Open-world continual learning via knowledge transfer

Li, Y.

Open-world Continual Learning via Knowledge Transfer

Proefschrift

ter verkrijging van
de graad van doctor aan de Universiteit Leiden,
op gezag van waarnemend rector magnificus,
volgens besluit van het college voor promoties
te verdedigen op dinsdag 27 januari 2026
klokke 11:30 uur

door

Yujie Li

geboren te Ganzhou, China

in 1996

Promotores:
Prof.dr. M.M. Bonsangue
Prof.dr. X. Yang (Southwestern University of Finance and Economics)

Promotiecommissie:
Prof.dr. M.S. Lew
Prof.dr.ir. F.J. Verbeek
Dr. D.M. Pelt
Prof.dr. T. Zhong (University of Electronic Science and Technology of China)
Dr. L. Huang (Southwestern University of Finance and Economics)

This thesis was performed at the Theory Group, Leiden Institute of Advanced Computer Science, Leiden University, The Netherlands.

Yujie Li was partially supported through by China Scholarship Council (CSC) to participate in the PhD program of Leiden University with grant number 202306980040.

ISBN 978-94-6522-966-9

To my grandmother, Peng Chenglan,
who taught me the meaning of love, forever in my heart.

Abstract

This thesis presents a systematic investigation into Open-world Continual Learning (OWCL), a newly emerging paradigm that fundamentally redefines how intelligent systems learn in non-stationary and uncertain environments. Rather than being a mere combination of continual learning and open-set recognition, OWCL establishes a new research problem that addresses the simultaneous needs for continuous adaptation and open-world awareness. It envisions learning agents that not only retain and transfer knowledge across evolving tasks, but also autonomously identify, characterize, and incorporate novel information beyond previously known categories. In doing so, OWCL departs from the traditional closed-world assumption underlying most deep learning frameworks and moves toward a truly dynamic, self-evolving learning paradigm. The overarching goal of this thesis is to construct the theoretical, methodological, and practical foundations of OWCL, thereby advancing the development of lifelong, adaptive, and trustworthy artificial intelligence.

The thesis begins by formalizing OWCL through a rigorous mathematical definition and identifying four representative real-world scenarios that capture its essential characteristics. This formalization provides a principled framework for understanding the inherent openness and non-stationarity of real-world data, distinguishing OWCL from traditional continual learning paradigms. Through extensive empirical analyses, the study reveals that OWCL is defined not merely by sequential learning but by its continuous exposure to the unknown—posing challenges related to open-set boundary expansion, knowledge retention, and adaptive generalization. These findings lay the groundwork for addressing the stability–plasticity dilemma within an open and evolving context.

Building upon this conceptual foundation, the thesis advances a knowledge transfer-centric perspective as a unifying principle for solving the challenges of OWCL. Within this framework, three novel models are introduced to address complementary aspects of continual adaptation. The first, Pro-KT, develops a prompt-based mechanism for representing and transferring both task-specific and task-agnostic knowledge. By dynamically encoding and reusing knowledge across tasks, Pro-KT enhances the system's ability to maintain adaptability while preserving previously learned information, achieving a balanced trade-off between stability and plasticity. Extending this direction, PEAK explores the few-shot open-world continual learning

setting, where labeled data are scarce and task boundaries are ambiguous. Through a joint architecture designed to mitigate representation collapse and boundary uncertainty, PEAK demonstrates strong resilience in learning under limited supervision while maintaining effective open-set detection. (Chapter 4 & Chapter 5)

To provide a unified theoretical and architectural foundation for OWCL, the thesis introduces HoliTrans, a holistic framework that integrates open-set recognition with continual learning through a shared decision space. By analyzing the intrinsic coupling between open-set dynamics and task progression, HoliTrans derives optimization principles that govern effective learning in open environments and embeds them into a transferable model design. This framework not only achieves state-of-the-art empirical performance but also offers provable guarantees on decision boundary stability and knowledge preservation, bridging the gap between theoretical understanding and algorithmic implementation. (Chapter 6)

Beyond theoretical and algorithmic contributions, this thesis demonstrates the practical value of OWCL through an applied study in cross-regional fraud detection. Real-world financial systems often face the dual challenges of distributional shifts and evolving fraudulent behavior across regions. To address these, the CCL framework is proposed, formulating fraud detection as a continual graph learning problem. CCL combines prototype-based knowledge replay, parameter smoothing, and heterogeneous graph modeling to adapt to region-specific fraud patterns while retaining critical knowledge from previous distributions. Experimental results show that CCL not only mitigates catastrophic forgetting but also enhances cross-regional adaptability, thus validating the real-world relevance and scalability of OWCL principles. (Chapter 7)

Collectively, the four research questions explored in this thesis chart a coherent trajectory from theoretical formulation to holistic methodology and real-world validation. The contributions of Pro-KT, PEAK, HoliTrans, and CCL together form an evolving narrative of how knowledge can be accumulated, transferred, and generalized in open, uncertain, and continually changing environments. This work establishes open-world continual learning as a distinct and necessary paradigm for next-generation artificial intelligence—one that is capable of embracing novelty, uncertainty, and continual evolution. The findings of this thesis not only deepen our theoretical understanding of open learning systems but also open new avenues for future research in lifelong, knowledge-centric, and robust AI.

Contents