

# Preface to the Proceedings of the 1st Streaming Continual Learning Bridge at AAI26

Heitor Murilo Gomes<sup>1</sup>, Andrea Cossu<sup>2</sup>, Federico Giannini<sup>3</sup>, Anton Lee<sup>1</sup>,  
Nuwan Gunasekara<sup>4</sup> and Emanuele Della Valle<sup>3</sup>

<sup>1</sup>Victoria University of Wellington

<sup>2</sup>University of Pisa

<sup>3</sup>Politecnico di Milano

<sup>4</sup>Halmstad University

## Abstract

Continual Learning and Streaming Machine Learning address complementary aspects of learning under non-stationary data, yet differ in objectives and assumptions, limiting interaction between the two communities. Streaming Continual Learning has recently emerged as a unifying paradigm that combines rapid adaptation to evolving data streams with selective knowledge retention. The Streaming Continual Learning Bridge brought together researchers working at the intersection of these areas to discuss open challenges in drift adaptation, forgetting, plasticity, and temporal dependencies. The contributions in this volume highlight recent methodological advances, practical tools, and application-driven perspectives, outlining future directions for learning systems operating under continuous distributional change.

## 1. Goal of the bridge

Continual Learning (CL) [1] and Streaming Machine Learning (SML) [2] both address the challenge of adapting learning agents to environments that evolve over time. Despite their shared goals, these research areas have had limited interaction, largely due to differences in their approaches, assumptions, and evaluation metrics.

CL primarily emphasizes long-term knowledge retention and the mitigation of forgetting, often without strict real-time constraints. It typically addresses situations where the new concept simply extends the previously observed learning task with a new input distribution, without invalidating the past task. In this case, retaining the past knowledge is crucial since the task associated with the previously observed input distribution is assumed to remain unchanged. Conversely, SML prioritizes rapid, efficient adaptation without making any assumptions about the type of change. This change may alter the mapping between the input and the desired target, invalidating a portion of previously acquired knowledge. Thus, SML focuses only on the current concept and ignores the problem of forgetting, prioritizing rapid, efficient adaptation to high-frequency streams.

Moreover, in many real-world scenarios (e.g., Internet of Things, sensor data), data points are collected over time with a specific order and timing, and the solution should model temporal dependence between data points. This problem is often underexplored by SML and CL and is addressed by Time Series Analysis (TSA) [3]. However, TSA is less commonly integrated to data streams.

Recently, **Streaming Continual Learning (SCL)** [4, 5, 6] is emerging as a unified learning paradigm for scenarios where neither Continual Learning nor Streaming Machine Learning is sufficient on its own. Real-world environments typically involve an interplay of *input drifts*, where new information appears in previously unseen feature subspaces, and *contradictory drifts*, where changes in the data invalidate parts of the learned decision boundary. SCL aims to reconcile CL and SML objectives by promoting five key properties: rapid adaptation to both types of drift; autonomous drift detection; efficient learning from single or few samples; the development of deep hierarchical representations; and the selective

---

1st Streaming Continual Learning Bridge at AAI26, January 21, 2026, Singapore.

✉ heitor.gomes@vuw.ac.nz (H. M. Gomes); andrea.cossu@unipi.it (A. Cossu); federico.giannini@polimi.it (F. Giannini); anton.lee@ecs.vuw.ac.nz (A. Lee); nuwan.gunasekara@hh.se (N. Gunasekara); emanuele.dellavalle@polimi.it (E. D. Valle)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

retention of relevant knowledge. From this perspective, avoiding forgetting is reframed as a dynamic resource-management problem, in which the model must determine which information remains useful and which becomes obsolete as the environment evolves. Consequently, learning systems must both *archive* a global representation of the data stream and *selectively reuse* past knowledge.

This bridge welcomed researchers at any level working on learning protocols and models for non-stationary environments where CL and SML ideas intersect. This also includes areas like online learning in non-stationary environments [7], transfer learning, domain/test-time adaptation, and TSA.

Participants in the bridge shared their ideas with researchers from different areas through dedicated sessions in the bridge program. In addition, participants learned from tutorials and invited talks about the key ideas behind CL and SML research, and how they differ from each other. The ultimate goal was to spark new collaborations in which experts in CL can contribute to long-standing challenges in SML, and vice versa.

## 2. Open research questions

The bridge focused on the interplay between rapid adaptation and knowledge consolidation (i.e., mitigating forgetting), two of the main objectives of CL and SML models. Recent studies provide preliminary evidence about how SML models address knowledge retention and how CL models can rapidly adapt in the presence of sudden drifts and high-frequency data streams [8, 9, 10].

The bridge encouraged participants to reason about key open questions, like:

- Can we design learning models that quickly adapt to new information (in the spirit of SML) without forgetting previous knowledge (in the spirit of CL)?
- What is the meaning of avoiding forgetting in the case of real drifts (i.e., the new classification problem changes the decision boundary in a portion of the previously observed feature space)?
- Is the loss of plasticity [11] commonly encountered in CL also present in SML? If not, how can we leverage insights from SML to mitigate this adverse phenomenon?
- Can we separate the concerns of continual knowledge representation and rapid task adaptation by combining CL and SML techniques?
- Can we leverage the temporal dependencies [12] usually present in a data stream to improve the learning experience?

## 3. Bridge Activities and Program

The bridge was held on January 21, 2026, in Singapore, collocated with the 40th AAAI Conference on Artificial Intelligence. The program combined theoretical discussions with practical sessions to foster shared understanding between the Continual Learning and Streaming Machine Learning communities. In particular, the bridge featured:

- **Technical tutorials** on widely adopted software libraries, providing participants with a common practical grounding for research in Streaming Continual Learning. Tutorials covered *Avalanche* [13] and *CapyMOA* [14], representative frameworks from the Continual Learning and Streaming Machine Learning communities, respectively.
- **Paper presentations and discussion sessions**, organized around poster spotlights followed by a dedicated poster session, which enabled in-depth technical exchanges and cross-community interaction. The bridge concluded with a plenary discussion synthesizing insights across contributions and identifying key open challenges and future research directions.
- **An invited talk** by Prof. Albert Bifet, co-leader of the MOA framework, who discussed drift-aware algorithms and scalable open-source tools for adaptive learning in non-stationary environments, with a focus on real-world applications in environmental monitoring.

## 4. Overview of the Papers

The bridge welcomed contributions on learning protocols and models for non-stationary environments at the intersection of Continual Learning, Streaming Machine Learning, and Time Series Analysis. Submissions were accepted in two tracks:

1. **Non-archival track**, including previously published works, preliminary results, and position papers. The bridge received five and two were presented.
2. **Archival track**, including short and full papers presenting original contributions, published in the post-proceedings on CEUR-WS.org. The bridge received 13 submissions and accepted seven.

The papers included in this volume correspond to revised post-proceedings versions of the accepted archival track contributions.

### 4.1. Foundation Models and Agents in Streaming Environments

A significant portion of the discussion focused on bridging the gap between static Foundation Models and dynamic data streams. Innovative ways to adapt Large Language Models (LLMs) and Large Tabular Models (LTMs) were proposed without incurring catastrophic forgetting. Proposals included bio-inspired frameworks like *Learn-Master-Teach Tuning* (LMT<sup>2</sup>), which simulates a human-like student-teacher lifecycle to resolve the stability-plasticity dilemma, and autonomous agents like *SOLAR*, which leverage parameter-level meta-learning to self-improve and reason in evolving environments. Vision papers further expanded this scope, exploring the use of LTMs as a bridge for SCL through in-context learning and summarization, and outlining the challenges of applying Foundation Models to Earth Observation data streams, where managing sensor degradation and environmental drift is critical.

### 4.2. Time Series Forecasting and Real-World Applications

The intersection of SCL and TSA proved to be a fertile ground for research. In the non-archival track, authors strengthened the theoretical connections between Online CL and TSA, proposing methods like *Natural Score-driven Replay* (NatSR) and unified frameworks that model temporal drift as a sequential domain shift process. On the application side, the bridge showcased works addressing concrete industrial challenges. These included energy consumption forecasting for heavy-duty electric vehicles using online incremental learning, and hybrid anomaly detection systems for industrial elevators. The latter notably introduced a "dual-learner" approach, combining fast online tree-based models with slower, pre-trained time-series foundation models to balance responsiveness and robustness.

### 4.3. Methodologies for Drift Adaptation and Evaluation

Finally, the bridge highlighted fundamental methodological contributions regarding how we evaluate performance in non-stationary environments. The crucial issue of evaluation was addressed by proposing new metrics that distinguish between a model's failure to adapt and the intrinsic difficulty of the data, offering a more nuanced view of robustness to temporal distribution shifts.

## Acknowledgments

We would like to thank our invited speaker, Albert Bifet, for his valuable contributions to the bridge program. We also express our gratitude to the Program Committee members for their time and expertise in reviewing the submissions, and to the AAAI organization for hosting this event.

## References

- [1] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, S. Wermter, Continual lifelong learning with neural networks: A review, *Neural Networks* 113 (2019) 54–71. doi:10.1016/j.neunet.2019.01.012.
- [2] H. M. Gomes, J. Read, A. Bifet, J. P. Barddal, J. a. Gama, Machine learning for streaming data: state of the art, challenges, and opportunities, *SIGKDD Explor. Newsl.* 21 (2019) 6–22. URL: <https://doi.org/10.1145/3373464.3373470>. doi:10.1145/3373464.3373470.
- [3] G. E. Box, G. M. Jenkins, G. C. Reinsel, G. M. Ljung, *Time series analysis: forecasting and control*, John Wiley & Sons, 2015.
- [4] F. Giannini, G. Ziffer, A. Cossu, V. Lomonaco, Streaming Continual Learning for Unified Adaptive Intelligence in Dynamic Environments, *IEEE Intelligent Systems* 39 (2024) 81–85. doi:10.1109/MIS.2024.3479469.
- [5] N. Gunasekara, B. Pfahringer, H. M. Gomes, A. Bifet, Survey on Online Streaming Continual Learning, in: *Thirty-Second International Joint Conference on Artificial Intelligence*, volume 6, 2023, pp. 6628–6637. doi:10.24963/ijcai.2023/743.
- [6] A. Cossu, F. Giannini, G. Ziffer, A. Bernardo, A. Gepperth, E. Della Valle, B. Hammer, D. Bacciu, A practical guide to streaming continual learning, *Neurocomputing* 674 (2026) 132951. URL: <https://www.sciencedirect.com/science/article/pii/S0925231226003486>. doi:<https://doi.org/10.1016/j.neucom.2026.132951>.
- [7] G. Ditzler, M. Roveri, C. Alippi, R. Polikar, Learning in Nonstationary Environments: A Survey, *IEEE Computational Intelligence Magazine* 10 (2015) 12–25. doi:10.1109/MCI.2015.2471196.
- [8] L. Chambers, M. M. Gaber, H. Ghomeshi, AdaDeepStream: Streaming adaptation to concept evolution in deep neural networks, *Applied Intelligence* (2023). doi:10.1007/s10489-023-04812-0.
- [9] Y. Ghunaim, A. Bibi, K. Alhamoud, M. Alfarra, H. A. Al Kader Hammoud, A. Prabhu, P. H. S. Torr, B. Ghanem, Real-Time Evaluation in Online Continual Learning: A New Hope, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 11888–11897.
- [10] H. Hu, O. Sener, F. Sha, V. Koltun, Drinking From a Firehose: Continual Learning With Web-Scale Natural Language, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45 (2023) 5684–5696. doi:10.1109/TPAMI.2022.3218265.
- [11] S. Dohare, J. F. Hernandez-Garcia, Q. Lan, P. Rahman, A. R. Mahmood, R. S. Sutton, Loss of plasticity in deep continual learning, *Nature* 632 (2024) 768–774. doi:10.1038/s41586-024-07711-7.
- [12] G. Ziffer, F. Giannini, E. Della Valle, Tenet: Benchmarking Data Stream Classifiers in Presence of Temporal Dependence, in: *2024 IEEE International Conference on Big Data (BigData)*, 2024, pp. 1187–1196. doi:10.1109/BigData62323.2024.10825670.
- [13] A. Carta, L. Pellegrini, A. Cossu, H. Hemati, V. Lomonaco, Avalanche: A pytorch library for deep continual learning, *J. Mach. Learn. Res.* 24 (2023) 363:1–363:6.
- [14] Y. Sun, H. M. Gomes, A. Lee, N. Gunasekara, G. W. Cassales, J. Liu, M. Heyden, V. Cerqueira, M. Bahri, Y. S. Koh, B. Pfahringer, A. Bifet, Machine learning for data streams with capymoa, in: *ECML/PKDD (10)*, volume 16022 of *Lecture Notes in Computer Science*, Springer, 2025, pp. 438–443.