



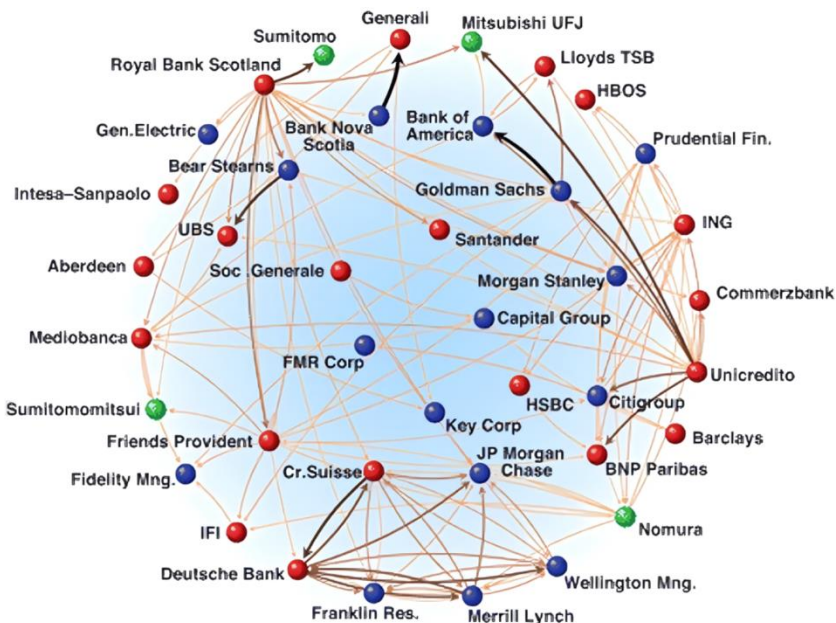
OblivGNN: Oblivious Inference on Transductive and Inductive Graph Neural Network

Zhibo Xu^{1,2}, Shangqi Lai², Xiaoning Liu³, Alsharif Abuadbba², Xingliang Yuan⁴, and Xun Yi³
¹Monash University, ²CSIRO's Data61, ³RMIT University, ⁴The University of Melbourne

GNN Tasks

Node Classification


(Graph Convolutional Network [Kipf et al. (ICLR'17)])



Bank

Graph Classification

(GraphSAGE [Hamilton et al. NIPS'17])



Halicin [3]

Powerful antibiotics discovered using AI
Machine learning spots molecules that work even against 'untargetable' bacteria.

Drug discovery

Link Prediction

(GraphSAGE [Hamilton et al. NIPS'17])

Pins: Visual bookmarks someone has saved from the internet to a board they've created.
Pin features: Image, text, links

Recommendation systems

Graph Neural Networks are specifically designed neural architectures operated on graph-structure data

--Graph Neural Networks: Foundations, Frontiers, and Applications by Lingfei et al.

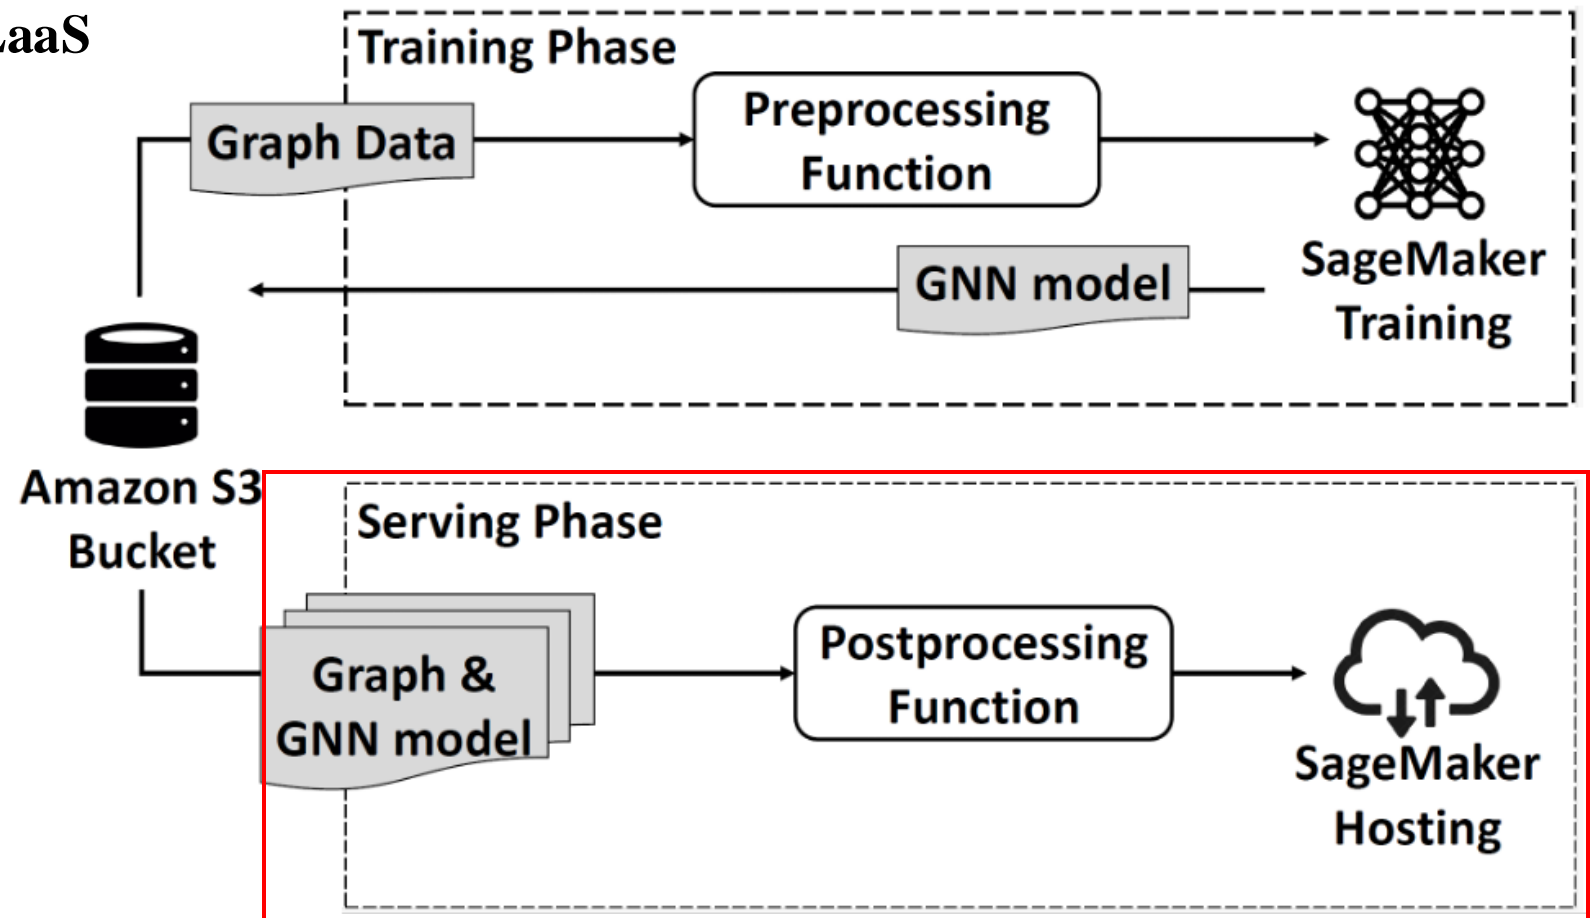
GNNs in Machine Learning as a Service (MLaaS)

GNN is increasingly featured on MLaaS platforms

Amazon: SageMaker Support for DGL

Google: Neo4j & Google Cloud Vertex AI

Microsoft: Azure ML Spektral

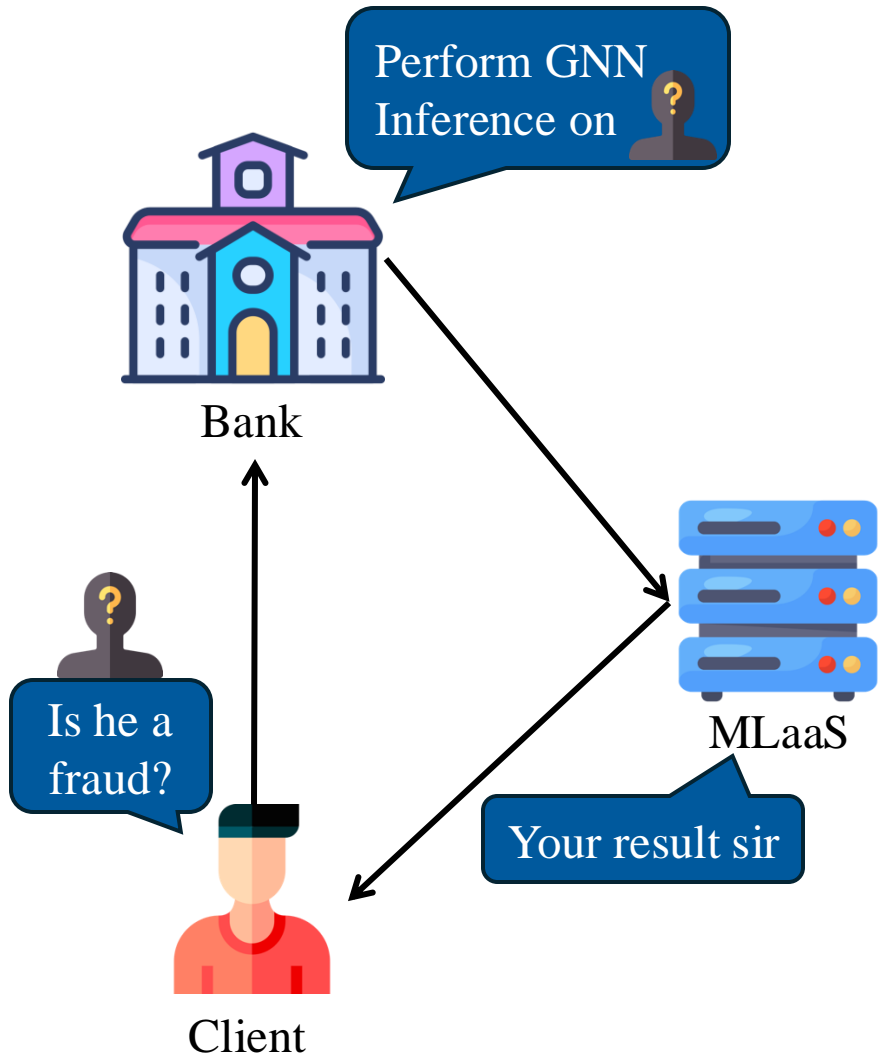


AWS SageMaker for GNN training and inference

Build a GNN-based real-time fraud detection solution using Amazon SageMaker, Amazon Neptune, and the Deep Graph Library:

<https://aws.amazon.com/blogs/machine-learning/build-a-gnn-based-real-time-fraud-detection-solution-using-amazon-sagemaker-amazon-neptune-and-the-deep-graph-library/>

Privacy Concerns



Privacy Concerns:

- Expose sensitive training/inference graph to MLaaS
 - Collecting training graphs often requires a large amount of human, computing, and economic resource
 - Graph data is sensitive by nature, e.g., users' financial transactions, private friendships
- Expose proprietary GNN model parameters and node features to MLaaS

Related Works in PPML

Traditional PPML Frameworks

Trident, Chameleon, Falcon, GAZELLE, MiniONN, Delphi, ABY³, SecureML, BLAZE, XONN, AriaNN, CryptGPU, SecureNN

Existing Secure GNN Frameworks

SecGNN, CryptoGCN, LinGCN

- Cannot support graph-structured data
- Do not offer full protection of graph structure information
- Leak degree information
- Not support the diverse settings of GNN deployment
- Heavy computation cost (via FHE), heavy communication cost (via ASS) due to the large size of the graph

GNNs

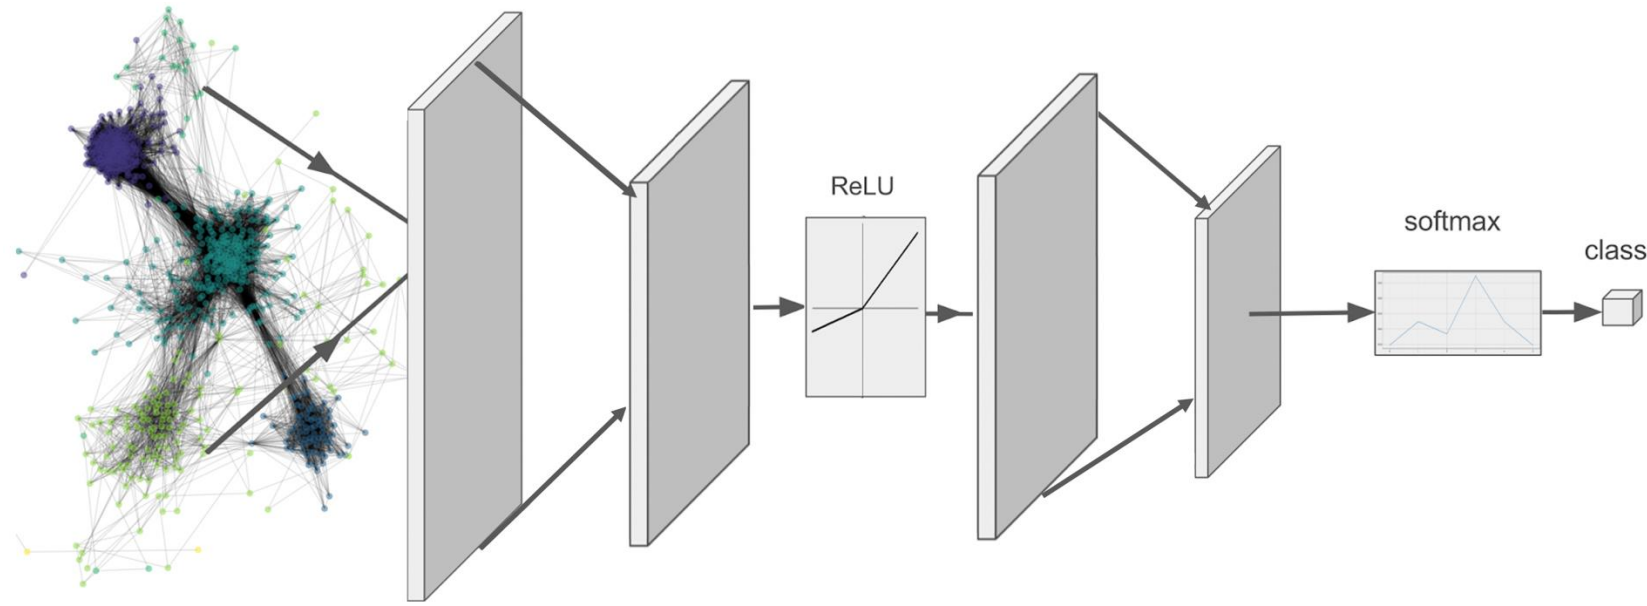
$$\mathbf{Z} = \text{Softmax}(\hat{\mathbf{A}} \text{ReLU}(\hat{\mathbf{A}}\hat{\mathbf{F}}\mathbf{W}_0)\mathbf{W}_1)$$

- \mathbf{W}_0 and \mathbf{W}_1 are two trainable weights
- $\hat{\mathbf{A}}$ is the symmetric normalized adjacency matrix
- $\hat{\mathbf{F}}$ is the normalized feature matrix

- Activation functions:

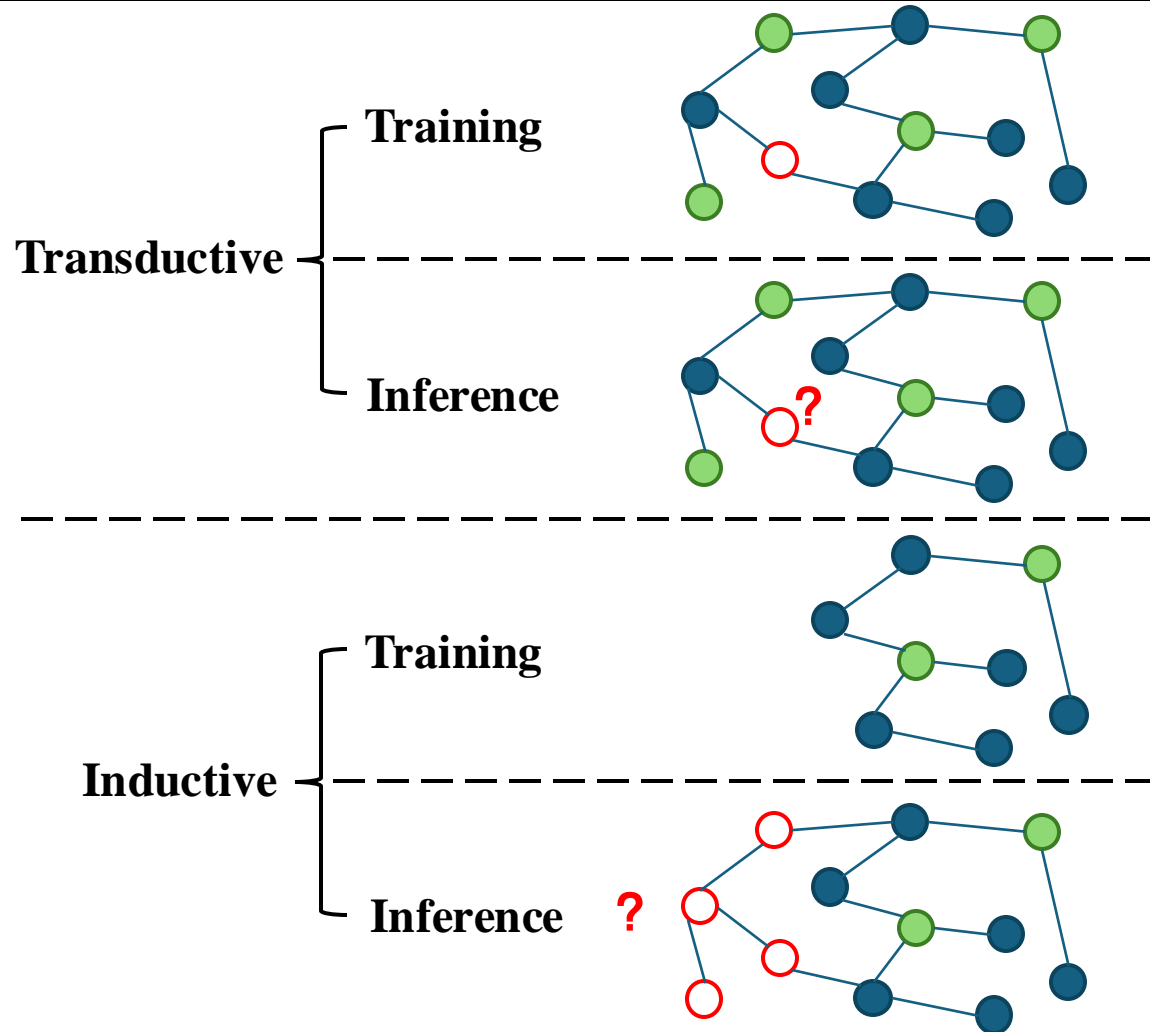
- $\text{ReLU}(x) = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$

- Softmax: $z_i = \frac{e^{x_i}}{\sum_{j \in [1, C]} e^{x_j}}, i \in [1, C]$



GNN Settings

Node Classification



Transductive:

- Unlabelled nodes and their connections exist in the training
- Graph for training and inference remains the same

Inductive:

- *New/updated* nodes, features, connections appear in the *inference*

Function Secret Sharing

Function Secret Sharing

Distributed Point Functions:

$\text{KeyGen}(\alpha, \beta) \rightarrow k_0, k_1$

$\text{Eval}(k_b, x) \rightarrow \llbracket y \rrbracket_b$

$$\text{Eval}(k_0, x) + \text{Eval}(k_1, x) = \begin{cases} \beta, & \text{if } x = \alpha \\ 0, & \text{otherwise} \end{cases}$$

Equality Test:

$\text{KeyGen}^=(\alpha = \gamma, \beta = 1) \rightarrow k_0^=, k_1^=$

$\text{Eval}^=(k_b^=, x) \rightarrow \llbracket y \rrbracket_b$

$$\text{Eval}^=(k_0, x') + \text{Eval}^=(k_1, x') = \begin{cases} y = 1, & \text{if } x' = \gamma \\ 0, & \text{otherwise} \end{cases}$$

Comparison:

$\text{KeyGen}^<(\alpha = \gamma, \beta = 1) \rightarrow k_0^<, k_1^<$

$\text{Eval}^<(k_b^<, x) \rightarrow \llbracket y \rrbracket_b$

$$\text{Eval}^<(k_0, x') + \text{Eval}^<(k_1, x') = \begin{cases} y = 1, & \text{if } x' \leq \gamma \\ 0, & \text{if } x' > \gamma \end{cases}$$

Arithmetic FSS:

Multiplication:

$\text{KeyGen}^\times(g^\circ, r_{in}^1, r_{in}^2, r_{out}) \rightarrow k_0^\times, k_1^\times$

$\text{Eval}^\times(k_b^\times, x'_1, x'_2) \rightarrow g_b^\circ(x_1 \times x_2) + r_{out}$

$$\text{Eval}^\times(k_0^\times, x'_1, x'_2) + \text{Eval}^\times(k_1^\times, x'_1, x'_2) = x_1 \times x_2 + r_{out}$$

Addition:

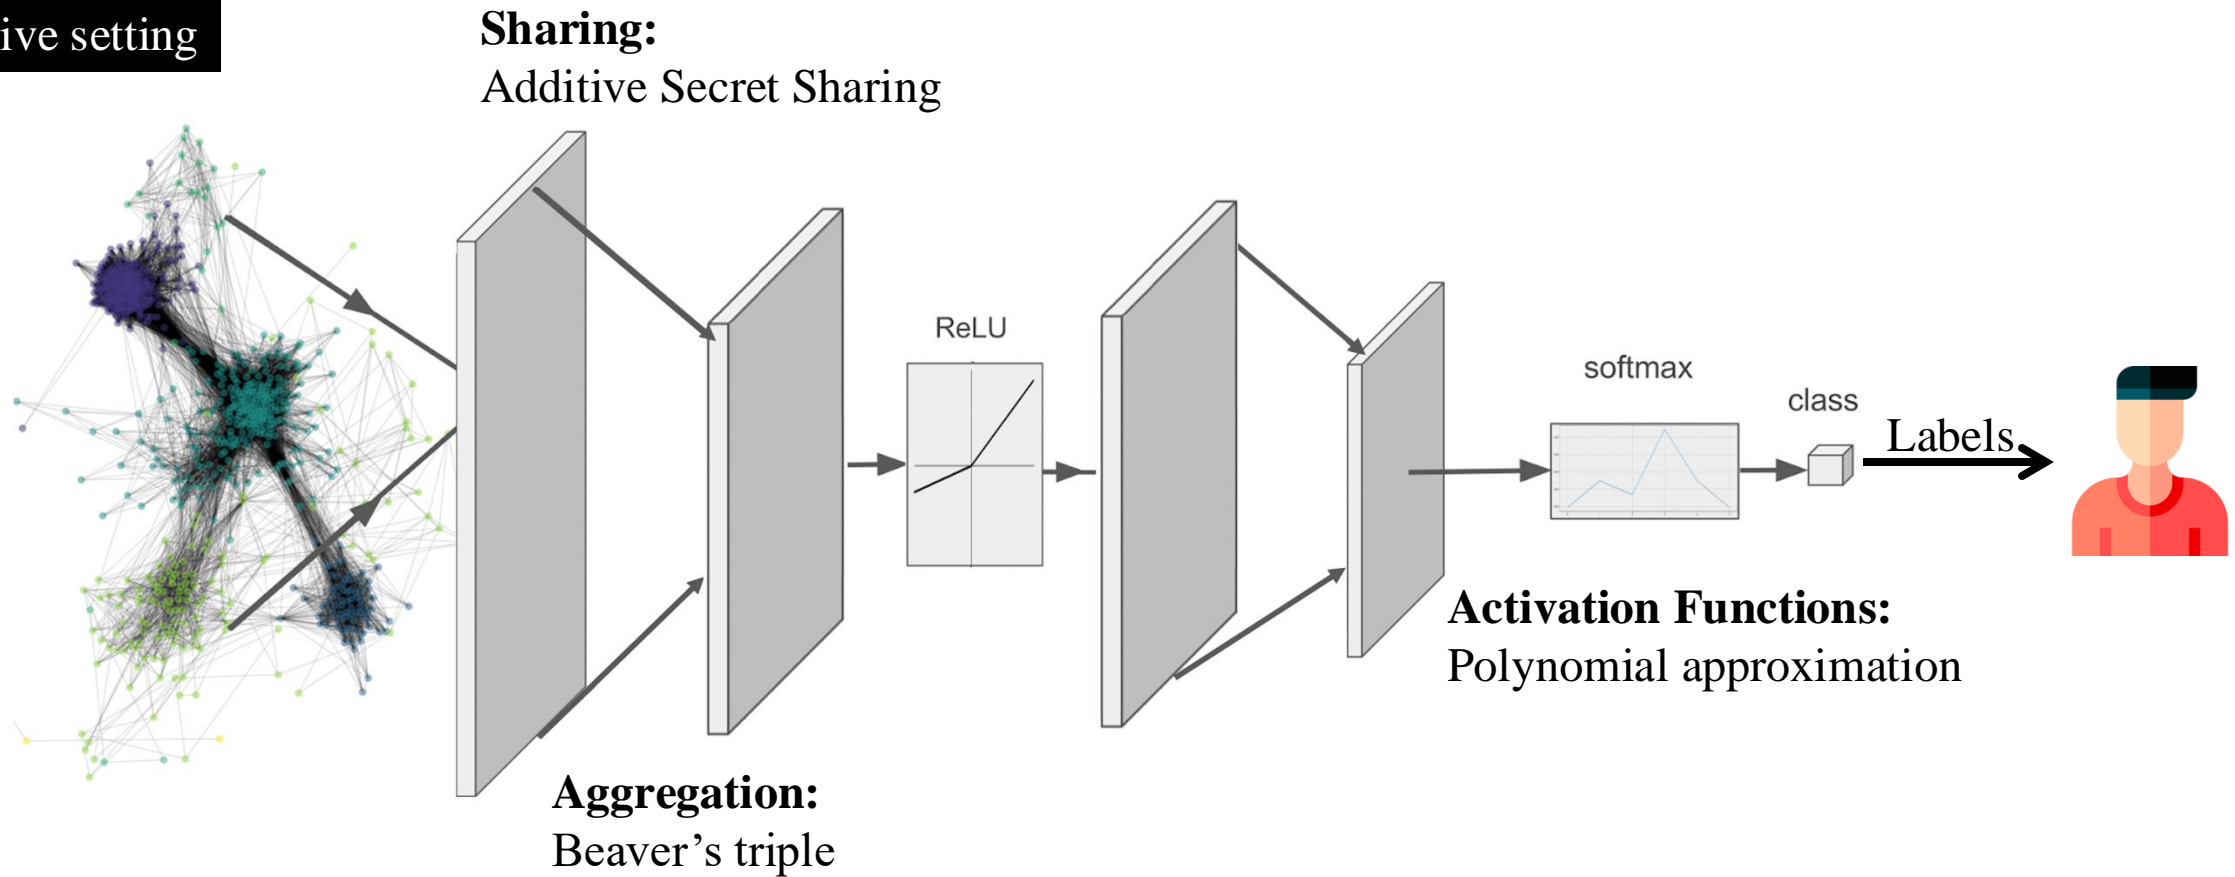
$\text{KeyGen}^+(g^\circ, r_{in}^1, r_{in}^2, r_{out}) \rightarrow k_0^+, k_1^+$

$\text{Eval}^+(k_b^+, x'_1, x'_2) \rightarrow g_b^\circ(x_1 + x_2) + r_{out}$

$$\text{Eval}^+(k_0^+, x'_1, x'_2) + \text{Eval}^+(k_1^+, x'_1, x'_2) = x_1 + x_2 + r_{out}$$

Strawman Approach


Transductive setting



Strawman Approach

Inductive setting


Graph update:
update the graph



Aggregation:
Beaver's triple



I can observe the update pattern



Activation Functions:
Polynomial approximation



class

Labels



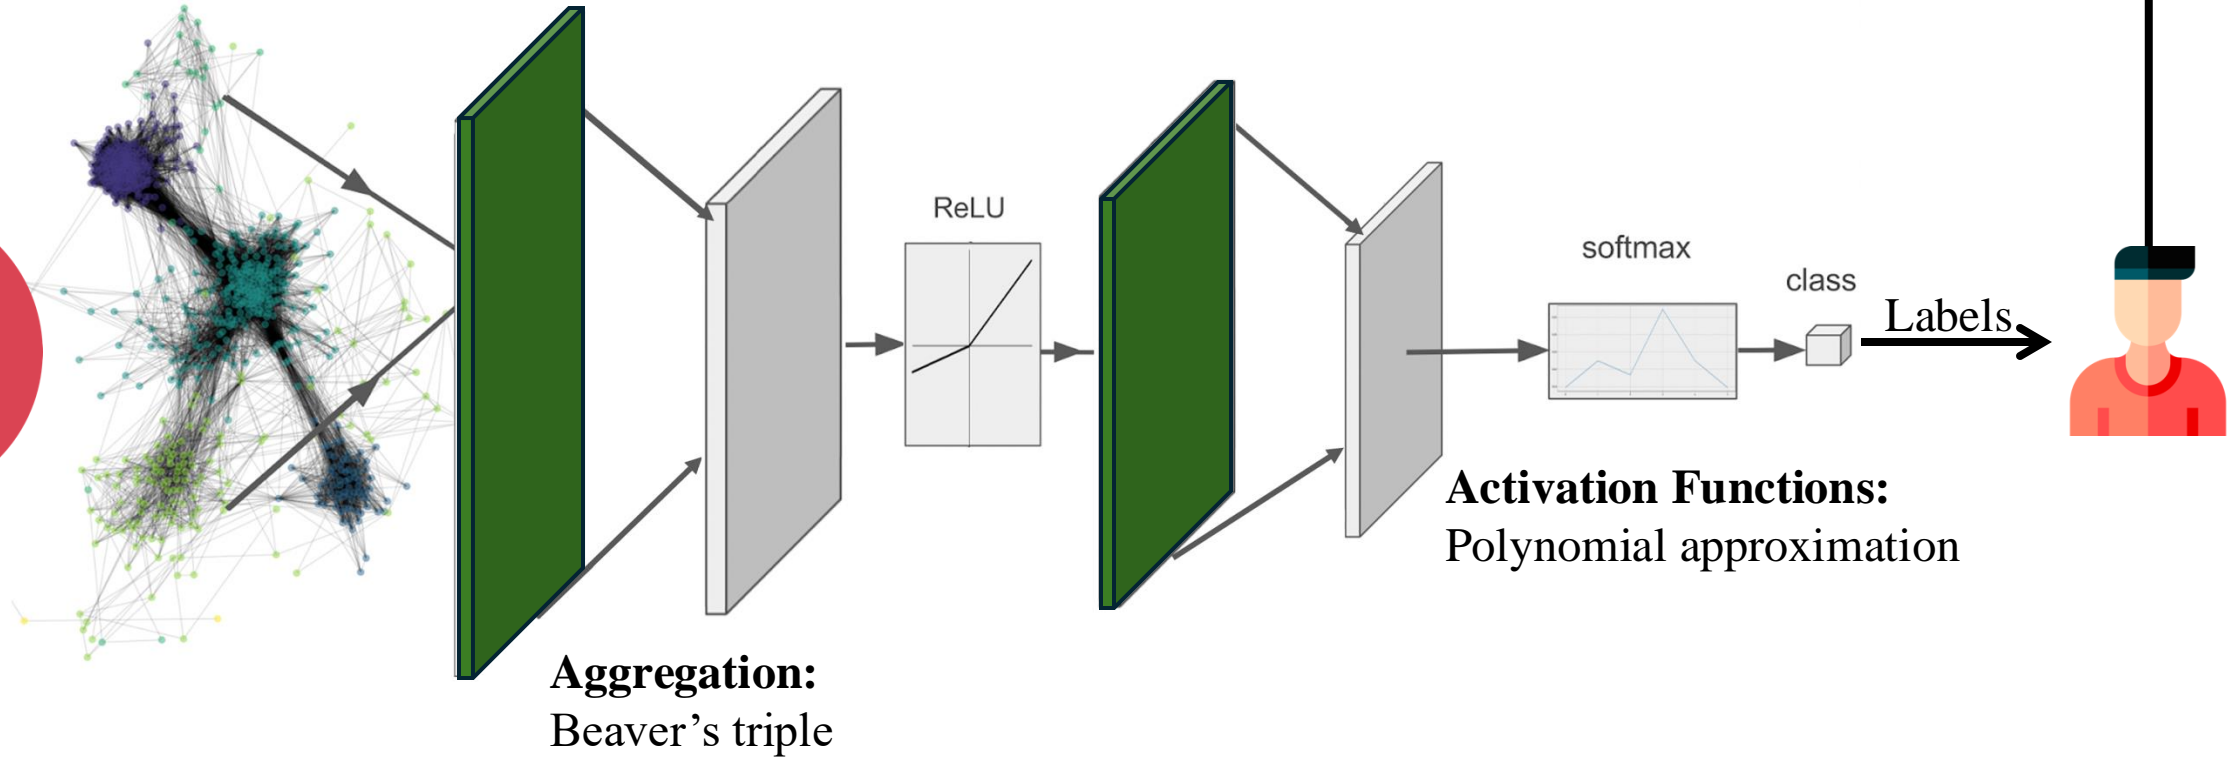
Problem:

1. Leak graph update access, suffering from leakage attack [Falzon and Paterson, ESORICS'22]

Strawman Approach

Inductive setting

Graph update:
reuploading the entire graph




Problems:

The communication/computation cost is significant when re-uploading the updated graph to update obliviously.

Research Questions

1. How to enable secure GNN inference for the *transductive* and *inductive* settings?
2. How to offer data *obliviousness* with semi-honest security?
3. How to *achieve* efficiency while achieving the above?

OblivGNN Architecture



- Semi-honest
- Outsource
- Non-colluding

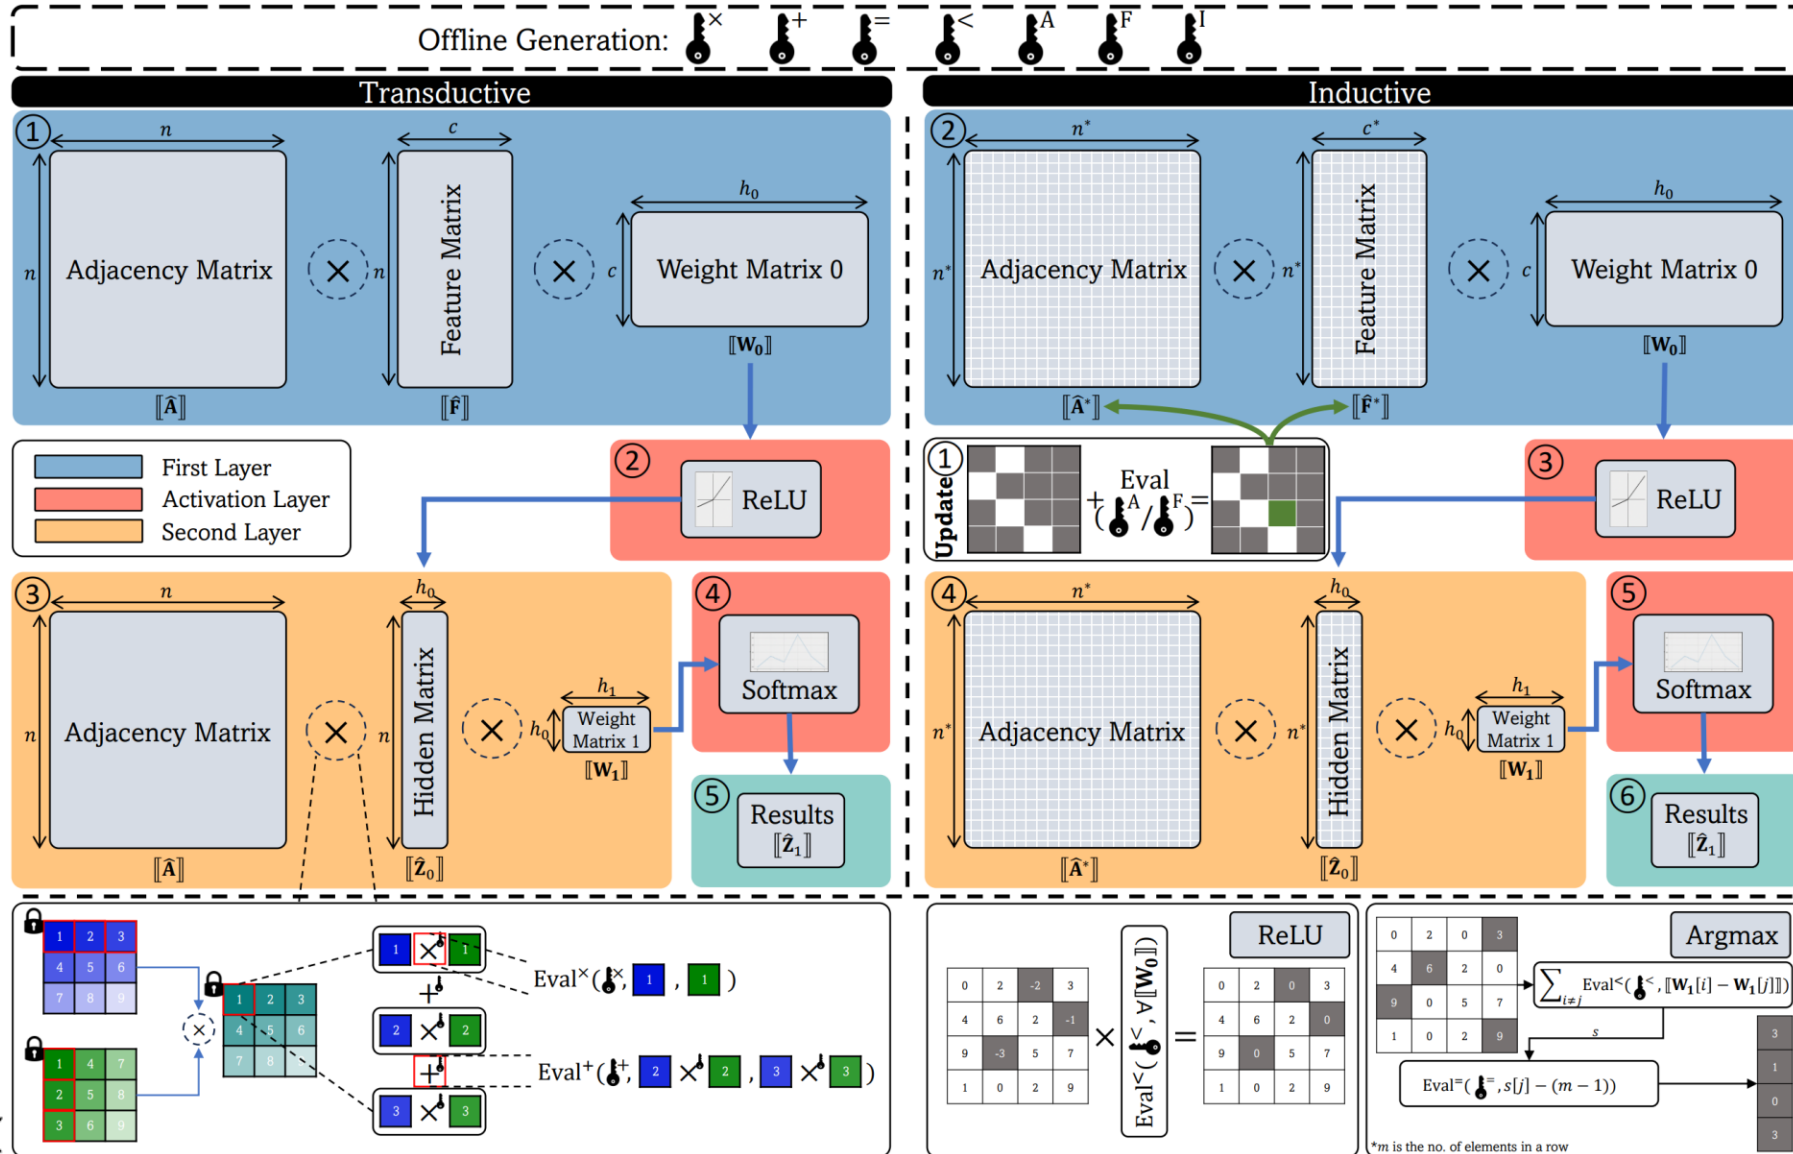
Security Guarantee

- Protect graph information
 - Adjacency Matrix $\hat{\mathbf{A}}$
 - Feature Matrix $\hat{\mathbf{F}}$
- Protect model information
 - Weight Matrix \mathbf{W}_0 and \mathbf{W}_1
- Protect access pattern to the graph structure and node feature
- Protect client query, intermediate results, and inference results

OblivGNN Approach

Offline

Online



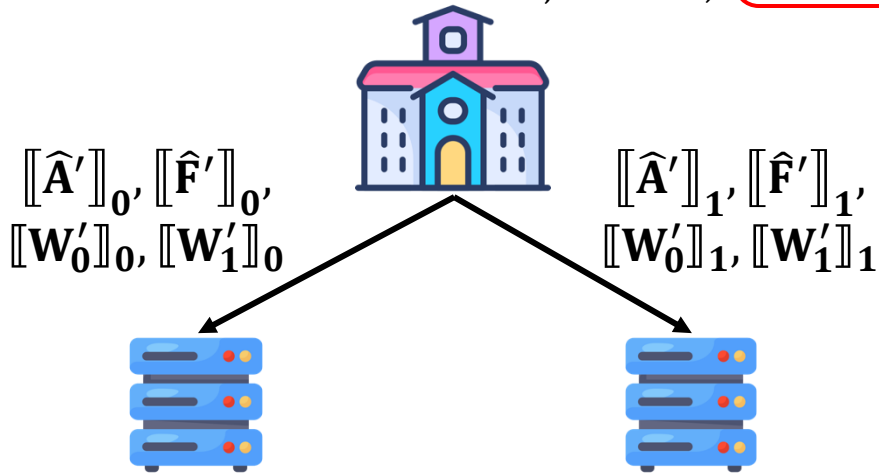
OblivGNN Approach

Offline

- Masking & secret share GNN model

Adjacency matrix: $\hat{\mathbf{A}}' \leftarrow \hat{\mathbf{A}} + r_{in}^1/r_{in}^2$
Feature matrix: $\hat{\mathbf{F}}' \leftarrow \hat{\mathbf{F}} + r_{in}^1/r_{in}^2$
Weight matrices: $\mathbf{W}'_{0,1} \leftarrow \mathbf{W}_{0,1} + r_{in}^1/r_{in}^2$

Masks for Arithmetic
FSS gates



- Two servers need to *recover* the ASS shares before operating FSS circuits
- The linear layer results (in *shares*) will be used in oblivious activation functions
- The oblivious updates are performed on the shares

OblivGNN Approach

Offline

- Key generation

FSS Key Pool Generation

Multiplication:

$\text{KeyGen}^\times(g^\circ, r_{in}^1, r_{in}^2, r_{out}) \rightarrow k_0^\times, k_1^\times : \text{FSS Multiplication keys}$

$\text{Eval}^\times(k_b^\times, x'_1, x'_2) \rightarrow g_b^\circ(x_1 \times x_2) + r_{out}$

Addition:

$\text{KeyGen}^+(g^\circ, r_{in}^1, r_{in}^2, r_{out}) \rightarrow k_0^+, k_1^+ : \text{FSS Addition keys}$

$\text{Eval}^+(k_b^+, x'_1, x'_2) \rightarrow g_b^\circ(x_1 + x_2) + r_{out}$

DPF Key Pool Generation

k^A : DPF Node Update keys

k^F : DPF Feature Update keys

k^I : DPF Client Inquiry keys

$k^=$: DPF Equality Test keys

$k^<$: DPF Comparison keys

Online keys

OblivGNN Approach

Online – Oblivious Aggregation


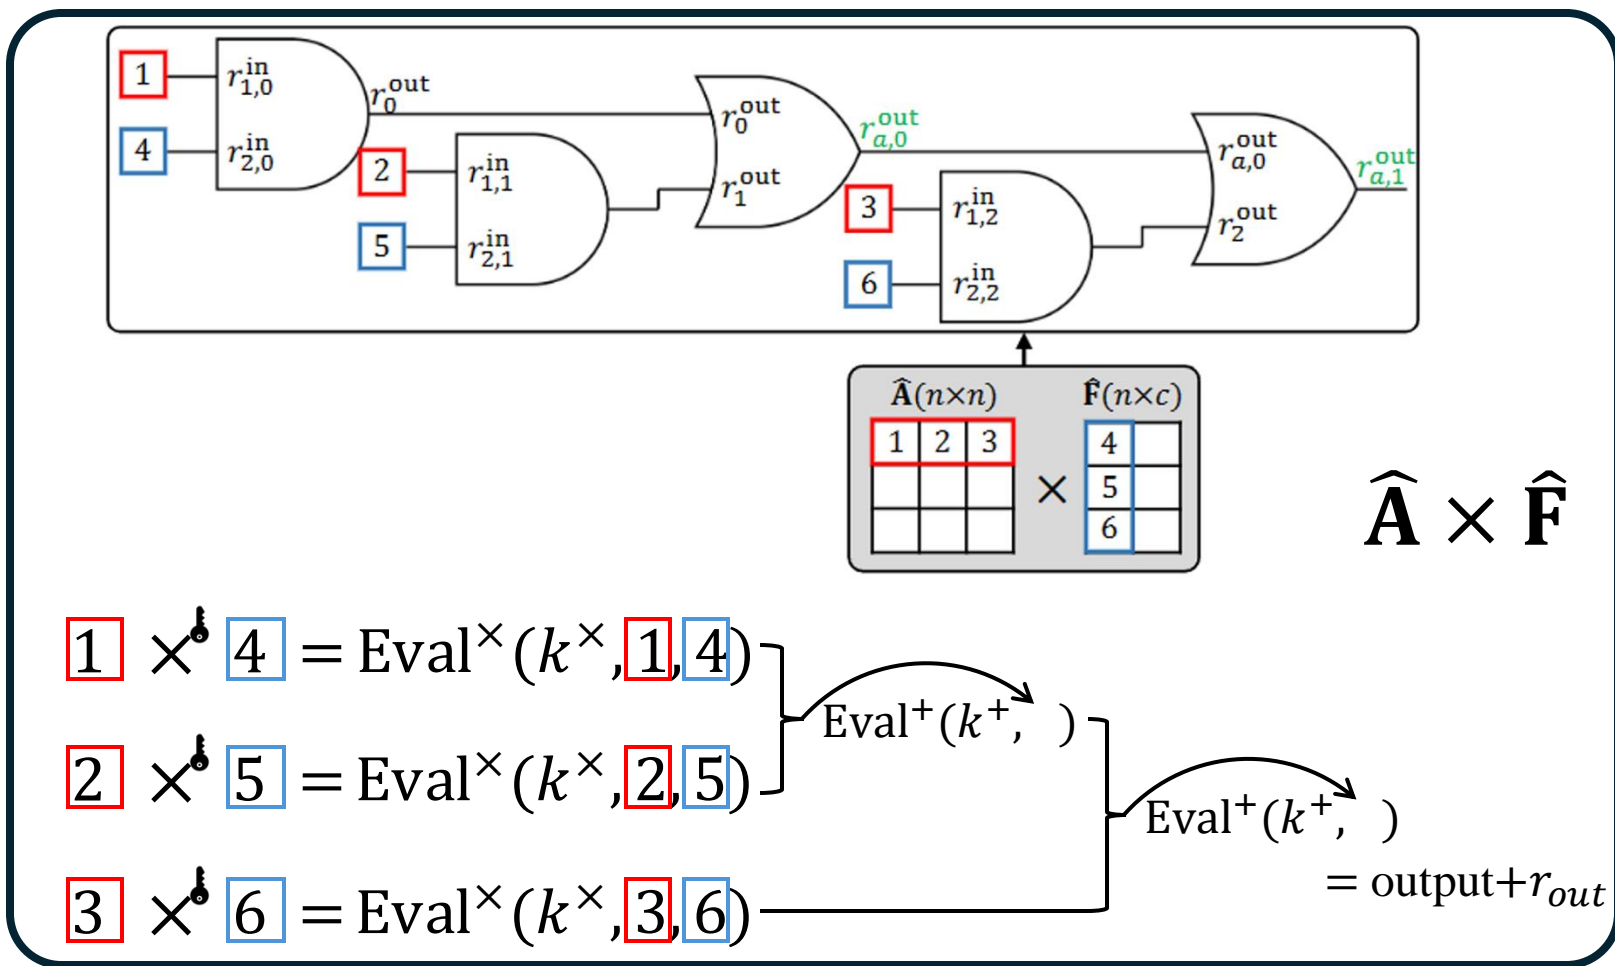
$$\text{Eval}^\times(k_0^\times, x'_1, x'_2) + \text{Eval}^\times(k_1^\times, x'_1, x'_2) = x_1 \times x_2 + r_{out}$$

$$\text{Eval}^+(k_0^+, x'_1, x'_2) + \text{Eval}^+(k_1^+, x'_1, x'_2) = x_1 + x_2 + r_{out}$$

Example:

$$x'_1 = x_1 + r_{in}^1$$

$$x'_2 = x_2 + r_{in}^2$$



OblivGNN Approach – *Inductive* Exclusive Ops

Online – Oblivious Graph Update

New Node Insertion

- Introduce *new* nodes
- Do NOT modify the existing graph

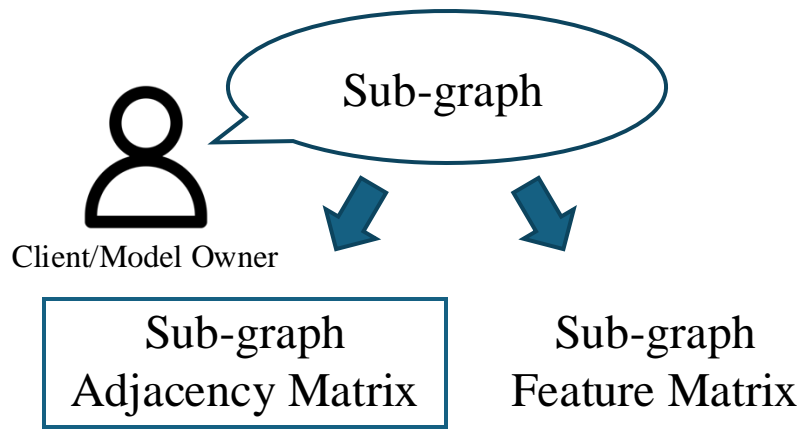
Existing Graph Update

- Modify the *existing* graph
 - Oblivious update adjacency matrix
 - Oblivious update feature matrix

OblivGNN Approach – *Inductive* Exclusive Ops

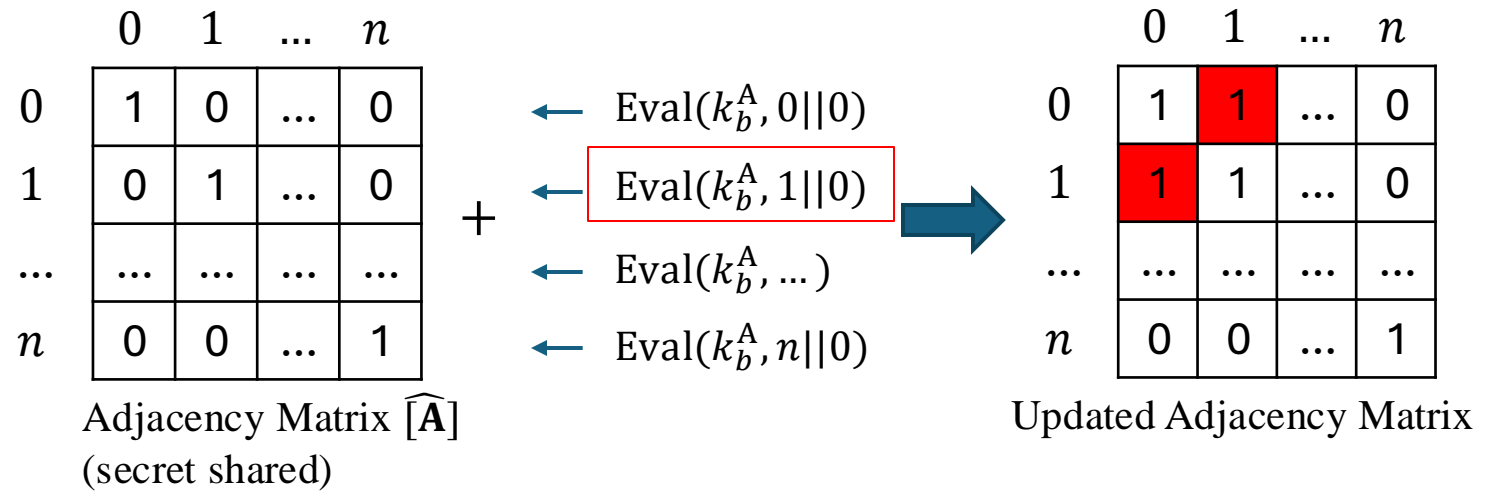
Online – Oblivious Graph Update

Existing Graph Update



KeyGen(0||1,1) $\rightarrow k_0^A, k_1^A$

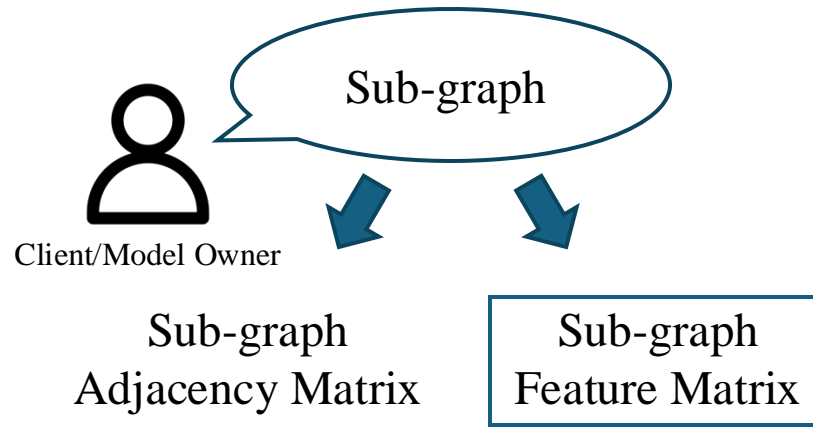
KeyGen(1||0,1) $\rightarrow k_0^A, k_1^A$



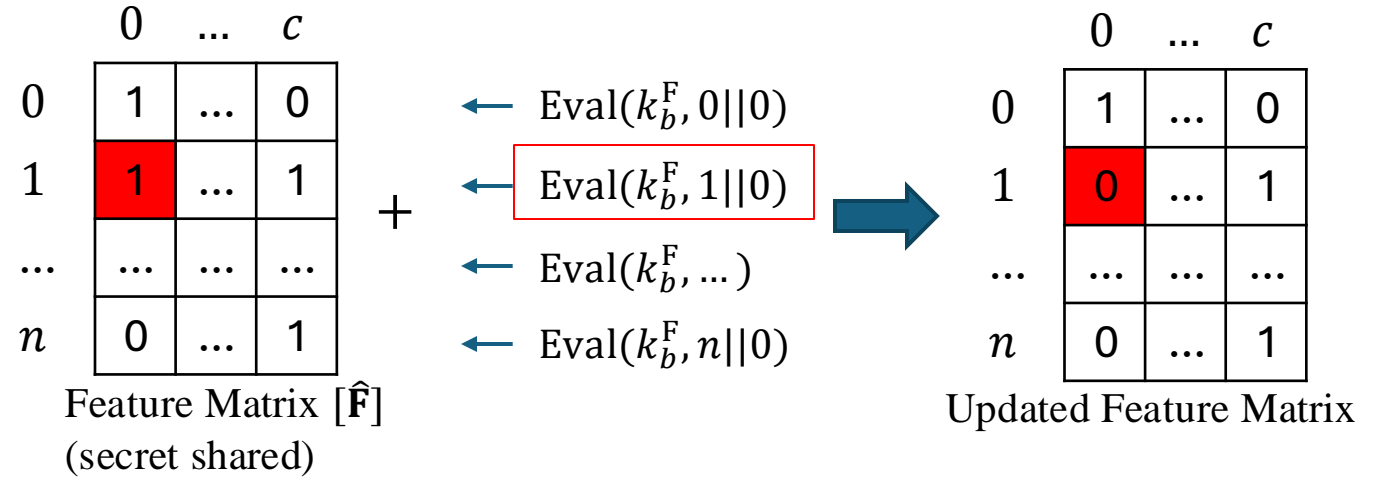
OblivGNN Approach – *Inductive Exclusive Ops*

Online – Oblivious Graph Update

Existing Graph Update



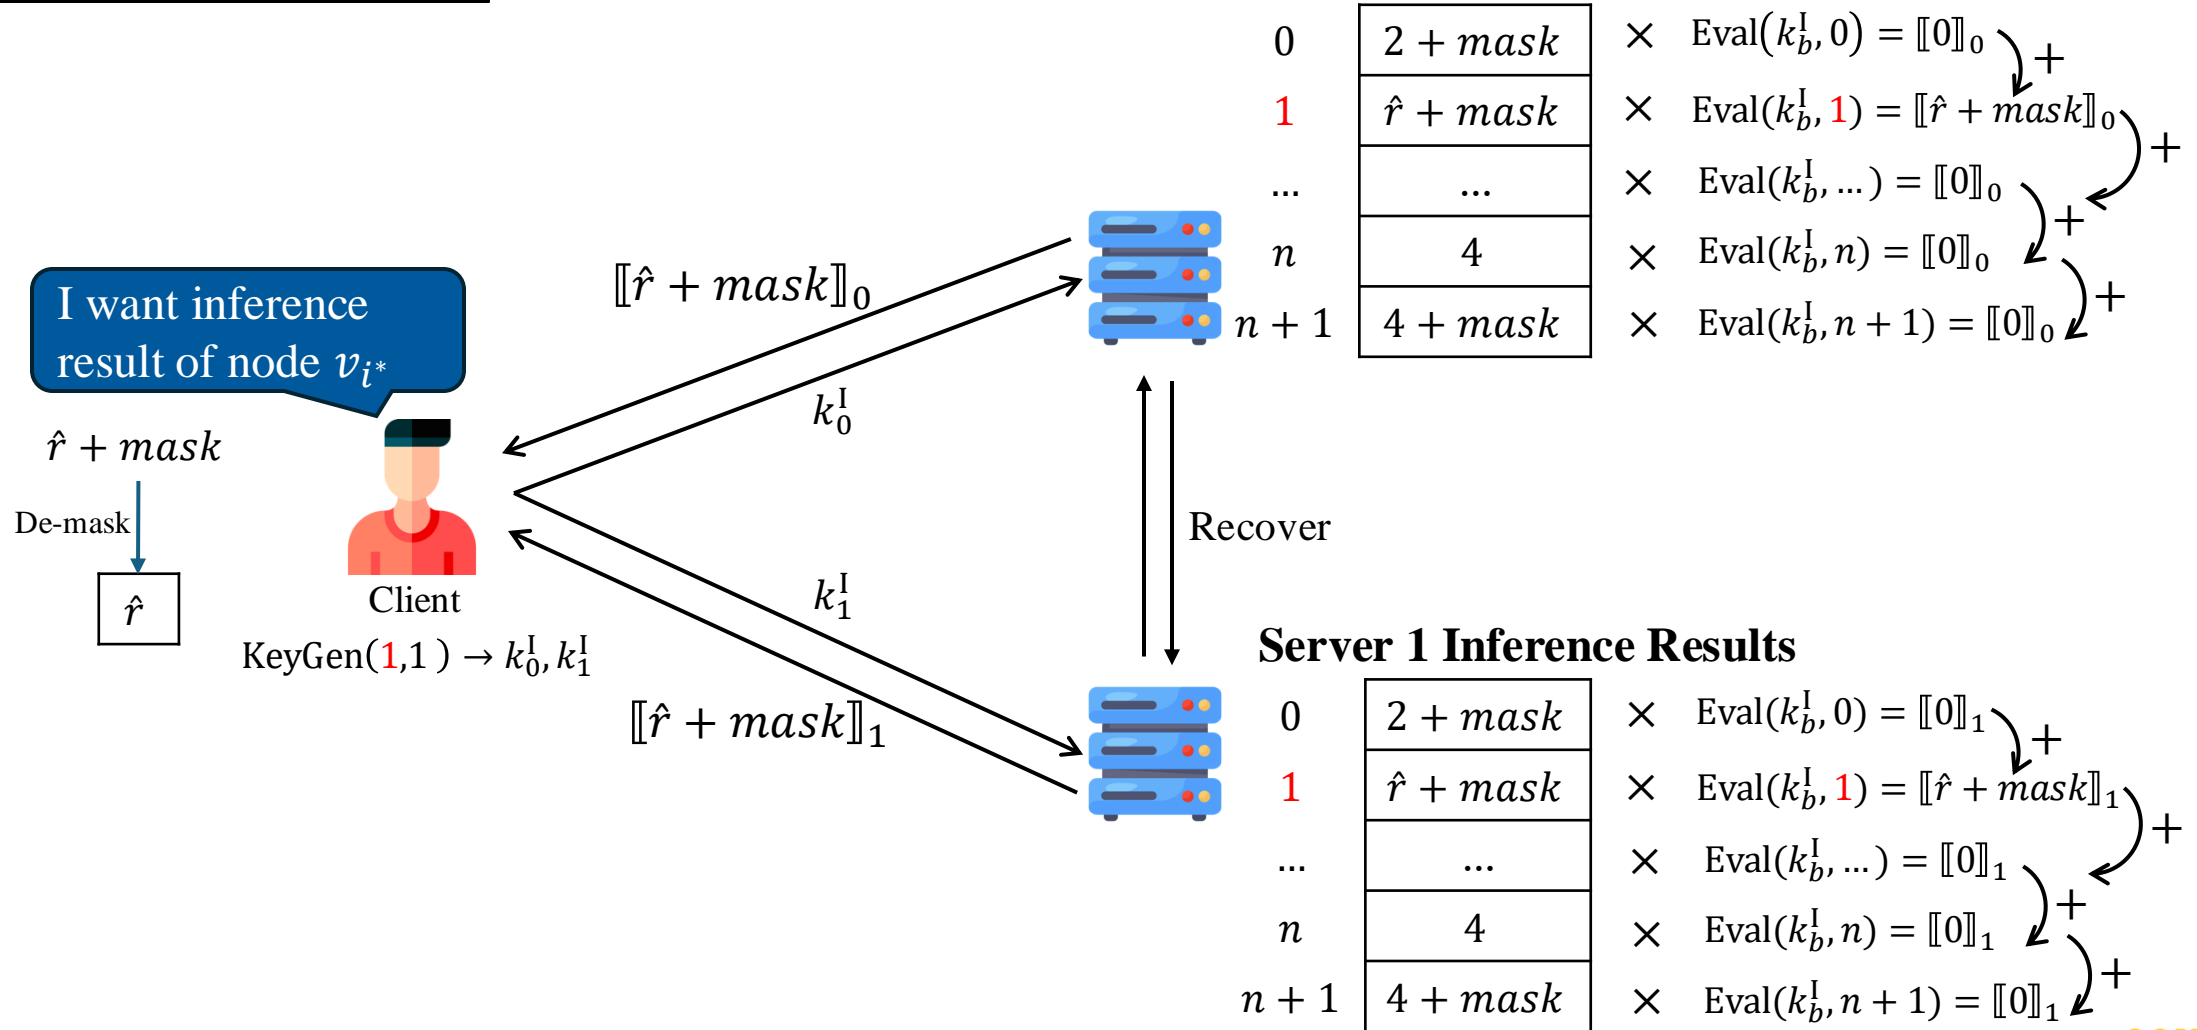
KeyGen($1||0, F_{\Delta}(v_i)$) $\rightarrow k_0^F, k_1^F$



By employing DPF keys to perform graph updates, we now achieve *oblivious* and *efficient* graph updates

OblivGNN Approach

Online – Client Inquiry



Experiments

- **Platform**

- Python and C++
- Server
 - 3.70GHz Intel(R) Xeon(R) E-2288G CPU
 - 64GB RAM and 128GB external storage
 - Ubuntu 20.04.5 LTS
- MP-SPDZ [Keller et al. (CCS'20)]

- **Datasets**

- Cora, Citeseer and Pubmed

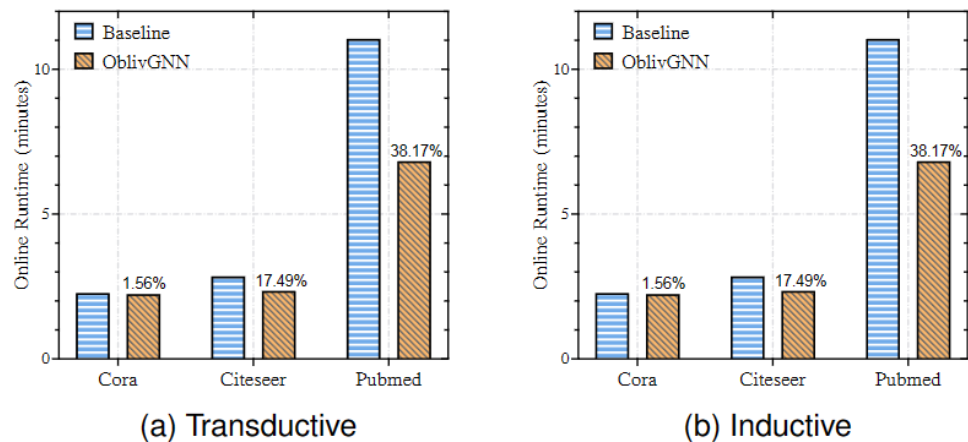
Dataset	Nodes	Feature	Edge	Classes
Cora	2708	1433	5429	7
Citeseer	3327	3703	4732	6
Pubmed	19717	500	44338	3

- **Baseline**

- Baseline: pure additive secret shares for inference.
- OblivGNN: additive secret shares with FSS for oblivious inference.

Experiments – System


System Runtime: Average Reduction: **38%**



System Communication (GB): Reduction: **10× - 151×**

	Baseline	OblivGNN
Cora	34.21	0.29
Citeseer	61.81	0.41
Pubmed	16.33	1.65

Graph Update Cost: Logarithm growth towards large graph



Special thanks to my supervisors, Prof Xingliang Yuan and Dr Shangqi Lai, for their continuous support!
Many thanks to my collaborators for their invaluable contributions and comments!



Q&A

Email: Zhibo.Xu@monash.edu

