



NeutronStream: A **Dynamic** GNN Training Framework with **Sliding Window** for Graph Streams

Chaoyi Chen, Dechao Gao, Yanfeng Zhang, Qiange Wang, Zhenbo Fu, Xuecang Zhang, Junhua Zhu, Yu Gu, Ge Yu

> Northeastern University, China Huawei Technologies Co., Ltd.



Dynamic Graph Neural Network

















Characteristic: Spatial-Temporal Dependency

- Spatial Dependency: a node status depends on its neighbors
- Temporal Dependency: a node status depends on its previous status



Characteristic: Spatial-Temporal Dependency

- Spatial Dependency: a node status depends on its neighbors
- Temporal Dependency: a node status depends on its previous status

Temporal Dependency



11

Key Observations : Events in the same area should update parameters collaboratively to better capture spatial dependencies

Update event stream

 $e_1 e_2 e_3 e_4 e_5 e_6 \cdots$





Six events are concentrated in one area

The traditional batched training method forcibly cuts off the event stream and ignores the spatial locality between events



The traditional batched training method forcibly cuts off the event stream and ignores the spatial locality between events



The traditional batched training method forcibly cuts off the event stream and ignores the spatial locality between events



• Temporal Dependency: a node status depends on its previous status

Existing frameworks sequential iterative processing



• Temporal Dependency: a node status depends on its previous status

Existing frameworks sequential iterative processing



- Without Temporal Dependency : Ignore time, execute events in parallel
- With Temporal Dependency : Keep time, execute events chronologically



$$e_1 e_2 e_3 \cdots$$

(b) With Temporal Dependency

- Without Temporal Dependency : Ignore time, execute events in parallel
- With Temporal Dependency : Keep time, execute events chronologically



- Without Temporal Dependency : Ignore time, execute events in parallel
- With Temporal Dependency : keep time, execute events chronologically





Challenge 1: The traditional batched training mode fail to capture the real-time structural evolution information





Challenge 2: The temporal dependency nature makes parallelism optimization hard to design



Our Contributions

We propose NeutronStream, a high-performance Dynamic GNN framework

 Solution 1: Propose a new incremental learning mode with a sliding window for training on graph streams

 Solution 2: Propose a fine-grained event parallel execution scheme
Challenge 2
Training performance improve

Incremental Mode with Sliding Window (denoted as NS-Slide)

We design a sliding window to select consecutive events from the input event stream

Update event stream





Incremental Mode with Sliding Window (denoted as NS-Slide)

We design a sliding window to select consecutive events from the input event stream

Update event stream





Incremental Mode with Sliding Window (denoted as NS-Slide)

We design a sliding window to select consecutive events from the input event stream

Update event stream





Incremental Mode with Sliding Window (denoted as NS-Slide)

We design a sliding window to select consecutive events from the input event stream

Update event stream



Dynamic social network graph



Incremental Mode with Sliding Window (denoted as NS-Slide)

We design a sliding window to select consecutive events from the input event stream

Update event stream



The window size determines how much data can be fed into the model Key Observations : A continuous segment of events is usually concentrated in a local area in the graph

Update event stream





The window size determines how much data can be fed into the model Key Observations : A continuous segment of events is usually concentrated in a local area in the graph

Update event stream





The window size determines how much data can be fed into the model Key Observations : A continuous segment of events is usually concentrated in a local area in the graph

Update event stream





The window size determines how much data can be fed into the model Key Observations : A continuous segment of events is usually concentrated in a local area in the graph

Update event stream





The window size should be dynamically adjustable

Adaptive Adjustment of Window Size (denoted as NS-AdaSlide) The window size can be adaptively adjusted for capturing the complete spatial dependencies

Update event stream

<i>e</i> ₁	<i>e</i> ₂	<i>e</i> ₃	e_4	e_{5}	e_{6}	<i>e</i> ₇	e_{8}	•••
-----------------------	-----------------------	-----------------------	-------	---------	---------	-----------------------	---------	-----





Adaptive Adjustment of Window Size (denoted as NS-AdaSlide) The window size can be adaptively adjusted for capturing the complete spatial dependencies

Update event stream

$e_1 e_2 e_3 e_3$	$e_4 e_5 e_6$	<i>e</i> ₇ <i>e</i> ₈
-------------------	---------------	---







Key Observations : An event only involves a subgraph. If there are no readwrite conflicts between two events, they can be executed in parallel.



A subgraph consists of event nodes and their neighbors, where

- Event nodes are write sets
- Event nodes and their neighbors are read sets

Key Observations : An event only involves a subgraph. If there are no readwrite conflicts between two events, they can be executed in parallel.



The purple area is the subgraph of events e_1

- Write set: {1, 2}
- Read set: {1, 2, 3, 4 }

Key Observations : An event only involves a subgraph. If there are no readwrite conflicts between two events, they can be executed in parallel.



The green area is the subgraph of events e_2

- Write set: {5, 8}
- Read set: {5, 8, 6, 7 }

Key Observations : An event only involves a subgraph. If there are no readwrite conflicts between two events, they can be executed in parallel.



Key Observations : An event only involves a subgraph. If there are no readwrite conflicts between two events, they can be executed in parallel.



Event Dependency Graph $\mathcal{G}_{\mathcal{E}} = (\mathcal{V}, \mathcal{E})$

Each node represents an event. Each edge represents a dependency.



Event dependency graph

Event Dependency Graph $\mathcal{G}_{\mathcal{E}} = (\mathcal{V}, \mathcal{E})$

Each node represents an event. Each edge represents a dependency.



Event dependency graph

NeutronStream

We deliver a dynamic GNN training framework NeutronStream , which consists of 10,336 LoC



NeutronStream architecture

- . Sliding-window-based Training Method
- 2. Dependency-Graph-Driven Event Parallelizing

NeutronStream

We provides a set of easy-to-use APIs



NeutronStream architecture



DyRep implementation with NeutronStream APIs $_{42}$

NeutronStream

We design a built-in graph storage structure that supports efficient dynamic updates



Experimental Setups

Baseline: Torch-Batch, NS-Batch (NeutronStream), NS-Slide (NeutronStream) and NS-AdaSlide (NeutronStream).

Platforms:

An Aliyun ECS cluster (ecs.g5.6xlarge instance, 24 vCPUs, 96GB RAM)

Models and datasets:

Dataset	V	E .init	E .final	evt.num	
Social Evolution [29]	84	575	794	54,369	'
Github [4]	284	298	4,131	20,726	
DNC [21]	2,029	0	5,598	39,264	
UCI [21]	1,899	0	20,296	59,835	
Reality [7]	6,809	0	9,484	52,052	
Slashdot [12]	51,083	0	131,175	140,778	

Table 2: Dataset statistics

3 Dynamic GNN
DyRep, LDG, DGNN

□ 6 real-world graphs.

Environment

- Ubuntu 18.04 LTS
- PyTorch 1.12.1

AUC Comparison

Model	Dataset	Training Method					
Model		Torch-Batch	NS-Batch	NS-Slide	NS-AdaSlide		
DyRep	Social	83.01%	83.01%	86.15%	89.32%		
	Github	73.46 %	73.46%	77.54%	79.28%		
	DNC	63.15%	63.15%	63.49%	66.18%		
	UCI	62.46%	62.46%	63.53%	65.68%		
LDG	Social	87.07%	87.07%	88.52%	92.98%		
	Github	74.34%	74.34%	78.10%	79.16%		
	DNC	64.62%	64.62%	66.50%	69.41 %		
	UCI	62.16%	62.16%	64.38 %	66.57%		
DGNN	Social	97.21%	97.21%	97.61%	97.67%		
	Github	81.94%	81.94%	82.54%	84.45%		
	DNC	86.04%	86.04 %	87.72%	88.81 %		
	UCI	78.45%	78.45%	81.14 %	82.09%		

Table 3: AUC comparison of three training methods

- The batch method (Torch-Batch and NS-Batch) has the lowest AUC because it cuts off the data stream, resulting in the loss of training information.
- (2) The adaptive sliding method (**NS-AdaSlide**) achieves the **highest AUC** by effectively capturing the spatial-temporal locality between events.

Runtime Comparison



Compared with the **Torch-Batch**, **NS-Batch** achieves **1.48x – 5.87x** on 3 Dynamic GNNs and six real datasets.

Runtime Comparison



Compared with the **Torch-Batch**, **NS-AdaSlide** have lowest parallelism. However, thanks to our system optimization, **NS-AdaSlide** can achieve **1.27X-4.44X** speedups on 3 Dynamic GNNs and six real datasets.

NeutronStream: A Dynamic GNN Training Framework with Sliding Window for Graph Streams.

Proposing a new incremental learning mode with a sliding window for training on graph streams

We design a sliding-window-based method to incrementally train models for capturing the spatialtemporal dependency of events.

NeutronStream: A Dynamic GNN Training Framework with Sliding Window for Graph Streams.

Proposing a new incremental learning mode with a sliding window for training on graph streams

We design a sliding-window-based method to incrementally train models for capturing the spatialtemporal dependency of events.

Proposing a fine-grained event parallel execution scheme

We build a dependency graph analysis method that identifies the events having no node-updating conflicts and processes them in parallel.

NeutronStream: A Dynamic GNN Training Framework with Sliding Window for Graph Streams.

Proposing a new incremental learning mode with a sliding window for training on graph streams

We design a sliding-window-based method to incrementally train models for capturing the spatialtemporal dependency of events.

Proposing a fine-grained event parallel execution scheme

We build a dependency graph analysis method that identifies the events having no node-updating conflicts and processes them in parallel.

Delivering a dynamic GNN system

We design and implement NeutronStream, a dynamic GNN system that achieves 1.48X-5.87X speedups over PyTorch.

NeutronStream: A Dynamic GNN Training Framework with Sliding Window for Graph Streams.

Proposing a new incremental learning mode with a sliding window for training on graph streams

We design a sliding-window-based method to incrementally train models for capturing the spatialtemporal dependency of events.

Proposing a fine-grained event parallel execution scheme

We build a dependency graph analysis method that identifies the events having no node-updating conflicts and processes them in parallel.

Delivering a dynamic GNN system

We design and implement NeutronStream, a dynamic GNN system that achieves 1.48X-5.87X speedups over PyTorch.

The codes are publicly available on github

https://github.com/iDC-NEU/NeutronStream