

Continual Learning

taming non-stationary data streams

Andrea Cossu – University of Pisa

[HTTPS://WWW.ANDREACOSSU.COM/](https://www.andreacossu.com/)

ANDREA.COSSU@DI.UNIPI.IT

Pervasive AI Laboratory



pai.di.unipi.it

Pervasive Computing + Artificial Intelligence

Designing AI for pervasive deployment
Pervasive computing for AI

- A joint initiative by DI.UNIPI & ISTI-CNR founded in 2019
- ~50 members, coordination of 4 EU Projects
- >15M Euro secured grants

ContinualAI

ACTIVITIES

 [2023 Seminar Series](#) [News & Events](#) [Research](#) [Lab](#)

ABOUT US

 [Supporters](#) [People & Governance](#)[Join Us!](#) [Contacts](#)


EXTERNAL LINKS

[Avalanche](#) [Wiki](#) [Continual Learning Course](#) [Open World Lifelong Learning Course](#) [Forum](#) [Newsletter](#) [Mailing-list](#)  Powered By GitBook

ContinualAI

A Non-profit Research Organization and Open Community on Continual Learning for AI



Humans have the extraordinary ability to *learn continually* from experience. Not only can we apply previously learned knowledge and skills to novel situations, but we can also use these as the foundation for later learning. One of the grand goals of AI is to build artificial "**continual learning**" agents that construct a sophisticated understanding of the world from their own experience through the incremental development of increasingly complex knowledge and skills.

 *ContinualAI is an official non-profit research organization and the largest open community on Continual Learning for AI. Our core mission is to fuel continual learning research by connecting researchers in the field and offering a platform to share, discuss, and produce original research on a topic we consider fundamental for the future of AI.*

We plan to fulfill our mission by:

- Organizing and sponsoring events and editorial publications on Continual Learning across the world such as workshops at major AI conferences, industrial panels, webinars, and special issues.
- Improving our aggregation and communication platforms: Slack, social media, mailing-list, newsletter, forum, etc.

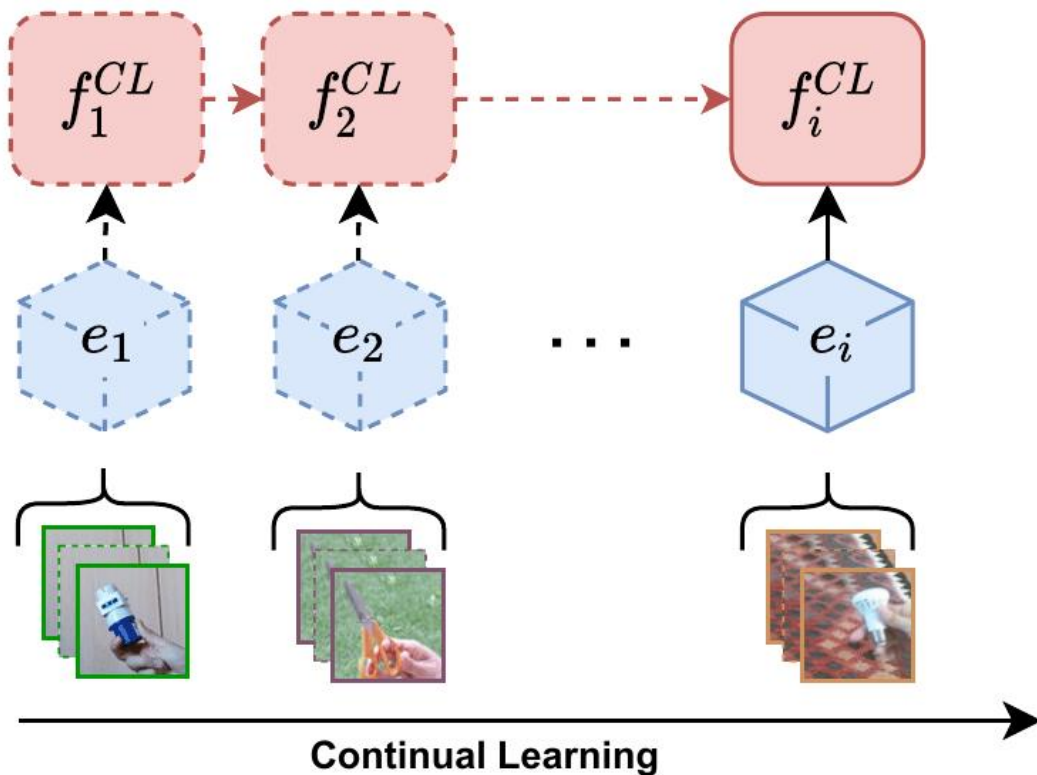
“Offline” Machine Learning



Why not?

- Gather large amount of data *in advance*
- Small drifts → performance drop
 - faults/anomalies, new hardware
 - requires *isolated* models
- Retraining from scratch (old + new) whenever data changes
 - ChatGPT-style → dataset = ?
 - Time + Compute → **Sustainability**
- Decide when to train and when to deploy (or do both in parallel / cycles)
 - A single model to rule them all

Continual Learning



- The stream cannot be stored/accessed at once
- *Drift* between experiences
- Knowledge retention: mitigation of forgetting
- Forward transfer
- Sub-linear (or better, constant!) memory and compute

Popular Continual Learning methods

- Architectural approaches
 - Expand dynamically, mixture of experts
 - What about memory / compute?
- Regularization approaches
 - Adaptive change of the model based on some constraint
- **Replay approaches**
 - train on new + *subset* of past data
 - *Memory buffer*
- Combine the best of all worlds!
 - ARR, Progress&Compress

Bonus: Streaming Continual Learning


- Fast adaptation
- Drift detection
 - e.g., ADWIN
 - passive approaches
- Slow (CL) + fast (SL) learning



An example application




Should *you* use it?



- Are you using multiple data sources over time?
 - Multiple detectors with different characteristics
 - Federated continual learning (for the future...)
- Do you need to react to alerts / anomalies?
 - Drift detection
 - Reduce false positives by *adapting*
- Real-time / fast adaptation?
 - Neuromorphic hardware (randomized networks, spiking networks)
 - Learning on-device is required

<https://bit.ly/ai4eic>



Avalanche overview

