Balancing Continual Learning and Fine-tuning for Human Activity Recognition

Chi Ian Tang[†*], Lorena Qendro[†], Dimitris Spathis[†], Fahim Kawsar[†], Akhil Mathur[†], Cecilia Mascolo[*]

[†]Nokia Bell Labs, UK [*]University of Cambridge, UK

Summary

- Explored two state-of-the-art **Continual Self-supervised Learning (CSSL)** methods, CaSSLe and Kaizen, for **human activity recognition**, one of the fundamental tasks in human-centric computing.
- Our evaluation indicates the **unified training scheme** proposed in Kaizen can perform better under realistic data assumptions.
- A **progressive importance coefficient** that adaptively adjusts the importance of knowledge retention and classification learning can better balance different learning objectives, reaching **higher performance compared to a fixed loss function**.

1. Motivation

- A key challenge in user modelling is the presence of shifts in human behaviours, where user behaviours can change over time.
- Even though many approaches have been proposed to mitigate catastrophic forgetting, most assume abundant labelled data for every new distribution, which is unrealistic for mobile sensing.

2. Continual Self-supervised Learning vs Continual Fine-tuning

We adapted the state-of-the-art Continual Self-supervised Learning schemes, CaSSLe, which adopts a train-and-then-fine-tune paradigm, and Kaizen, which proposes a unified continual fine-tuning paradigm to human activity recognition.

Train-and-then-fine-tune (CaSSLe)

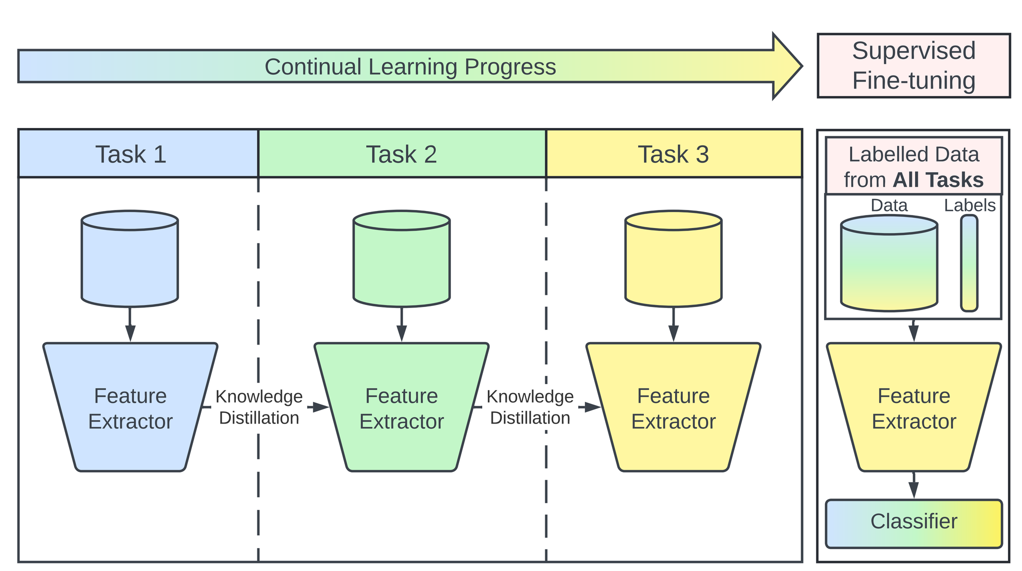

Continual Fine-tuning (Kaizen)

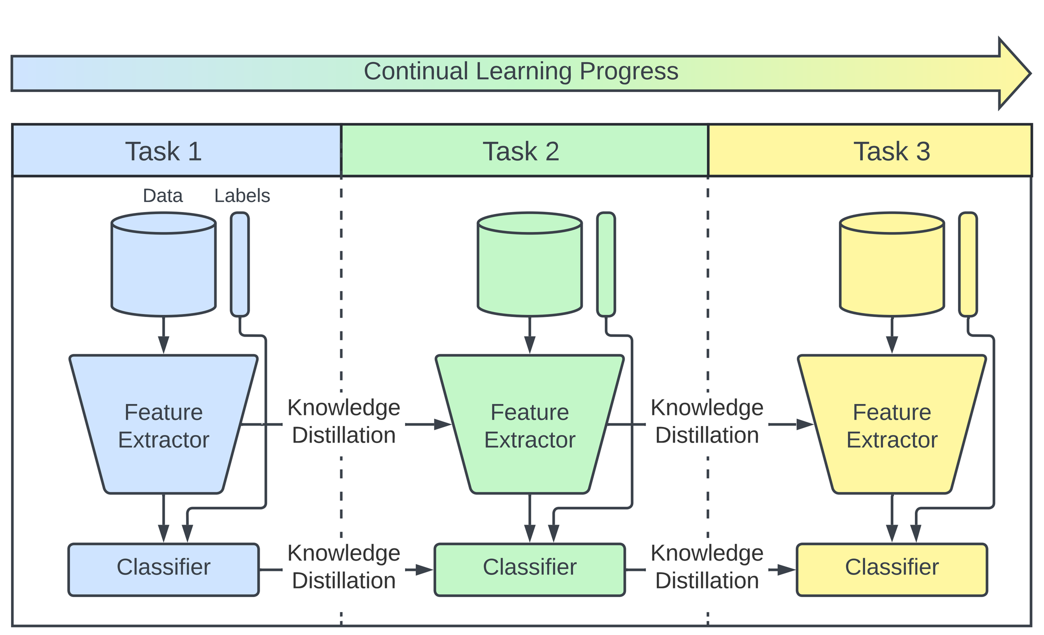

3. Overall Performance Comparison

We compare the overall performance of **Kaizen**, **CaSSLe** and the ***No distill*** setup conducted on the WISDM2019 dataset.
Data availability is standardized across methods, with each having access to current task data and 1% replay data from previous tasks.

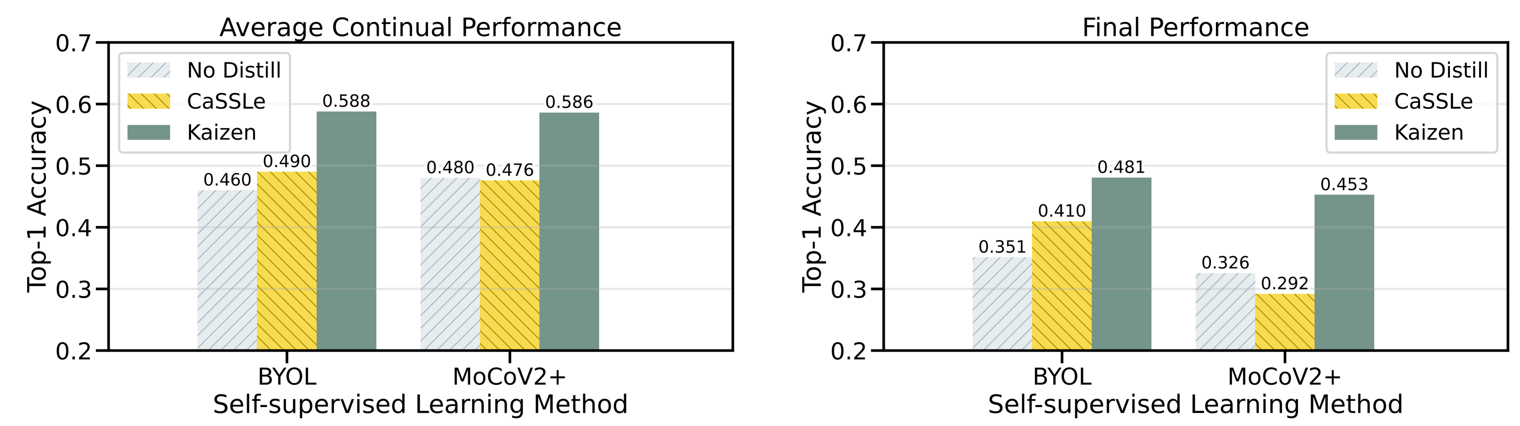

Kaizen consistently surpasses the two baseline models. Specifically, Kaizen achieved the highest continual accuracy and final accuracy, exceeding CaSSLe by **9.8%** and **7.1%**.

4. Balancing Continual Learning and Fine-tuning

We hypothesise that the relative importance of the knowledge distillation task compared to learning from new data in classification learning can have a direct impact on the performance of the classifier across time.
Therefore, we introduce an importance coefficient λ_C to the loss function which allows us to change the weighting of the learning objectives:

$$\mathcal{L}_{\text{Kaizen(adaptive)}} = (\mathcal{L}_{\text{FE}}^{\text{CT}} + \mathcal{L}_{\text{FE}}^{\text{KD}}) + (\mathcal{L}_{\text{C}}^{\text{CT}} + \lambda_{\text{C}}\mathcal{L}_{\text{C}}^{\text{KD}})$$

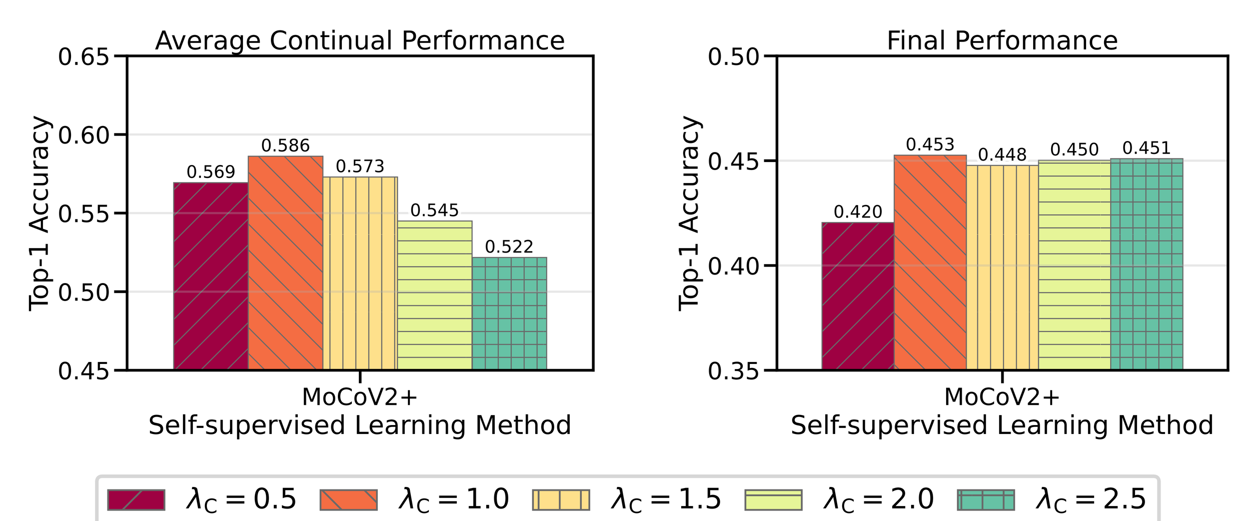

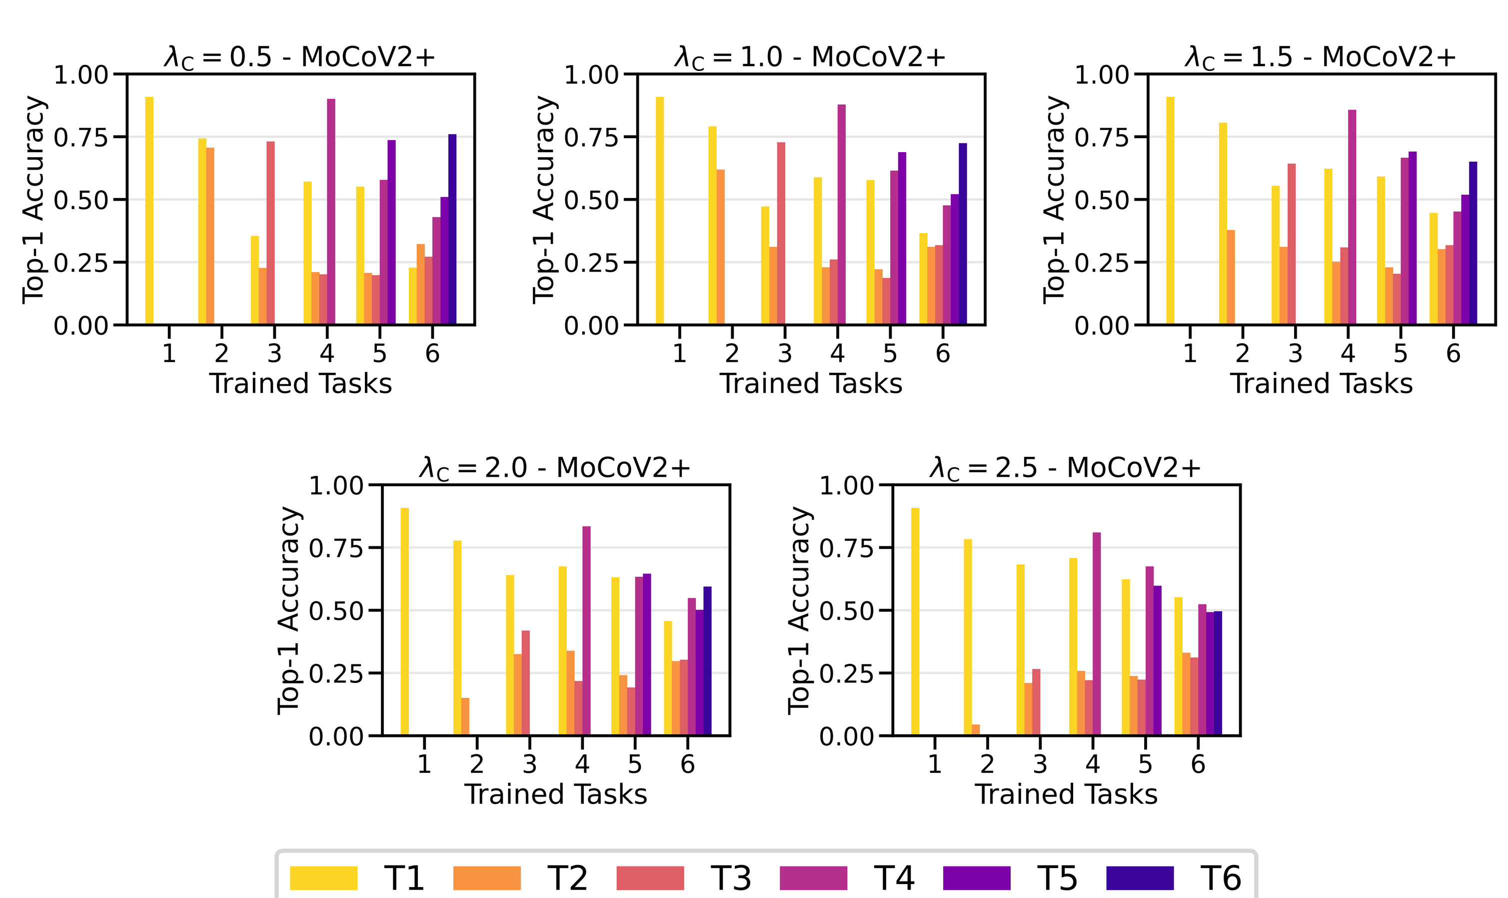

5. Progressive Importance Coefficient

In addition, we hypothesise that a *progressive* importance coefficient can allow the model to balance knowledge retention and new task learning better than a constant value. We denote the setting where λ_C is set to a initially, and increased by b after each task as $\lambda_C = a \oplus b$.

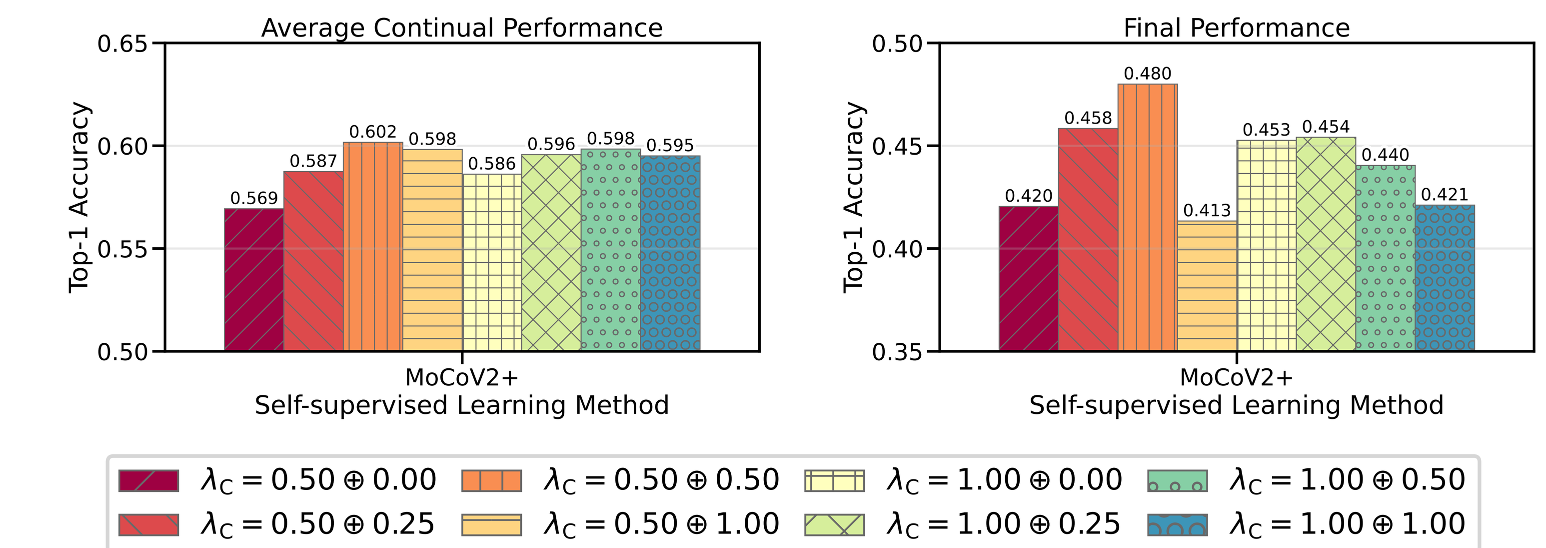

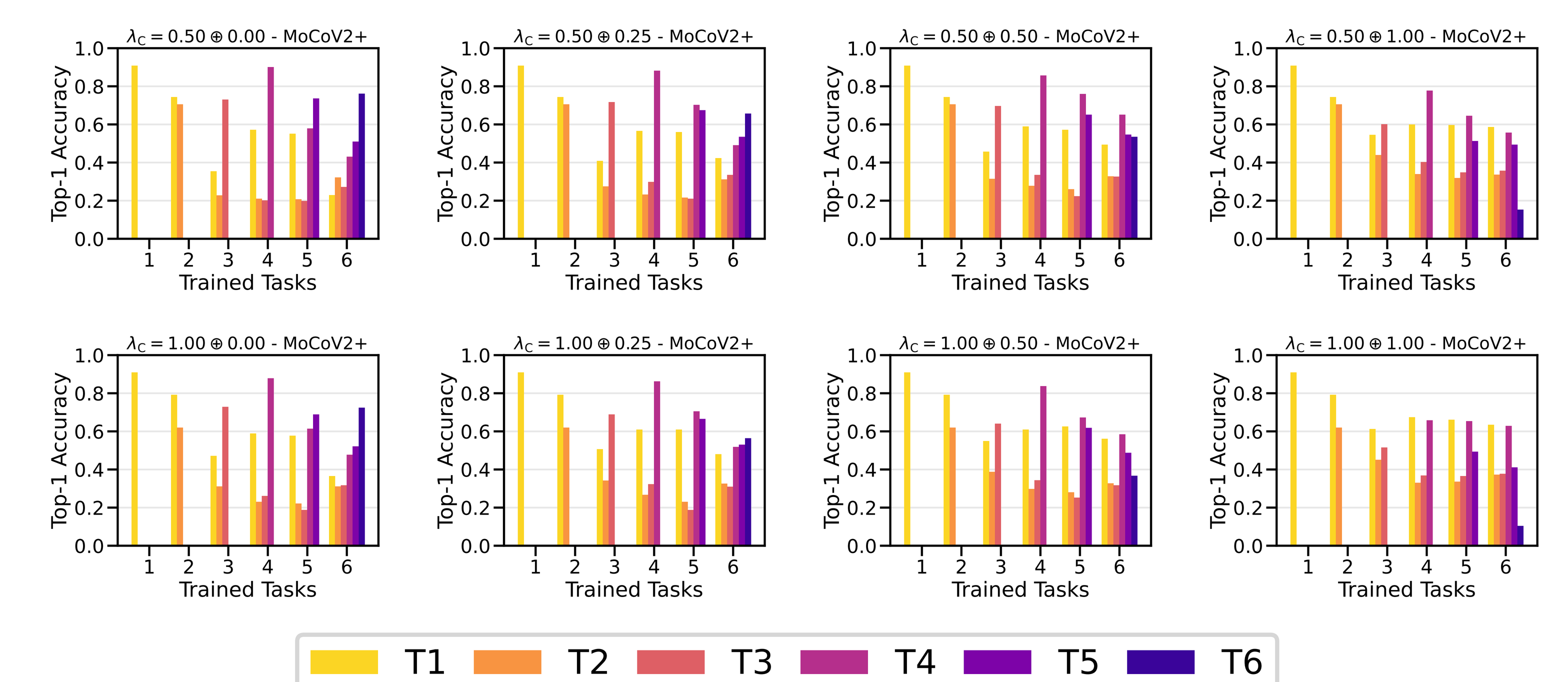

- With $\lambda_C = 0.50 \oplus 0.50$, which corresponds to scaling the importance of knowledge retention proportional to the number of tasks learned, the model achieved the highest performance.
- The use of a progressive factor allows the model to shift the focus from new task learning to knowledge retention over time.

6. Conclusion

- We adapted two state-of-the-art CSSL frameworks, CaSSLe and Kaizen, from visual representation learning to human activity recognition.
- A unified training scheme handling both representation learning and classification learning performs better under realistic data assumptions.
- A progressive importance coefficient that adaptively adjusts the importance of knowledge retention and classification learning can offer better trade-off between different learning objectives.

More results are available in our paper.

arxiv.org/abs/2401.02255

AAAI 2024 Human-Centric Representation Learning Workshop