insight is that feature distillation and category adaptation are complementary rather than competing mechanisms. The resulting method successfully balances representation stability with the flexibility required to discover new semantic categories.

- The third contribution (AdaGauss) studies the evolution of class distributions in latent space during continual learning. The author identifies covariance drift and dimensional collapse as important sources of forgetting and task-recency bias. Building on this observation, AdaGauss introduces explicit adaptation of class statistics and a dedicated anti-collapse regularization strategy, leading to strong empirical results.

Concerning the "new" chapters, I have a few comments:

1. A more comprehensive state of the art on continual learning (beyond a presentation of the most relevant baselines) would have improved the thesis considerably. The broader landscape of rehearsal-based methods, parameter-efficient adaptation methods, dynamic architectures, and foundation-model-based continual learning is only briefly discussed.

2. For Chapter 2, some technical elements such as the Fisher Information Matrix are introduced with limited detail. A more formal treatment or additional references would improve readability for readers not already familiar with the area.

3. Several mathematical formulations appear simplified to the point of becoming formally incomplete. For example, the EWC objective is missing an explicit summation over model parameters, and some notation in the continual category discovery formulation is inconsistent with the cumulative-class setting described in the text.

4. A few minor typographical and notation inconsistencies are present (e.g., "ADAM" instead of "Adam", inconsistent capitalization of method names and abbreviations, and a small number of grammatical issues).

5. Some cross-references appear incorrect or incomplete (e.g., references to figures where a table is intended), suggesting that a final proofreading pass would be beneficial.

These observations, however, concern presentation rather than substance. Overall, the content of the thesis is of very high scientific quality, with the author having contributed substantially as first author on all three publications. The work addresses relevant and difficult problems, proposes technically sound solutions, and demonstrates impact through publication at the highest level of the field. In my opinion, the thesis clearly satisfies the requirements for the award of the Ph.D. degree. A minor revision of the introductory and background chapters would improve the quality of the manuscript, but it does not affect my positive evaluation.

Simone Scardapane
Associate Professor
Sapienza University of Rome
Rome, 01/06/2026