Pre-trained Online Contrastive Learning for Insurance Fraud Detection

Rui Zhang[1,2], Dawei Cheng[1,3,2*], Jie Yang[1], Yi Ouyang[4], Xian Wu[4], Yefeng Zheng[4], Changjun Jiang[1,2]

[1]Department of Computer Science and Technology, Tongji University, Shanghai, China
[2]Shanghai Artificial Intelligence Laboratory, Shanghai, China
[3]Key Laboratory of Artificial Intelligence, Ministry of Education, Shanghai, China
[4]Jarvis Research Center, Tencent YouTu Lab, Shenzhen, China
{2050271,dcheng,2153814,cjjiang}@tongji.edu.cn, {yiouyang,kevinxwu,yefengzheng}@tencent.com

Abstract

Medical insurance fraud has always been a crucial challenge in the field of healthcare industry. Existing fraud detection models mostly focus on offline learning scenes. However, fraud patterns are constantly evolving, making it difficult for models trained on past data to detect newly emerging fraud patterns, posing a severe challenge in medical fraud detection. Moreover, current incremental learning models are mostly designed to address catastrophic forgetting, but often exhibit suboptimal performance in fraud detection. To address this challenge, this paper proposes an innovative online learning method for medical insurance fraud detection, named POCL. This method combines contrastive learning pre-training with online updating strategies. In the pre-training stage, we leverage contrastive learning pre-training to learn on historical data, enabling deep feature learning and obtaining rich risk representations. In the online learning stage, we adopt a Temporal Memory Aware Synapses online updating strategy, allowing the model to perform incremental learning and optimization based on continuously emerging new data. This ensures timely adaptation to fraud patterns and reduces forgetting of past knowledge. Our model undergoes extensive experiments and evaluations on real-world insurance fraud datasets. The results demonstrate our model has significant advantages in accuracy compared to the state-of-the-art baseline methods, while also exhibiting lower running time and space consumption. Our sources are released at https://github.com/finint/POCL.

Introduction

Medical insurance fraud poses a severe detriment to society and is drawing increasing attention from the public. In 2017, the United States expended a staggering $3.5 trillion on healthcare (Sisko et al. 2019), with over 20%, or $720 billion (Cubanski, Neuman, and Freed 2019), dedicated to medical insurance. Yet, amidst these vast expenditures, fraudulent organizations and individuals have found opportunities for exploitation. Studies indicated that an alarming 3-10% of medical insurance funds (Morris 2009), equating to $21-70 billion, were squandered due to deceitful activities. Such fraudulent actions not only heighten the operational costs of the healthcare system, but also burden consumers indirectly (NHCAA 2021). The ramifications of these acts echo throughout society, emphasizing the imperative nature of combating insurance fraud.

Moreover, medical insurance fraud tactics are constantly evolving (Thornton et al. 2013). Fraudsters adapt, devising increasingly clandestine tactics to reduce their exposure to detection (Timofeyev and Jakovljevic 2022). Such evolving stratagems exacerbate the challenges in fraud detection. Contemporary static health insurance fraud detection systems are not fit for this situation (Thornton et al. 2013), overlooking many fraudulent activities that cause massive losses. This underscores the imperative to devise a medical insurance verification system characterized by rapid adaptability and online learning capabilities.

Fraud detection has a long history, tracing its origins back to research in the 1980s (McDowell 1987). Traditional methodologies often pivoted around rule-based approaches as depicted in (Dua and Bais 2014), or embraced the realm of machine learning, evidenced by the work of (Fiore et al. 2019). These approaches garnered significant attention due to their efficacy in fraud detection. Parallelly, as deep learning method gained traction in the broader scientific community, they began to be applied in the fraud detection domain, ushering in a novel era of sophisticated detection mechanisms, such as (Roy et al. 2018). Pioneering works such as (Dou et al. 2020; Cheng et al. 2023; Ma et al. 2023; Gao et al. 2023) highlight the capability of graph neural networks (GNNs) to learn complex relationships and patterns, thereby signifying the burgeoning potential of this approach in elevating the state of fraud detection. However, the models delineated above predominantly cater to the offline learning paradigm. Transposing these methods directly to online learning contexts frequently results in suboptimal outcomes. Retraining a complete model to circumvent these limitations (Lebichot et al. 2020) often introduces significant challenges, given the substantial computational and temporal resources required (Wu, Dobriban, and Davidson 2020).

Fortunately, in recent years, various of fraud detection models employing incremental or online learning have emerged. Notably, in specific dynamic settings, these models have demonstrated commendable efficacy. For instance, (Sadreddin and Sadaoui 2022) leverages an innovative adaptive learning approach, melding transfer learning with incremental learning. (Anowar and Sadaoui 2021) introduces

*Corresponding Author is Dawei Cheng.

a block-based incremental learning framework tailored to combat auction fraud, while (Bayram, Köroğlu, and Gönen 2020) employs gradient boosted trees, addressing the dynamic nuances of credit card fraud. Nevertheless, these methods still have limitations. Firstly, current techniques often overlook the intricate structural and temporal nuances inherent to fraud patterns. Secondly, given the stringent data storage restrictions posed by financial contexts, historical data becomes infeasible.

To address these challenges, in this study, we propose **P**re-trained **O**nline **C**ontrastive **L**earning model (POCL), an innovative online learning graph neural network tailored for medical insurance fraud detection. Firstly, we divide the historical dataset into positive and negative medical node graphs. We pre-train an upstream model by contrasting the difference between these two types of graphs. As illuminated in (Le-Khac, Healy, and Smeaton 2020), contrastive learning serves as a powerful tool in discovering complex patterns, subsequently improving the precision of fraud detection. Simultaneously, our empirical observations suggest that the model proficiently projects evolving fraud characteristics into congruent spaces, which can reduce the difficulty of updates and decrease the likelihood of forgetting. Next, we introduce a downstream fraud detection network and combine it with the pre-trained model to create an end-to-end fraud detection model. In the online learning scenario, we use a novel Temporal Memory Aware Synapses (MAS) (Aljundi et al. 2018) method to update the model. This method incorporates temporal features and calculates the importance weights of parameters for each update. By using the momentum technique to integrate the historical importance weights, this approach determines the extent of forgetting and retaining parameters during updates without storing any historical data. In summary, our work has the following contributions:

- To the best of our knowledge, this is the first work to introduce an online learning model within the domain of insurance fraud detection. This model adeptly amalgamates structural features, continually adapting to the evolving paradigms of fraud patterns.
- We propose a novel online learning GNN model based on contrastive learning pre-training. This model, when paired with the Temporal MAS method, accurately identifying fraudulent claims and significantly attenuates the occurrence of forgetting previously learned knowledge.
- Through extensive experimentation on a real-world medical insurance fraud dataset, we demonstrate evidence of our model's impressive accuracy during protracted online updates. Moreover, our model consistently exhibits lower time and space complexity.

Related Work

Fraud Detection

In recent years, the realm of fraud detection has expanded its influence across various sectors, prominently within credit card operations (Varmedja et al. 2019), the broader financial sphere (Ashtiani and Raahemi 2021), and the insurance domain (Aslam et al. 2022). Machine learning algorithms have proven instrumental in stymieing fraudulent endeavors. In the pursuit of fraud detection, three salient deep learning methodologies emerge. Foremost, Deep Structure Embedding, facilitated by the FraudNE algorithm (Zheng et al. 2018), stands out for its adeptness in retaining intricate, nonlinear structural nuances, offering the capability to embed heterogeneous vertex types into a congruent latent space. Next, GNNs, typified by the PCGNN (Liu et al. 2021), CAREGNN (Dou et al. 2020) and GTAN (Xiang et al. 2023), exhibit robustness when confronted with unstructured datasets. Finally, the deployment of Long Short-Term Memory Networks (LSTM) (Choi et al. 2016; Haque and Tozal 2022; Wiese and Omlin 2009; Jurgovsky et al. 2018) emerges as an effective approach for sequential data interpretation. Notably, these fraud detection techniques are within offline learning paradigms. Consequently, shift in fraud patterns results in diminished accuracy, necessitating a comprehensive model retraining. In contrast, our model circumvents this issue by dynamically updating parameters during online learning, significantly mitigating both time and computational overheads.

Online Learning

Online learning methods can be broadly divided into three categories: (1) weight regularization and optimization strategies, (2) memory management strategies, and (3) network structure adaptation strategies. Weight regularization and optimization strategies maintain the performance of tasks already learned by constraining the weight updates of the model. This intentional limitation restricts weight changes during the learning of new tasks (Liu et al. 2020; Li et al. 2020; Liu et al. 2018). Memory management strategies, assist the model in remembering old tasks by saving and revisiting some of the data from these tasks. Specific applications of this include memory replay (Zhou and Cao 2021), and pseudo-rehearsal (Atkinson et al. 2021; Pomponi, Scardapane, and Uncini 2020; Kase, Tateishi, and Ogata 2022). Network structure adaptation strategies, alleviate the problem of catastrophic forgetting by adjusting the structure of the neural network to accommodate new tasks. Specific applications of this include dynamic expandable networks (Yoon et al. 2017; Yang, Chen, and Liu 2022), progressive neural networks (Rusu et al. 2016), MAS (Aljundi et al. 2018), and PackNet (Mallya and Lazebnik 2018). In terms of combination with GNNs, the memory replay method has more notable achievements, such as ERGNN (Zhou and Cao 2021), and Stream GNN based on generative replay (Wang et al. 2022b). However, many models neglect the structural nuances of fraud and its temporal evolution. Additionally, the mandated storage of historical data is impractical. In contrast, our model seamlessly integrates time and structure information, updates parameters online, and eschews the need for storage of historical data.

Methodology

As shown in Fig. 1, the pipeline of our model can be divided into three stages: *pre-training*, *task-learning*, and *online learning*. In this section, we will firstly formulate our

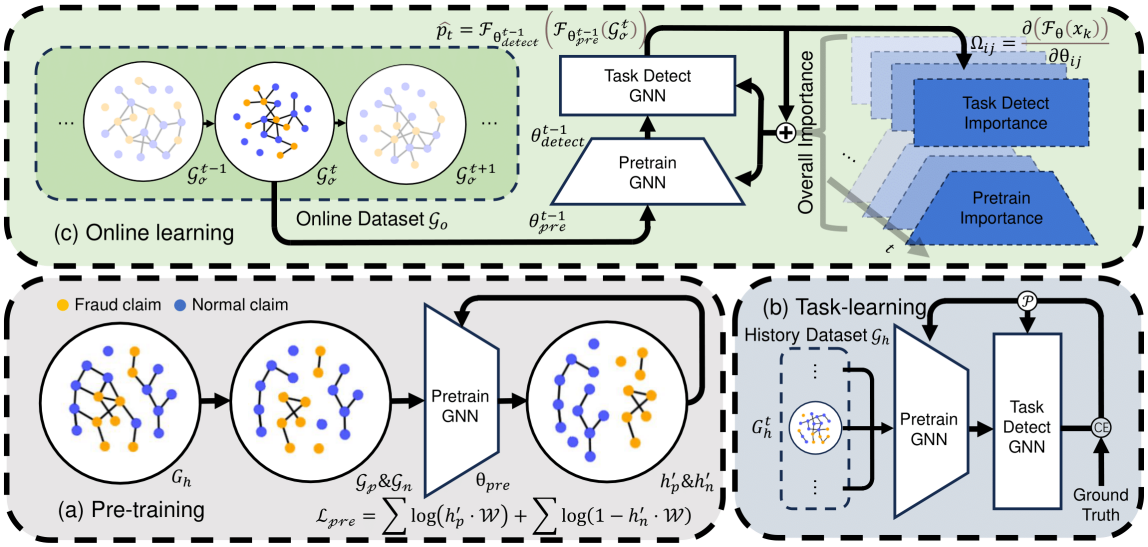

Figure 1: The illustration of our Pre-trained Online Contrastive Learning (POCL) model. (a) Pre-training stage utilizes contrastive learning to learn deep features. (b) Tasking learning stage involves training an offline model using the pre-trained model and the fraud detection model. (c) Online learning stage employs Temporal MAS to continuously update the model.

problem. Then, in the *pre-training stage*, we will introduce the pre-trained model and its optimization strategy. In the *task-learning stage*, we will describe how to combine the pre-trained model and the tasking learning model to train an offline model. Finally, in the *online learning stage*, we will introduce the Temporal MAS online update method, which is used to determine the degree of change in variables.

Problem Formulation

In online medical fraud detection, we define a medical claim graph as $\mathcal{G}(\mathcal{C}, \mathcal{E})$, where $\mathcal{C} = (h_1, h_2, ..., h_{\mathcal{N}_C})$ denotes the set of medical claims, where each medical claim's features is h_i, and $\mathcal{E} = (e_1, e_2, ..., e_{\mathcal{N}})$ represents the edges between claims, in which two connected claims have the same medical provider or beneficiary. Here, $\mathcal{N}_C$ is the number of claims, and $\mathcal{N}$ is the number of edges in the graph. In the online learning scenario, given two part of data: history dataset $\mathcal{G}_h = (\mathcal{G}_h^1, \mathcal{G}_h^2, ..., \mathcal{G}_h^T)$ and online learning dataset $\mathcal{G}_o = (\mathcal{G}_o^1, \mathcal{G}_o^2, ..., \mathcal{G}_o^T)$, our aim is to learn $(\theta_0, \theta_1, ..., \theta_T)$, where θ_0 is the parameter of GNN trained using the traditional training method based on the whole history dataset $\mathcal{G}_h$, while the parameter $\theta_t, t > 0$ is trained on the online learning data $\mathcal{G}_o^t$, in particular, it is incrementally trained based on the parameter of θ_{t-1} using an online learning method. We want to keep the accuracy of model θ as high as possible without retraining the whole model, and expect it to detect situations where some patterns have changed.

Pre-trained Model for Enhanced Robustness

During the pre-training stage, we partition claims in the medical claim graph $\mathcal{G}$ into two sets: the positive set h_p and the negative set h_n. Subsequently, we use the scheme provided by the dataset to build edges in the two graphs. By doing so, we create two distinct graphs: the positive graph $\mathcal{G}p$ and the negative graph $\mathcal{G}_n$. Both graphs are then utilized to train the pre-trained model θ_{pre}.

We use GraphSAGE (Hamilton, Ying, and Leskovec 2017) as our pre-trained model, where the k-th layer is defined as:

$$h_k = \sigma(W_k \cdot \text{MEAN}(\{h_v^{k-1}\} \cup \{h_u^{k-1}, \forall u \in \mathcal{N}(v)\})), \quad (1)$$

where the $\mathcal{N}(v)$ is the neighborhood of node v, and σ is the activation function. Then we obtain two feature embeddings updated by the model, $h_p^{'} = \mathcal{F}_{\theta_{pre}}(h_p)$ and $h_n^{'} = \mathcal{F}_{\theta_{pre}}(h_n)$, which will be used for fraud detection.

We refer to (Veličković et al. 2019), use the binary cross entropy loss function to help recognize the fraud pattern:

$$\mathcal{L}_{pre} = \sum \log(\mathcal{D}(h_p^{'}, s)) + \sum \log(1 - \mathcal{D}(h_n^{'}, s)), \quad (2)$$

where $\mathcal{D}$ is the distance function, $\mathcal{D} = h_p^{'} \cdot \mathcal{W} \cdot s$, $\mathcal{W}$ is the learnable matrix and s is the logic center, which can be set as average of h_p or E. In practice, we set s as E, so that Eq.2 can be simplified as

$$\mathcal{L}_{pre} = \sum \log(h_p^{'} \cdot \mathcal{W}) + \sum \log(1 - h_n^{'} \cdot \mathcal{W}), \quad (3)$$

where $\mathcal{W}$ is the learnable matrix. The loss function is designed to effectively separate positive and negative instances, enabling the model to focus on learning specific patterns associated with medical fraud. It can pull nodes with the same label towards each other and push nodes with different label apart in the feature space so that the model can focus on the relationship between similar nodes and find the difference of fraud nodes and non-fraud nodes.

Detecting Network and Task Learning

The output of pre-trained model modifies feature matrix instead of making predictions, and in downstream task, we use

GAT (Veličković et al. 2018) network θ_{detect} as our detection network, where in k-th layer, the formula for calculating attention coefficients is as follows:

$$\alpha_k^{(i,j)} = \frac{\exp\left(\sigma\left(\mathbf{a}^T[\mathbf{W}_k h_k^{(i)} || \mathbf{W}_k h_k^{(j)}]\right)\right)}{\sum_{m\in\mathcal{N}(i)} \exp\left(\sigma\left(\mathbf{a}^T[\mathbf{W}_k h_k^{(i)} || \mathbf{W}_k h_k^{(m)}]\right)\right)}, \tag{4}$$

where W_k is the weight matrices in layer k, $\mathbf{a}$ is the attention parameter, $||$ is the concatenation operation and σ as the $LeakyReLU$ function. After obtaining the attention coefficients between nodes, we define the k-th layer of GAT as:

$$\text{GATLayer}_k(h_k) = \sigma\left(\sum_{j\in\mathcal{N}(i)} \alpha_k^{(i,j)} \mathbf{W}_k h_k^{(j)}\right), \tag{5}$$

where $\mathcal{N}(i)$ is the set of neighbor node i.

Then, we proceed to combine the pre-trained model with the detection network, creating an end-to-end model. We train the entire model using the historical dataset $\mathcal{G}_h$, which involves training the detection model θ_{detect} and fine-tuning the pre-trained model θ_{pre}. However, a critical challenge arises during this step—ensuring that the pre-trained network θ_{detect} can learn from the labels, while θ_{pre} does not forget the valuable information it has already acquired. To address this issue, we propose a novel loss function that effectively mitigates catastrophic forgetting. We build upon the standard cross-entropy loss function by introducing a penalty term $\mathcal{P} = \sum \log(h'_{p_{\text{pre}}} \cdot \mathcal{W}) + \sum \log(1 - h'_{n_{\text{pre}}} \cdot \mathcal{W})$, so the overall loss $\mathcal{L}_{detect}$ is:

$$\begin{aligned}\mathcal{L}_{detect} = &-\sum (y \cdot \log(p) + (1-y)\cdot \log(1-p)) \\ &+ \sum \log(h'_{p_{\text{pre}}} \cdot \mathcal{W}) + \sum \log(1 - h'_{n_{\text{pre}}} \cdot \mathcal{W}),\end{aligned} \tag{6}$$

where $h'_{p_{\text{pre}}}$ is the positive part of the output of the pre-trained model $\mathcal{F}_{\theta_{pre}}(h)$ and $h'_{n_{pre}}$ is the negative part. Our modified binary cross-entropy loss builds upon the previously mentioned BCE loss $\mathcal{L}_{pre}$ in Eq. 2. By incorporating the the penalty term $\mathcal{P}$ into the loss function, we incentivize the pre-trained network θ_{pre} to preserve its previously learned features while simultaneously helping the whole model learning from the new data. This approach ensures a balanced learning process, preventing the model from overly emphasizing the new data at the expense of forgetting important information obtained during the pre-training stage.

Online Learning and Forgetting Control

Online learning method. To address the challenge of combating constantly changing fraud patterns without retraining the entire model, we employ an online learning method to incrementally update the model. Due to the uniqueness of the financial scenario, we can only access a portion of the online dataset $\mathcal{G}_o$ at each time. At time t, our model can be represented as follows:

$$\hat{p_t} = \mathcal{F}_{\theta_{detect}^{t-1}}(\mathcal{F}_{\theta_{pre}^{t-1}}(\mathcal{G}_o^t)). \tag{7}$$

Here, $\hat{p_t}$ denotes the predicted fraud possibility, $\mathcal{F}_{\theta_{detect}^{t-1}}$ represents the pre-trained model function and $\mathcal{F}_{\theta_{pre}^{t-1}}$ is the detection model function at time $t-1$. These functions learn from the data collected from time 1 to $t-1$.

Temporal MAS approach. In the medical insurance fraud detection scenario, we have observed that new fraud patterns emerge alongside existing ones. This means that failing to promptly learn new patterns or forgetting old ones will both lead to a decrease in accuracy. In order to control what to learn and remember at the period of online learning, we design a novel method based on MAS (Aljundi et al. 2018), called Temporal MAS. Firstly, compared to other incremental learning or online learning methods, this approach does not require storing historical data, which is crucial for our model. Its time efficiency is remarkably high, as it only needs to calculate the gradients of each parameter and merge them, resulting in a fixed space and time overhead. Furthermore, it takes advantage of the evolving nature of fraud patterns in medical insurance fraud detection. By using the momentum method to combine historical importance weights and individual task importance weights, the model achieves a balance between forgetting and retaining knowledge while learning new fraud patterns. This allows the model to adapt to changes in fraud patterns over time, making it more robust and effective in detecting medical fraud.

Firstly, we need to calculate the importance weights of the model's parameters within a single graph. Here, we use the gradients of each parameter as the importance weights for the current task. This is because larger gradients indicate that modifying a particular parameter will have a greater impact on the model's output, making it more susceptible to catastrophic forgetting:

$$\mathcal{F}_\theta(h_k + \delta) - \mathcal{F}_\theta(h_k) \approx \sum_{i,j} \Omega_{ij}(h_k)\delta_{ij}, \tag{8}$$

where h_k is the graph feature and δ is the small change in the parameter θ. We use gradient to calculate these importance weights:

$$\Omega_{ij} = \frac{\partial(\mathcal{F}_\theta(x_k))}{\partial\theta_{ij}}. \tag{9}$$

In the context of online learning for medical insurance fraud detection, new fraud patterns emerge gradually and may replace old patterns. Simply summing up the importance weights of multiple graphs is insufficient. To address this issue, we adopt the concept of momentum, considering the trend of importance weight changes and gradually attenuating the significance of previous weights. By incorporating the momentum idea, we adaptively adjust the importance weights over time, allowing the model to gradually forget less relevant patterns that have not appeared for an extended period. This adaptive mechanism ensures that the model can continuously learn and adapt to new fraud patterns, while still retaining the capability to forget outdated information. We calculate the adjusted importance weights by

$$\mathcal{I}_t = \lambda \times \mathcal{I}_{t-1} + (1-\lambda) \times \Omega_t, \tag{10}$$

where λ is the momentum parameter and Ω_t is the importance weight of Graph $\mathcal{G}_o^t$

After computing the global importance weights, we incorporate them as penalty terms into the loss function $\mathcal{L}$,

$$\begin{aligned}\mathcal{L}_{online} = &-\sum(y_t \cdot \log(p_t) + (1-y_t)\cdot\log(1-p_t)) \\ &+\sum\log(h^{'}_{p^t}\cdot\mathcal{W}) + \sum\log(1-h^{'}_{n^t}\cdot\mathcal{W}) \\ &+\omega\sum_{i,j}\mathcal{I}_t(\theta_{t-1}-\theta_t)^2. \end{aligned} \tag{11}$$

Here y_t is the ground truth of the input, p_t is the prediction of the model, $h^{'}_{p^t}$ is the positive output of the pre-trained model θ^{t-1}_{pre}, ω is the weight of the penalty term and $h^{'}_{n^t}$ is the negative one at time t.

Experiments

We demonstrate the empirical results of our model on a real-world healthcare insurance fraud dataset (we will abbreviate it as medical fraud dataset) and other fraud datasets in two scenarios using simulated timestamps. We first introduce the scale of the medical fraud dataset and experimental setup, then elaborate on the experimental framework and results. Subsequent ablation studies further prove the effectiveness of each model component. Finally, through a case study, we show that the performance of the model is consistent with the expected results.

Dataset and Experimental Setup

To assess the efficacy of our model in real-world medical insurance fraud detection, we leveraged the medical fraud dataset as detailed in (Ma et al. 2023). This extensive medical insurance dataset encompasses real-world information on approximately 200,000 beneficiaries, over 5,000 providers, and around 550,000 medical insurance claims. About 38.1% of these claims were identified and labeled as fraudulent. The most cherish part of this dataset is the meticulous fraud labels and timestamps, which have been annotated by experts. Thus, it has typicality and authority.

For the evaluation process, we select one year of data, subsequently construct many distinct medical claim graphs, each represents a single day. Among them, the initial 15-day period is reserved as a historical dataset, and the rest are online learning datasets. To ensure consistency, every model is trained on this comprehensive historical dataset to create offline models. In the online learning phase, models are adapted dynamically, with the online learning datasets being introduced incrementally. Ideally, in this phase, each model should have access only to a specific segment of the online learning datasets at any point in time. Nevertheless, to mirror real-world scenarios and practical constraints, we allow certain models to either retain or revisit whole information from prior graphs.

Also, to analyze more deeply the effectiveness of our experimental model in different fraud scenarios, we selected two common fraud detection datasets, Amazon (McAuley and Leskovec 2013) and YelpChi (Rayana and Akoglu 2015), for a series of experiments. These experiments were designed by simulating timestamps to generate corresponding online data, thereby mimicking real user behavior and transaction processes.

Baseline. In our experiments, we select several state-of-the-art streaming or online graph neural network models, fraud detection models and OnlineGNN, RetrainGNN, and OfflineGNN, for comparison. Below is a detailed introduction to each of them:

- **OnlineGNN**: This model specifically focuses on online learning scenarios, where it continuously learns from incoming data without revisiting historical data.
- **RetrainGNN**: This model retrains whenever new data is introduced. It discards previous knowledge and starts training with a random weight with whole data.
- **OfflineGNN**: This model is only trained on the entire historical dataset, with no subsequent exposure to any of new data during the online learning phase.
- **ERGNN** (Zhou and Cao 2021): This model uses experience replay to continuously learn a sequence of tasks, addressing the catastrophic forgetting problem.
- **FGN** (Wang et al. 2022a): This model bridges graph learning and lifelong learning by converting continual graph learning into regular graph learning.
- **ContinuesGNN** (Wang et al. 2020): This model is based on continuous learning to learn and maintain patterns through data replay and model regularization.
- **CAREGNN** (Dou et al. 2020): This model uses multi-relation graph to detect fraud, it focuses on fighting against cheaters who use disguises.
- **PCGNN** (Liu et al. 2021): This model uses the pick and choose method to solve the problem of category imbalance in the topological graph of financial relationships.

Experimental Results

We have conducted a comprehensive comparison of our model against advanced baseline models. As depicted in Figure 2(a) and 2(c), our model exhibits remarkably high average monthly accuracy, while Table 1 showcases the impressive results on other evaluation metrics on medical fraud dataset. In other datasets, where experiments were conducted with simulated time, we selected several of the best-performing baselines for experimentation. These experiments demonstrated trends similar to those observed in the medical fraud dataset, as shown in Figure 3. The findings demonstrate that our model achieves nearly the level of RetrainGNN in terms of average monthly accuracy, outperforming other baseline models by a significant margin of 1%-4%.

To further visually observe the performance in a long-term online learning scenario, we plot for the average accuracy decline rate in the first six months and the last six months in Figure 2, which were obtained through data fitting. For better visualization, we have flipped the y-axis. It can be observed that although CAREGNN and PCGNN have impressive performance in the beginning, the accuracy drops quickly as shown in Figure 2(c) and 2(d). Other models exhibit similar decline rates in the first six months, but in the last six months, except for our model and RetrainGNN, other models also experience significant declines. Considering other evaluation metrics like average AUC, and average F1, we also approach the accuracy of RetrainGNN

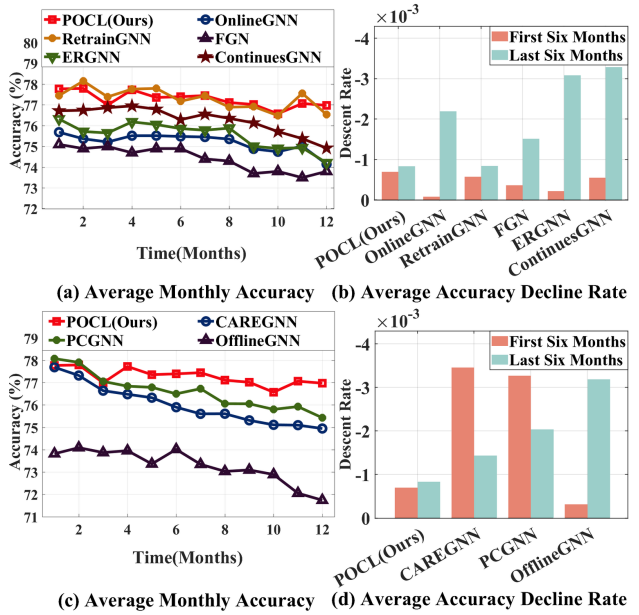

Figure 2: Comparison of different fraud detection methods on average monthly accuracy and average accuracy decline rate for online learning within a year in medical fraud dataset. The average accuracy decline rate is divided into the first six months and the last six months.

and significantly outperform other models. Furthermore, when comparing the excessive time consumption and memory consumption of RetrainGNN, our model demonstrates a training time similar to OnlineGNN, indicating a significant time advantage over RetrainGNN. Additionally, our model exhibits a comparable space cost to OfflineGNN because we do not need to store any historical data and only require storage for a small number of importance weights. This highlights our model's performance at a remarkably high level while showcasing its notable time and space efficiency, thereby attesting to its effectiveness.

Furthermore, a significant decline in accuracy is observed for PCGNN and CAREGNN model without any optimization for online learning. This confirms the existence of dynamic evolution patterns in medical insurance data and underscores the necessity of optimizing models for online learning. It further validates the crucial role of online learning in adapting to the evolving fraud patterns. Additionally, our model's advantages become increasingly apparent over time compared to other models, demonstrating its seamless adaptation to pattern changes with the aid of pre-trained models and the significant mitigation of catastrophic forgetting, facilitated by Temporal MAS method.

It is worth noting that, some models require access to all previous data (e.g., RetrainGNN) or partial data storage (e.g., ERGNN), which is a crucial reason why they need more space in the training process. However, our model overcomes the need for data storage, striking an optimal balance between user privacy protection and exceptional performance.

Model	AUC	F1	Time	Space
OnlineGNN	70.48	48.32	3.54	2.21
OfflineGNN	68.72	46.70	3.16	2.13
ERGNN	73.11	52.58	3.64	2.98
FGN	73.85	56.18	4.35	2.67
ContinuesGNN	78.82	57.64	12.4	10.4
CAREGNN	75.91	54.03	8.74	8.90
PCGNN	78.71	60.45	10.4	9.34
RetrainGNN	80.58	63.77	97.4	11.4
POCL(ours)	80.34	63.53	3.62	2.35

Table 1: The average AUC, average F1 score, the total time (minutes) in the learning process and the average memory consumption (GB) on medical fraud dataset.

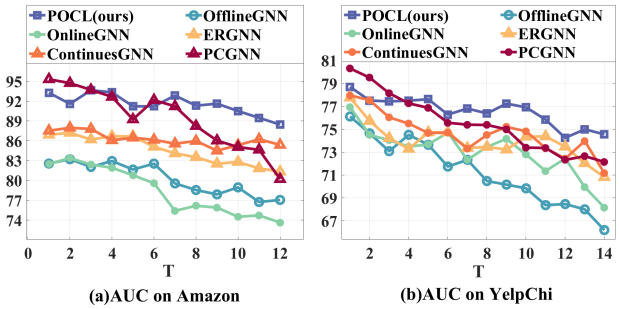

Figure 3: Experiments conducted on common fraud datasets using simulated timestamps. We use AUC as the evaluation metric, because these two datasets are highly imbalanced.

Ablation Experiment

To assess the impact of pre-trained contrastive learning model and online learning with Temporal MAS on the insurance fraud detection model, we conduct a series of ablation experiments. In these experiments, we separately test the pre-trained contrastive learning, the online learning with Temporal MAS, the combined model POCL, and the OnlineGNN for comparison.

- **POCL w/o OL:** We remove Temporal MAS part and employ pre-trained contrastive learning part with direct parameter updates in the following online stages.
- **POCL w/o PCL:** We remove pre-trained contrastive learning part and using a normal GraphSAGE instead, using Temporal MAS to update the whole model.

According to Figure 4, the **POCL w/o OL** model obtained promising results during the early stages. However, as training progresses, this model shows a certain degree of performance decline compared to the original model, although the decline is less pronounced than using the GAT model alone. This suggests that contrastive learning aids the model in identifying fraud patterns but may lead to some performance degradation during later stages of training.

Next, the **POCL w/o PCL** model shows a significant reduction in the later-stage decline, demonstrating the effectiveness of Temporal MAS in mitigating catastrophic forgetting. However, its accuracy shows a considerable decrease

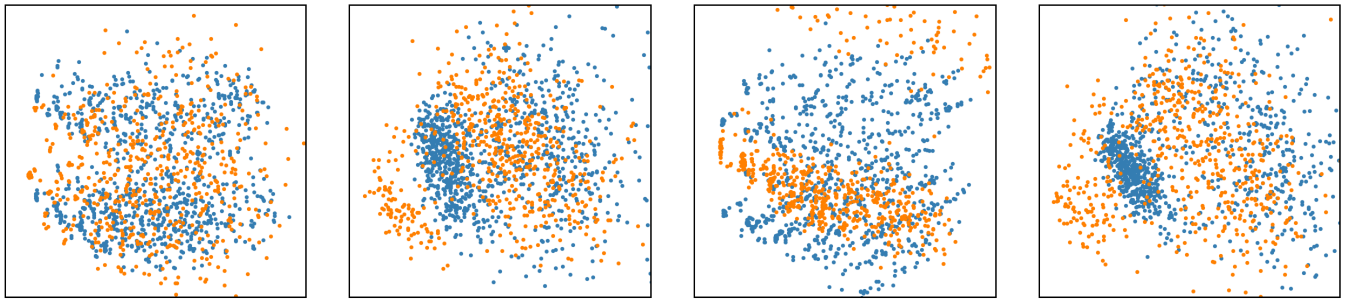

(a) h_{100} (b) h'_{100} (c) h_{300} (d) h'_{300}

Figure 5: The distributions of original features and output features on day 100 and day 300 on medical fraud dataset. In the visual representation, yellow denotes fraudulent nodes, whereas blue denotes normal nodes.

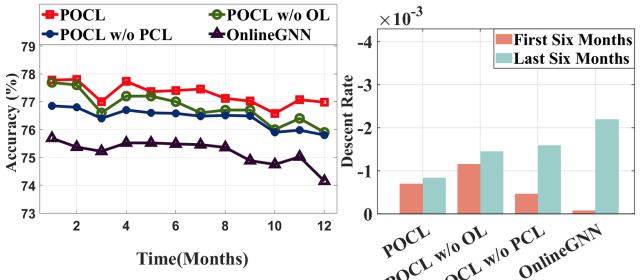

(a) Average Monthly Accuracy (b) Average Decline Rate

Figure 4: An ablation experiment of the proposed POCL w.r.t. average monthly accuracy (a) and average accuracy decline rate (b) on medical fraud dataset.

compared to the original model, indicating weaker fraud recognition capabilities.

Finally, combining contrastive learning with online learning using Temporal MAS method yields satisfactory results. This joint model further reduces the decline observed in both individual models, substantiating the crucial role of contrastive learning in aiding fraud pattern recognition and reducing model update difficulties, as well as the advantage of online learning with Temporal MAS method in preventing forgetting previous knowledge.

Case Study

In the POCL model, the incorporation of a pre-trained model, realized via contrastive learning, plays a pivotal role in fraud representations learning, fortifying the model's robustness in the process. The Temporal MAS online update methodology adeptly navigates the delicate equilibrium between information retention and obsolescence.

We show a case study, focusing on two distinct medical claim graphs at day 100 and day 300, as presented in Figure 5. The nodes colored in yellow represent fraudulent claims, while those in blue denote legitimate transactions. After the dimensionality reduction, these nodes are mapped onto a 2D space. Figure 5(a) and 5(c) depict the original distribution of the features. In contrast, Figure 5(b) and 5(d) showcase the attribute outputs of the pre-trained model.

The pre-trained model amplifies the interspace between nodes of different labels, and fosters a clustering of nodes with an identical-label. This enhancement facilitates the fraud detection network's capacity to discern between the two distinct categories. We notice a significant change in the node distribution patterns between Figure 5(a) and 5(c). Figure 5(a) presents a balanced dispersion of both fraudulent and legitimate nodes. In contrast, Figure 5(c) shows a denser congregation of fraudulent nodes, signaling an evolution in fraud patterns over the 200-day span. Nonetheless, an inspection of Figure 5(b) and 5(d) suggests that, in spite of the modifications in input features, the output feature patterns conferred by the pre-trained model sustain a marked consistency. Such persistence contrasts starkly with the notable shifts observed in Figure 5(a) and 5(c). This observation underscores that the Temporal MAS online learning method reduces the impact of the catastrophic forgetting, thus attesting to the efficacy of the Temporal MAS. Furthermore, the congruence in the output features rendered by the pre-trained model increases robustness. This aids in curbing the magnitude of parameter adjustments during online learning and diminishing the propensity for catastrophic forgetting.

In summation, the POCL model is an effective countermeasure against catastrophic forgetting during extended online updates and exhibits a nuanced capacity to distinguish between the features of fraudulent and legitimate nodes.

Conclusion

In this study, we presented POCL, which combined contrastive learning and online learning to address the evolving landscape of insurance fraud. Our model, which employs contrastive learning to extract and map features, also integrates the Temporal MAS method for online parameter updates. This ensures high performance in long-term online learning scenarios while reducing the computational overhead and training time. An extensive assessment on real-world datasets, pitting our proposed model against many online learning and fraud detection models, revealed the efficacy of our approach. We believe that our model can make a contribution in combating insurance fraud.

Acknowledgments

This work was supported by the National Key R&D Program of China (Grant no. 2022YFB4501704), the National Natural Science Foundation of China (Grant no. 62102287), the Oriental Talents Program Youth Project, the foundation of Key Laboratory of Artificial Intelligence, Ministry of Education, P.R. China and the Shanghai Science and Technology Innovation Action Plan Project (Grant no. 22YS1400600 and 22511100700).

References

Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision*, 139–154.

Anowar, F.; and Sadaoui, S. 2021. Incremental learning framework for real-world fraud detection environment. *Computational Intelligence*, 37(1): 635–656.

Ashtiani, M. N.; and Raahemi, B. 2021. Intelligent fraud detection in financial statements using machine learning and data mining: a systematic literature review. *IEEE Access*, 10: 72504–72525.

Aslam, F.; Hunjra, A. I.; Ftiti, Z.; Louhichi, W.; and Shams, T. 2022. Insurance fraud detection: Evidence from artificial intelligence and machine learning. *Research in International Business and Finance*, 62: 101744.

Atkinson, C.; McCane, B.; Szymanski, L.; and Robins, A. 2021. Pseudo-rehearsal: Achieving deep reinforcement learning without catastrophic forgetting. *Neurocomputing*, 428: 291–307.

Bayram, B.; Köroğlu, B.; and Gönen, M. 2020. Improving Fraud Detection and Concept Drift Adaptation in Credit Card Transactions Using Incremental Gradient Boosting Trees. In *19th IEEE International Conference on Machine Learning and Applications*, 545–550.

Cheng, D.; Ye, Y.; Xiang, S.; Ma, Z.; Zhang, Y.; and Jiang, C. 2023. Anti-money laundering by group-aware deep graph learning. *IEEE Transactions on Knowledge and Data Engineering*.

Choi, E.; Bahadori, M. T.; Schuetz, A.; Stewart, W. F.; and Sun, J. 2016. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine Learning for Healthcare Conference*, 301–318. PMLR.

Cubanski, J.; Neuman, T.; and Freed, M. 2019. The facts on medicare spending and financing. *Henry J. Kaiser Family Foundation, San Francisco*.

Dou, Y.; Liu, Z.; Sun, L.; Deng, Y.; Peng, H.; and Yu, P. S. 2020. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 315–324.

Dua, P.; and Bais, S. 2014. Supervised learning methods for fraud detection in healthcare insurance. *Machine Learning in Healthcare Informatics*, 261–285.

Fiore, U.; De Santis, A.; Perla, F.; Zanetti, P.; and Palmieri, F. 2019. Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479: 448–455.

Gao, Y.; Wang, X.; He, X.; Feng, H.; and Zhang, Y. 2023. Rumor detection with self-supervised learning on texts and social graph. *Frontiers of Computer Science*, 17(4): 174611.

Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30.

Haque, M. E.; and Tozal, M. E. 2022. Identifying health insurance claim frauds using mixture of clinical concepts. *IEEE Transactions on Services Computing*, 15(4): 2356–2367.

Jurgovsky, J.; Granitzer, M.; Ziegler, K.; Calabretto, S.; Portier, P.-E.; He-Guelton, L.; and Caelen, O. 2018. Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100: 234–245.

Kase, K.; Tateishi, A.; and Ogata, T. 2022. Robot task learning with motor babbling using pseudo rehearsal. *IEEE Robotics and Automation Letters*, 7(3): 8377–8382.

Le-Khac, P. H.; Healy, G.; and Smeaton, A. F. 2020. Contrastive representation learning: A framework and review. *IEEE Access*, 8: 193907–193934.

Lebichot, B.; Paldino, G. M.; Bontempi, G.; Siblini, W.; He-Guelton, L.; and Oblé, F. 2020. Incremental learning strategies for credit cards fraud detection. In *IEEE 7th International Conference on Data Science and Advanced Analytics*, 785–786. IEEE.

Li, Y.; Zhang, R.; Lu, J. C.; and Shechtman, E. 2020. Few-shot image generation with elastic weight consolidation. *Advances in Neural Information Processing Systems*, 33: 15885–15896.

Liu, L.; Kuang, Z.; Chen, Y.; Xue, J.-H.; Yang, W.; and Zhang, W. 2020. Incdet: In defense of elastic weight consolidation for incremental object detection. *IEEE Transactions on Neural Networks and Learning Systems*, 32(6): 2306–2319.

Liu, X.; Masana, M.; Herranz, L.; Van de Weijer, J.; Lopez, A. M.; and Bagdanov, A. D. 2018. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *24th International Conference on Pattern Recognition*, 2262–2268. IEEE.

Liu, Y.; Ao, X.; Qin, Z.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2021. Pick and choose: a GNN-based imbalanced learning approach for fraud detection. In *Proceedings of The Web Conference 2021*, 3168–3177.

Ma, J.; Li, F.; Zhang, R.; Xu, Z.; Cheng, D.; Ouyang, Y.; Zhao, R.; Zheng, J.; Zheng, Y.; and Jiang, C. 2023. Fighting against organized fraudsters using risk diffusion-based parallel graph neural network. In *International Joint Conference on Artificial Intelligence*, 6138–6146.

Mallya, A.; and Lazebnik, S. 2018. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7765–7773.

McAuley, J. J.; and Leskovec, J. 2013. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd International Conference on World Wide Web*, 897–908.

McDowell, T. N. 1987. The medicare-medicaid anti-fraud and abuse amendments: Their impact on the present health care system. *Emory Law Journal*, 36: 691.

Morris, L. 2009. Combating fraud in health care: an essential component of any cost containment strategy. *Health Affairs*, 28(5): 1351–1356.

NHCAA. 2021. The challenge of health care fraud. https://www.nhcaa.org/tools-insights/about-health-care-fraud/the-challenge-of-health-care-fraud. Accessed: 2023-7-19.

Pomponi, J.; Scardapane, S.; and Uncini, A. 2020. Pseudo-rehearsal for continual learning with normalizing flows. *arXiv preprint arXiv:2007.02443*.

Rayana, S.; and Akoglu, L. 2015. Collective opinion spam detection: Bridging review networks and metadata. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 985–994.

Roy, A.; Sun, J.; Mahoney, R.; Alonzi, L.; Adams, S.; and Beling, P. 2018. Deep learning detecting fraud in credit card transactions. In *2018 Systems and Information Engineering Design Symposium*, 129–134. IEEE.

Rusu, A. A.; Rabinowitz, N. C.; Desjardins, G.; Soyer, H.; Kirkpatrick, J.; Kavukcuoglu, K.; Pascanu, R.; and Hadsell, R. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671*.

Sadreddin, A.; and Sadaoui, S. 2022. Incremental feature learning for fraud data stream. In *ICAART (3)*, 268–275.

Sisko, A. M.; Keehan, S. P.; Poisal, J. A.; Cuckler, G. A.; Smith, S. D.; Madison, A. J.; Rennie, K. E.; and Hardesty, J. C. 2019. National health expenditure projections, 2018–27: economic and demographic trends drive spending and enrollment growth. *Health Affairs*, 38(3): 491–501.

Thornton, D.; Mueller, R. M.; Schoutsen, P.; and Van Hillegersberg, J. 2013. Predicting healthcare fraud in medicaid: a multidimensional data model and analysis techniques for fraud detection. *Procedia Technology*, 9: 1252–1264.

Timofeyev, Y.; and Jakovljevic, M. 2022. Fraud and corruption in healthcare. *Frontiers in Public Health*, 10.

Varmedja, D.; Karanovic, M.; Sladojevic, S.; Arsenovic, M.; and Anderla, A. 2019. Credit card fraud detection-machine learning methods. In *18th International Symposium Infoteh-Jahorina*, 1–5. IEEE.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph attention networks. In *International Conference on Learning Representations*.

Veličković, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep graph infomax. In *International Conference on Learning Representations*.

Wang, C.; Qiu, Y.; Gao, D.; and Scherer, S. 2022a. Lifelong graph learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13719–13728.

Wang, J.; Song, G.; Wu, Y.; and Wang, L. 2020. Streaming graph neural networks via continual learning. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, 1515–1524.

Wang, J.; Zhu, W.; Song, G.; and Wang, L. 2022b. Streaming graph neural networks with generative replay. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1878–1888.

Wiese, B.; and Omlin, C. 2009. Credit card transactions, fraud detection, and machine learning: Modelling time with LSTM recurrent neural networks. In *Innovations in Neural Information Paradigms and Applications*, 231–268. Springer.

Wu, Y.; Dobriban, E.; and Davidson, S. 2020. Deltagrad: Rapid retraining of machine learning models. In *International Conference on Machine Learning*, 10355–10366. PMLR.

Xiang, S.; Zhu, M.; Cheng, D.; Li, E.; Zhao, R.; Ouyang, Y.; Chen, L.; and Zheng, Y. 2023. Semi-supervised credit card fraud detection via attribute-driven graph representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14557–14565.

Yang, Y.; Chen, B.; and Liu, H. 2022. Bayesian compression for dynamically expandable networks. *Pattern Recognition*, 122: 108260.

Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2017. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*.

Zheng, M.; Zhou, C.; Wu, J.; Pan, S.; Shi, J.; and Guo, L. 2018. FraudNe: A joint embedding approach for fraud detection. In *International Joint Conference on Neural Networks*, 1–8. IEEE.

Zhou, F.; and Cao, C. 2021. Overcoming catastrophic forgetting in graph neural networks with experience replay. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4714–4722.