



Universiteit
Leiden
The Netherlands

Open-world continual learning via knowledge transfer

Li, Y.

Citation

Li, Y. (2026, January 27). *Open-world continual learning via knowledge transfer*. Retrieved from <https://hdl.handle.net/1887/4287955>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4287955>

Note: To cite this publication please use the final published version (if applicable).

Chapter 1

Introduction

1.1 Background

It is not the most intellectual of the species that survives; it is not the strongest that survives; but the species that survives is the one that is able best to adapt and adjust to the changing environment in which it finds itself.

— Charles Darwin ¹

Learning serves as the fundamental mechanism enabling intelligent systems to adapt to dynamic environments. Throughout biological evolution, humans and other organisms have developed exceptional adaptive capabilities, allowing them to continuously acquire, refine, and apply knowledge in response to environmental changes [1, 2, 3, 4, 5]. Inspired by these natural learning processes, researchers aim to enrich artificial intelligence (AI) systems with analogous continual learning abilities. This pursuit has led to the development of continual learning (CL), a paradigm focused on training models to incrementally process sequences of tasks while maintaining performance on previously learned ones (see Figure 1.1). Continual learning enables knowledge acquisition across diverse domains, including new skill learning, skill refinement, environmental adaptation, and contextual understanding [6, 7, 8].

In the literature, continual learning is often synonymous with incremental learning and lifelong learning [4, 8, 9, 10]. While these terms are frequently used interchangeably, they collectively embody the overarching objective of developing artificial intelligence systems capable of dynamic evolution, preserving acquired knowledge while avoiding strict dependence on static data distributions.

Unlike traditional machine learning models that assume static data distributions, continual learning specifically addresses evolving, dynamic distributions. A cen-

¹This quote is widely attributed to Charles Darwin, but it does not appear verbatim in his writings. The phrasing is believed to originate from Professor Leon C. Megginson, who paraphrased Darwin’s ideas. Despite its frequent misattribute, the quote effectively captures the essence of Darwinian evolution and has since been popularized in both scientific and managerial literature.

tral challenge in continual learning is catastrophic forgetting (CF) [11, 12], where learning new information severely degrades a system’s ability to retain and utilize previously acquired knowledge. This phenomenon reveals the fundamental tension between two competing objectives: **learning plasticity** (the capacity to acquire new knowledge rapidly) and **memory stability** (the ability to preserve existing knowledge). Achieving an optimal balance proves challenging, as excessive plasticity compromises stability, and vice versa [2, 12, 13].

Furthermore, effective continual learning systems must demonstrate strong **generalizability** to handle distributional shifts both within individual tasks and across multiple task domains (see Figure 1.1 **b**). While one potential solution involves re-training models on all previous data, this approach presents substantial limitations, including prohibitive computational costs, extensive storage requirements, and potential privacy issues [14].

The core objective of continual learning is consequently to enable efficient model updates that learn new data incrementally while preserving knowledge from past experiences, maintaining computational and storage efficiency, and scaling effectively with continuous changes.




Figure 1.1: A conceptual framework of continual learning. **a.** Continual learning requires adapting to incremental tasks with dynamic data distributions. **b.** A desirable solution should ensure an appropriate trade-off between stability (orange arrow) and plasticity (green arrow), as well as a generalizability to intra-task (blue arrow) and inter-task (red arrow) distribution differences.

Despite significant advancements in continual learning, most current approaches operate under a **closed-world assumption**, limiting systems to recognizing and processing only those object classes explicitly encountered during training [8, 15, 16]. This assumption implicitly excludes the possibilities of unexpected or novel entities emerging during testing or deployment, thus falling short in real-world scenarios, as the environment is inherently open and constantly presents unknown objects.

Human cognition, by contrast, treats novelty as a powerful *intrinsic motivator* for learning. The human brain responds to new stimuli with distinct neural activity and behavioral changes, driving exploration and enabling flexible encoding of unfamiliar information [8]. Over time, however, this novelty response attenuates as the brain rapidly adapts to repeated exposure to the same stimuli [17, 18], demonstrating an adaptive balance between exploration and habituation.

To operate effectively in an **open-world environment**, artificial intelligence systems must simulate these human capabilities. Specifically, an ideal system would autonomously identify and respond to novel entities, incrementally integrate new knowledge while maintaining existing competencies, and adapt dynamically without requiring a periodic offline retraining. Such capabilities would enable artificial intelligence systems to continuously evolve through a sequence of new tasks, allowing real-time adaptation to novel scenarios in practical applications.

In this context, Open-world Continual Learning (**OWCL**) [8, 19] has emerged as both a highly practical yet profoundly challenging machine learning paradigm. This framework requires artificial intelligence systems to continuously adapt to unbounded task sequences in a dynamic open environment [15, 20], where novel elements might appear unpredictably during testing phases [5, 21, 22].

The core objective of the open-world continual learning framework includes operating under the open-world assumption, recognizing unseen/open samples, and incrementally acquiring knowledge from new tasks while preventing catastrophic forgetting [23, 24, 25]. These requirements introduce two critical challenges. Due to the potential occurrence of novelties in continual learning [14], open-world continual learning requires detecting unknown samples to prevent misclassification into known categories. At the same time, open-world continual learning requires the model to retain previously learned knowledge without forgetting during continuous open detection. Consequently, the open-world continual learning paradigm faces a dual complexity. On the one hand, the presence of unknown samples complicates the fundamental tension between *knowledge stability* [1, 14] and *knowledge plasticity* [26, 27]. On the other hand, the continuous expansion of the knowledge space through incremental learning of new tasks complicates open detection in the embedding space.

Although existing continual learning methods have made progress in addressing catastrophic forgetting, they fundamentally operate under three limiting assumptions [28, 29]: (1) task boundaries are always clearly defined, (2) all test-time samples belong to known classes, and (3) data distributions remain relatively stable [24]. These assumptions break down in open real-world environments where models encounter novel categories, shifting contexts, and ambiguous task transitions, precisely the challenges our work addresses.




Figure 1.2: An illustration of the open-world continual learning problem and comparison between continual learning. In open-world continual learning, the test set of each task may include samples from classes unseen during training, deviating from the closed-world assumption of traditional continual learning. As shown in the figure, this setting imposes two key challenges: (1) preserving knowledge from previously learned tasks to mitigate catastrophic forgetting, and (2) performing robust open-set recognition to correctly identify and reject instances from unknown classes. While continual learning models tend to misclassify open samples as one of the known classes (e.g., the green pentagon and orange square), open-world continual learning models are explicitly designed to distinguish between known and unknown classes, thereby enabling more reliable and adaptive decision-making in dynamic environments.

To further illustrate the unique challenges of open-world continual learning, we refer to the scenario depicted in Figure 1.2. Unlike continual learning, where all test-time samples are assumed to belong to previously seen classes, open-world continual learning confronts a more realistic setting where unknown classes may emerge during testing. This introduces the first key challenge: open-set detection, i.e., accurately identifying and rejecting samples from unseen categories rather than misclassifying them into existing known classes. As shown in Figure 1.2, while continual learning models tend to misclassify such unknowns (e.g., the green pentagon and orange square) into existing classes, open-world continual learning explicitly distinguishes them as open samples.

Another challenge lies in the need for stronger generalization and boundary stability. Open-world continual learning models must not only acquire new knowledge incrementally and retain prior learning—addressing the classical stability–plasticity dilemma—but also maintain a consistent open-set recognition capability across tasks. This requires the ability to form and update well-defined task or class boundaries over time, ensuring that the decision boundaries between known and unknown categories remain reliable despite ongoing changes in the data distribution.

Lastly, and most critically, open-world continual learning raises a fundamental question that goes beyond mere classification: how can models exploit knowledge from unknown samples? While traditional continual learning gains knowledge by learn-

ing from labeled data across successive tasks, the presence of unlabeled and unpredictable unknowns in open-world continual learning complicates this process. Effectively mining and leveraging knowledge from open-set instances, without ground truth supervision, becomes essential for sustained learning in dynamic and open environments.

To provide a more detailed description of the diverse scenarios that may arise in open-world continual learning, we construct four representative open-world continual learning settings based on the three classical continual learning paradigms [6]: task-incremental learning, domain-incremental learning, and class-incremental learning. It is important to note, however, that a critical distinction in open-world continual learning lies in the absence of task identity information during testing, as unknown samples can appear unpredictably at any time [30]. This fundamental characteristic renders conventional task-incremental and domain-incremental settings insufficient for fully capturing the complexity of open-world continual learning scenarios.

Therefore, we propose a novel scenario designed for open-world continual learning, called knowledge-incremental learning (**KIL**), characterized by the ability to solve tasks previously seen without knowing which task is being performed (no task identifiers available), non-strictly disjoint training sets across different tasks, and potential distribution shifts among known categories.

Knowledge-incremental learning both subsumes class-incremental learning and accommodates category-specific distribution shifts, thereby providing a more realistic and consequently more challenging framework for open-world continual learning. Based on this foundation, we then introduce four distinct scenarios for open-world continual learning of increasing difficulty: class-incremental with non-repetitive open samples (**CINO**), class-incremental learning with repetitive open samples (**CIRO**), knowledge-incremental learning with non-repetitive open samples (**KINO**), knowledge-incremental learning with repetitive open samples (**KIRO**).

Table 1.1: Overview of different OWCL scenarios.

<i>Scenario</i>	CINO	CIRO	KINO	KIRO
Known Classes Repeatedly Appear	✗	✗	✓	✓
Unknown Samples Repeatedly Appear	✗	✓	✗	✓

Having established four representative open-world continual learning scenarios based on classical continual learning settings, we now turn to the practical significance of developing robust open-world continual learning systems. Recent advances in deep representation learning for computer vision and natural language processing have catalyzed growing interest in open-world continual learning. However, many current approaches [25, 31] still treat open-world continual learning as a loose combination of open-set recognition and continual learning, rather than embracing it as a unified and

principled learning paradigm. This fragmented view leads to two core limitations: (1) knowledge transfer is typically confined to labeled data, overlooking the rich semantic information that may be embedded in unknown samples; and (2) the lack of standardized problem formulations and evaluation protocols impedes consistent benchmarking and method selection.

In response, this thesis takes a knowledge transfer perspective to investigate the open-world continual learning problem. Rather than treating unknown-class detection and known-class classification as isolated objectives, we aim to understand and enhance the underlying transfer mechanisms across both labeled and unlabeled samples in dynamic, open environments. Throughout the thesis, we explore multiple model architectures, training strategies, and evaluation protocols to advance the understanding and practical utility of knowledge transfer in open-world continual learning, which contributes not only to methodological innovation but also to enabling practical open-world continual learning systems that exhibit robust adaptability and generalization across a wide range of real-world scenarios.

For instance, in financial security, open-world continual learning models can support real-time detection of emerging fraud patterns, such as AI-generated deepfake scams, while maintaining accuracy on previously encountered attack types, thus reducing reliance on costly periodic retraining. In healthcare, open-world continual learning systems could incrementally incorporate novel disease markers (e.g., viral mutations) without compromising diagnostic performance for known conditions. These examples highlight open-world continual learning’s potential to enhance operational robustness in dynamic environments, while also addressing pressing societal challenges.

Beyond immediate applications, the broader societal implications of open-world continual learning are equally profound. By mimicking human-like continual learning, open-world continual learning systems can reduce the environmental footprint of artificial intelligence by eliminating repeated retraining cycles, lower the barrier for adoption by organizations with limited access to large-scale labeled data, and foster more trustworthy artificial intelligence systems through transparent novelty detection. Nonetheless, these benefits come with ethical considerations, such as ensuring privacy in continuous surveillance settings or maintaining accountability in medical diagnostics—issues we explicitly address through built-in uncertainty quantification mechanisms in our framework.

While the empirical studies in this thesis primarily focus on image-based data, the proposed open-world continual learning framework is designed to be modality-agnostic. Its core principles—adaptive knowledge transfer, novelty awareness, and dynamic representation updating—can naturally extend to other data modalities such as video, text, speech, and multimodal sensor streams. In particular, applying this framework to sequential modalities (e.g., video surveillance or dialogue systems) could reveal new insights into temporal transfer and incremental context modeling

under a realistic open-world environment. Future research may further explore how the same open-world principles can be instantiated across heterogeneous domains, paving the way toward a unified and generalizable paradigm of continual learning beyond visual tasks.

1.2 Research Questions and Contributions

Our approach advances open-world continual learning beyond current limitations through three key innovations: (a) a dynamic knowledge representation that evolves with both known and unknown samples, (b) a unified optimization framework that jointly preserves stability and enables plasticity, and (c) task-agnostic detection mechanisms that operate without explicit task identifiers. This combination theoretically enables infinite learning capacity while maintaining practical deployability. Note that all subsequent investigations of the problems addressed in this thesis are conducted within the context of image tasks, and all experimental validations in subsequent chapters are carried out using image benchmark datasets.

Research Question 1 (RQ 1) *Can we develop an effective open-world continual learning model to accumulate the knowledge gained in the past and to use the knowledge to help divide the test data into knowns and unknowns/opens, especially for enlarging the open-set boundary between knowns and unknowns?*

Existing methods fail to exploit previous unknowns from former tasks to help new tasks, leading to performance degradation on new tasks. Hence, the main motivation of RQ 1 here is to learn knowledge for both knowns and unknowns from previous tasks and later use the knowledge to formulate an open-set boundary between knowns and unknowns for new tasks. This question targets the foundational challenge of open-world continual learning: maintaining previously acquired knowledge while remaining adaptable to new, unseen classes. A robust open-world continual learning system must not only classify known samples accurately but also detect and separate unknowns with high confidence.

The key challenge of **RQ 1** lies in addressing two critical issues: (1) *Knowledge Stability*, where newly acquired knowledge from a task that is significantly different from previous ones might conflict with existing knowledge; and (2) *Knowledge Plasticity*, which requires the model to update its understanding by reclassifying unknown samples as knowns once their ground-truth data becomes available in subsequent tasks. Our primary contribution to addressing this research question is the development of a novel **Prompt-enhanced Knowledge Transfer** approach, referred to as **Pro-KT**. Pro-KT is introduced in Chapter 4 and based on the following publication [30]:

- **Li, Y.**, Yang, X., Wang, H., Wang, X., & Li, T.: Learning to Prompt Knowledge Transfer for Open-World Continual Learning. In: The 38th Annual AAAI Conference on Artificial Intelligence. AAAI '24 (2024) 38(12), 13700-13708.

Our main ideas are *knowledge transfer* and *prompt learning* [32, 33, 34]. Pro-KT delineates an innovation of prompt learning for open-world continual learning, a novel plug-and-play prompt bank for knowledge transfer, and two adaptive threshold selection strategies for determining the open-set boundary. Specifically, to address *knowledge stability*, we create a prompt bank designed to encode knowledge through the use of prompts. These prompts serve as instructions for directing the model in task execution. By flexibly selecting prompts from the proposed prompt bank, Pro-KT can facilitate effective knowledge transfer with both task-generic and task-specific knowledge across diverse tasks. To address *knowledge plasticity*, we design two adaptive threshold-selection strategies for determining the open-set detection boundary. Through these strategies, the open-set decision boundary will be updated according to the newly learned knowledge continually, so as to handle knowledge plasticity.

While RQ1 addresses the stability-plasticity trade-off, its success heavily depends on access to labeled data. In practice, such labels are often scarce. This motivates the investigation in RQ2, which explores how to achieve strong open-world continual learning performance under severe label constraints by leveraging transfer learning and few-shot techniques.

Research Question 2 (RQ 2) *How can we enhance the performance of existing open-world continual learning models under the constraint of limited labeled data by employing an improved knowledge transfer approach?*

Existing open-world continual learning still requires a large amount of labeled data for training, which is often impractical in real-world applications. Given that new categories/entities typically come with limited annotations and are in small quantities, a more realistic scenario is open-world continual learning with scarce labeled data, i.e., few-shot training samples, which addresses our second research question **RQ 2**.

In RQ 2, we investigate the problem of open-world few-shot continual learning (OFCL), challenging in (i) learning unbounded tasks without forgetting previous knowledge and avoiding overfitting, (ii) constructing compact decision boundaries for open detection with limited labeled data, and (iii) transferring knowledge about knowns and unknowns and even update the unknowns to knowns once the labels of open samples are learned. To address RQ 2, we introduce a novel framework, also referred to as *PEAK*, in Chapter 5 based on the following publication [35], that, to the best of our knowledge, is the first work to study the open-world few-shot continual learning problem:

- Li, Y., Wang, X., Yang, X., Bonsangue, M., Zhang, J., Li, T.: Improving Open-World Continual Learning under the Constraints of Scarce Labeled

Data. In: Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining. KDD '25 (2025).

Specifically, we first propose an instance-wise token augmentation (ITA) aimed at acquiring additional ‘knowledge’ to mitigate the inadequate representation caused by scarce labeled data. Moreover, instance-wise token augmentation can facilitate knowledge transfer by matching learnable tokens to each sample embedding. Additionally, given the scarce labeled data where certain exemplar points (hubs) appear among the nearest neighbors of many other points, test samples will be assigned to it regardless of their true label, resulting in low accuracy [36]. To mitigate this, inspired by the embedding representations on a hypersphere [37], we introduce a novel and compact margin-based open boundary (MOB) and an adaptive knowledge space (AKS) consisting of learnable hyperspheres, where each hypersphere is characterized by a class centroid and an associated radius. In particular, enables the formulation of compact decision boundaries between known and unknown samples, thereby enhancing open detection. Simultaneously, the adaptive knowledge space encourages the model to learn incrementally from unknowns and classify previously encountered unknown samples in new tasks, effectively transforming unknowns into knowns over time.

Even with enhanced few-shot learning (RQ2), existing open-world continual learning methods often disregard the latent information embedded in unknown samples. RQ3 challenges this orthogonality assumption by asking whether knowledge from the open world—especially unlabeled or outlier data—can be actively incorporated into future tasks through holistic knowledge transfer.

Research Question 3 (RQ 3) *How can we leverage knowledge from both known and unknown samples to enable knowledge transfer in open-world continual learning, effectively improving open-set detection with incremental classification at both theoretical and architectural levels?*

Unlike traditional learning models that operate in a closed and fixed set of classes, open-world continual learning aims at learning on the job in an open-world assumption with the goal of recognizing unseen/open samples and incrementally acquiring knowledge from new tasks without forgetting [23, 24, 25]. Due to the potential occurrence of novelties in continual learning [14], open-world continual learning models need to accurately detect unknowns to prevent unknown samples from being incorrectly classified into known categories. At the same time, open-world continual learning requires the model to retain previously learned knowledge without forgetting while continually performing open detection. In summary, the open detection for unknowns and classification for knowns are *interdependent* in open-world continual learning: the incremental learning of new tasks makes open detection in an embedding space more challenging with the expanding knowledge space, and vice versa.

However, current approaches [25, 31] still treat open-world continual learning as a simple combination of open-set recognition and continual learning, rather than as an integrated paradigm, making it only effective in knowledge transfer related to known samples, while neglecting the knowledge derived from unknown samples. Therefore, a promising open-world continual learning model must be capable of **knowns-unknowns knowledge transfer**, i.e., effectively transferring knowledge both for known categories and unknown samples. Besides, there is a lack of problem formulation and thorough empirical explorations of potential issues in open-world continual learning, making it difficult to compare the performance of existing methods and making it unclear how to choose one method over another. In Chapter 6, we explore the issues arising in open-world continual learning and adopt an integrative perspective to deal jointly with unknown samples’ detection and known samples’ classification, particularly the knowledge transfer for both unknowns and knowns, based on the following paper [38]:

- **Li, Y.**, Lai, G., Yang, Y., Li, Y., Bonsangue, M., & Li, T.: Exploring Open-World Continual Learning with Knowns-Unknowns Knowledge Transfer. arXiv preprint arXiv:2502.20124.

Specifically, we provide a formal and unified formulation of the open-world continual learning problem, delineating four distinct yet interrelated scenarios that capture the multifaceted nature of open-world dynamics. Our empirical investigations reveal a strong and non-trivial interaction between the tasks of open-set detection and incremental prediction, challenging the prevailing assumption that these components can be independently optimized. This insight calls for a more holistic approach to open-world continual learning, one that integrates both discovery and adaptation within a unified learning process.

From a theoretical standpoint, we rigorously define the decision space of open-world continual learning and uncover the core optimization objectives that govern learning under open-world conditions. Building upon this foundation, we propose a novel framework, **HoliTrans**, which effectively transfers knowledge from both known and previously unseen (unknown) instances. HoliTrans unifies open-set recognition and continual adaptation through a single principled mechanism, thereby enabling robust generalization in the presence of evolving data distributions and task boundaries.

Extensive empirical evaluations across all four open-world continual learning scenarios corroborate our theoretical insights. HoliTrans consistently outperforms state-of-the-art baselines, demonstrating not only superior performance in isolated metrics but also stable adaptation across open and incremental learning tasks. These results affirm the necessity of treating open-world continual learning as an inherently integrated problem and highlight HoliTrans as a promising step toward a unified, theoretically based open-world continual learning framework.

Having established methods for knowledge retention, transfer, and open-sample utilization, RQ4 shifts focus to practice. It examines whether continual learning and open-world continual learning frameworks can go beyond benchmarks and offer tangible value in complex, real-world tasks characterized by data drift, concept emergence, and geographical heterogeneity.

Research Question 4 (RQ 4) *How can we improve the model performance by formulating practical problems with an open-world continual learning framework and developing effective models for challenging real-world applications?*

Building upon the aforementioned works, we further investigate the potential of continual learning and open-world continual learning to address real-world challenges. As emerging and increasingly influential paradigms in machine learning, continual learning and open-world continual learning offer the promise of reshaping how practical systems adapt to dynamic environments. In contrast to conventional static models, these frameworks inherently accommodate both the continuity of data streams and the openness of task boundaries, making them particularly well-suited for applications characterized by evolving knowledge spaces.

For instance, in the domain of dynamic financial transactions, OWCL-enabled systems can incrementally learn from newly observed fraud types (e.g., deepfake-based voice scams) while simultaneously preserving detection capabilities for previously encountered fraud patterns (e.g., credit card cloning or identity theft). Similarly, in smart healthcare, continual learning has been leveraged to enhance pathological diagnosis systems, empowering them to incrementally incorporate rare or emerging disease categories—such as novel variants of infectious diseases—without requiring complete retraining or catastrophic memory loss.

Motivated by these promising applications, we aim to extend the frontier of open-world continual learning by formulating a novel cross-regional fraud detection problem under the continual learning paradigm. Unlike conventional fraud detection approaches that are confined to static datasets within isolated geographical regions, our proposed setting reflects a more realistic scenario wherein financial service providers operate across multiple cities or countries, each exhibiting unique transaction patterns and fraud behaviors.

The inherent challenges in such settings include high retraining costs, data privacy constraints, and the risk of catastrophic forgetting when adapting models to new regional data. In light of these challenges, we propose in Chapter 7 a unified framework that treats cross-regional fraud detection as an incremental learning task. This formulation enables the development of robust, scalable fraud detection systems that can generalize across diverse regions while effectively retaining prior knowledge and adapting to novel threats. This chapter is based on the following publications [39, 40]:

- **Li, Y.**, Yang, Y., Gao, Q., & Yang, X: Cross-Regional Fraud Detection via Continual Learning (Special Program). In: Proceedings of the AAAI Conference on Artificial Intelligence, 37(13), 16260-16261.
- **Li, Y.**, Yang, X., Gao, Q., Wang, H., Zhang, J., & Li, T.: Cross-Regional Fraud Detection via Continual Learning With Knowledge Transfer. In: IEEE Transactions on Knowledge and Data Engineering. TKDE'24 (2024) 38(12), 7865-7877.

1.3 Thesis Outline

The overall organization of this thesis is as follows. In this chapter, we give a brief introduction to the thesis’s background and motivation, the main research questions, and contributions. Chapter 2 provides a literature review for existing methods in continual learning and open-set recognition, and focuses more on open-world continual learning methods. Chapter 3 presents some general definitions, which are used throughout the whole dissertation, problem definitions of open-world continual learning, and different scenarios, respectively.

Subsequently, we present investigations corresponding to the four research questions, detailed respectively in Chapters 4 through 7. Specifically, in Chapter 4, we address **RQ 1** by introducing *Pro-KT*, a prompt-based knowledge transfer method for open-world continual learning. We design a plug-and-play prompt bank to balance knowledge stability and plasticity by enabling flexible transfer of task-specific and task-generic knowledge.

Chapter 5 addresses **RQ 2** and introduces a novel framework, called *PEAK*, which is designed to tackle the challenges of open-world continual learning under few-shot learning scenarios. Within this framework, we propose three core modules that collectively enable robust performance when labeled data is severely limited, offering an exploration of the unique difficulties posed by limited supervision in open-world continual learning.

Chapter 6 addresses **RQ 3** and introduces *HoliTrans*, a unified framework for open-world continual learning. We provide a formal theoretical construction of the open-world continual learning decision space and identify four interrelated scenarios that reflect the complexity of open-world dynamics. Our findings reveal a strong interaction between open-set detection and incremental prediction, challenging the assumption of their separability and motivating a holistic approach that unifies discovery and adaptation within a holistic paradigm.

Chapter 7 addresses **RQ 4** by introducing *CCL*, a continual learning framework for cross-regional fraud detection. Motivated by real-world deployments, we extend open-world continual learning to a more practical setting where financial institutions operate across diverse regions with varying transaction patterns and fraud behaviors.

Unlike traditional methods limited to static, region-specific datasets, CCL enables adaptive learning across regions, capturing both shared knowledge and local nuances in a unified model.

At last, Chapter 8 concludes the contributions of the thesis and discusses possible future work.