Deep Learning on Graphs with Differential Privacy

Sina Sajadmanesh
Idiap Research Institute
Swiss Federal Institute of Technology Lausanne (EPFL)

Joint work with Ali Shahin Shamsabadi, Aurélien Bellet, and Daniel Gatica-Perez

Sony AI PPML Job Talk, May 2023

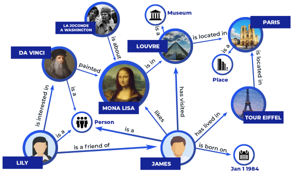

Knowledge Graphs

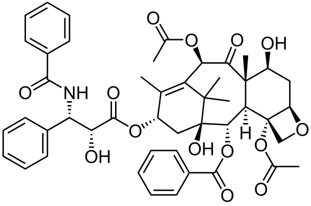

Molecules

Social Networks

Image source (from left to right): https://yashuseth.blog/2019/10/08/introduction-question-answering-knowledge-graphs-kgqa/, https://en.wikipedia.org/wiki/Terpenoid, https://yashuseth.blog/2019/10/08/introduction-question-answering-knowledge-graphs-kgqa/

Graph Neural Networks

- We need to learn representation of nodes in a low-dimensional space
  - Similar nodes in the graph should be mapped close to each other in the embedding space

- **Graph Neural Networks** (GNNs) are **state-of-the-art** representation learning algorithms for graphs.

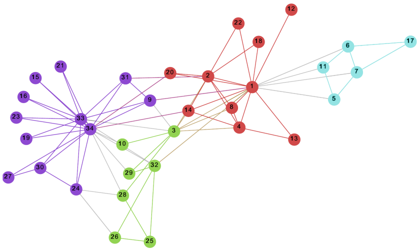

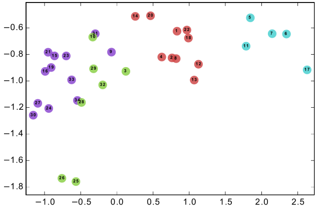

Image source: [Perozzi et al., 2014]

Privacy Concerns

- Graph data could be **privacy-sensitive**
  - e.g., social relationships, financial transactions, medical/biological networks, . . .

Privacy Concerns

- Graph data could be privacy-sensitive
  - e.g., social relationships, financial transactions, medical/biological networks, . . .
- Graph-based ML algorithms are vulnerable to privacy attacks
  - e.g., **link stealing attack** [He et al., 2021] or **membership inference attack** [Olatunji et al., 2021]

Privacy Concerns

- Graph data could be privacy-sensitive
  - e.g., social relationships, financial transactions, medical/biological networks, . . .
- Graph-based ML algorithms are vulnerable to privacy attacks
  - e.g., **link stealing attack** [He et al., 2021] or **membership inference attack** [Olatunji et al., 2021]

We need privacy-preserving machine learning algorithms for graph data!

Differential Privacy [Dwork et al., 2006]

Randomized algorithm A is ϵ-DP if for all neighboring graphs $G \simeq G'$ and all sets of outputs S:

$$\Pr[A(G) \in S] \leq e^{\epsilon} \Pr[A(G') \in S]$$

- **Edge-Level DP**
  Neighboring graphs differ by at most one edge
- **Node-Level DP**
  Neighboring graphs differ by at most one node (and all adjacent edges)

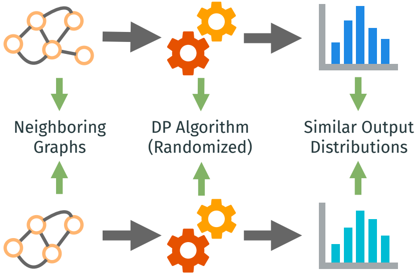

$\mathbf{A}$: Adjacency matrix

$\mathbf{X}$: Input node features

$\mathbf{Y}$: Predicted node labels

$\mathbf{H}^{(i)}$: Hidden node representations of layer i

Agg: Aggregation function

- e.g., summation: $\textsc{Agg}(\mathbf{H}, \mathbf{A}) = \mathbf{A}^T \cdot \mathbf{H}$

Upd: Learnable update function

- e.g., an MLP

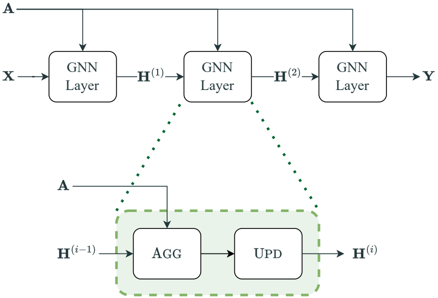

DP GNN Challenges: Why Not DP-SGD?

- Exploding Sensitivity
  - With a K-layer GNN, each node affects the embedding of all the nodes in its K-hop neighborhood
  - $O(D^K)$ gradient terms change at once (D is maximum degree)

DP GNN Challenges: Why Not DP-SGD?

- Exploding Sensitivity
  - With a K-layer GNN, each node affects the embedding of all the nodes in its K-hop neighborhood
  - $O(D^K)$ gradient terms change at once (D is maximum degree)
- Inference Privacy
  - GNNs query the graph structure during inference
  - Private information leaks at inference, even with a private model

DP GNN Challenges: Why Not DP-SGD?

- Exploding Sensitivity
  - With a K-layer GNN, each node affects the embedding of all the nodes in its K-hop neighborhood
  - $O(D^K)$ gradient terms change at once (D is maximum degree)
- Inference Privacy
  - GNNs query the graph structure during inference
  - Private information leaks at inference, even with a private model

DP-SGD cannot be directly applied to GNNs

Our Approach: Aggregation Perturbation

- **Aggregation Perturbation:** adding noise to output of the aggregation step
    - Prevents the exploding sensitivity problem by composing differentially private aggregation steps
    - Ensures inference privacy

Our Approach: Aggregation Perturbation

- **Aggregation Perturbation:** adding noise to output of the aggregation step
  - Prevents the exploding sensitivity problem by composing differentially private aggregation steps
  - Ensures inference privacy
- Applying aggregation perturbation to the conventional GNNs is **costly**
  - Every forward pass of the model consumes privacy budget
  - The excessive noise results in poor performance

Our Approach: Aggregation Perturbation

- **Aggregation Perturbation:** adding noise to output of the aggregation step
  - Prevents the exploding sensitivity problem by composing differentially private aggregation steps
  - Ensures inference privacy
- Applying aggregation perturbation to the conventional GNNs is **costly**
  - Every forward pass of the model consumes privacy budget
  - The excessive noise results in poor performance

We decouple the aggregation operations from the model parameters to maintain the privacy budget

GNN with Aggregation Perturbation (GAP)

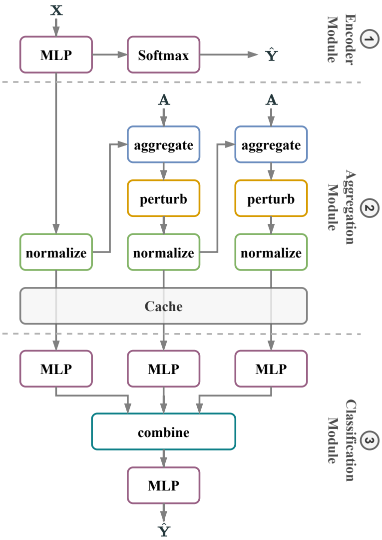

✓ Edge-level DP

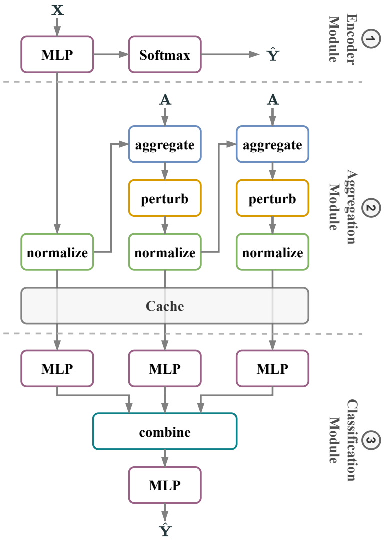

Advantages of GAP Architecture

- ✓ Edge-level DP
- ✓ Node-level DP through combination with DP-SGD
  - For bounded-degree graphs

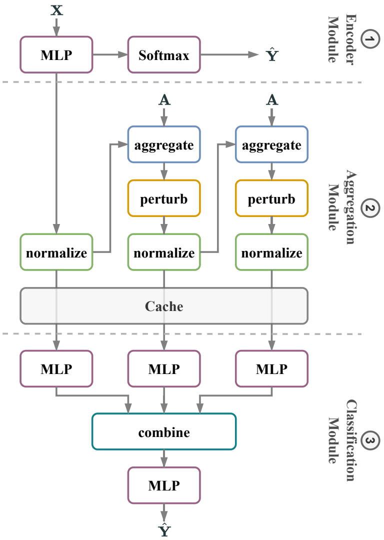

Advantages of GAP Architecture

✓ Edge-level DP

✓ Node-level DP through combination with DP-SGD

- For bounded-degree graphs

✓ Multi-hop aggregations

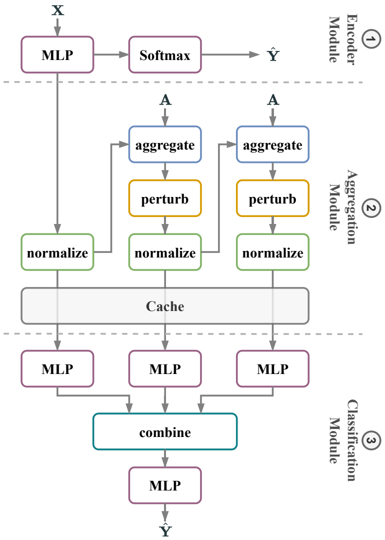

Advantages of GAP Architecture

- ✓ Edge-level DP
- ✓ Node-level DP through combination with DP-SGD
  - For bounded-degree graphs
- ✓ Multi-hop aggregations
- ✓ Zero-cost inference privacy

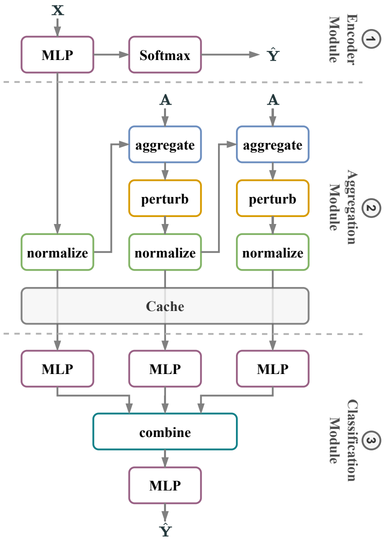

Experiment Settings

- **Task:** Node Classification
- **Baselines:** MLP, GraphSAGE

Dataset	Classes	Nodes	Edges	Features	Med. Degree
Facebook	6 Year	26,406 User	2,117,924 Friendship	501	62
Reddit	8 Community	116,713 Post	46,233,380 Mutual User	602	209
Amazon	10 Category	1,790,731 Product	80,966,832 Mutual Purchase	100	22

Accuracy of Non-Private Methods

Method	Facebook	Reddit	Amazon
GAP-∞	80.0 ± 0.48	**99.4 ± 0.02**	91.2 ± 0.07
SAGE-∞	**83.2 ± 0.68**	99.1 ± 0.01	**92.7 ± 0.09**

Edge-Level DP Accuracy-Privacy Trade-Off

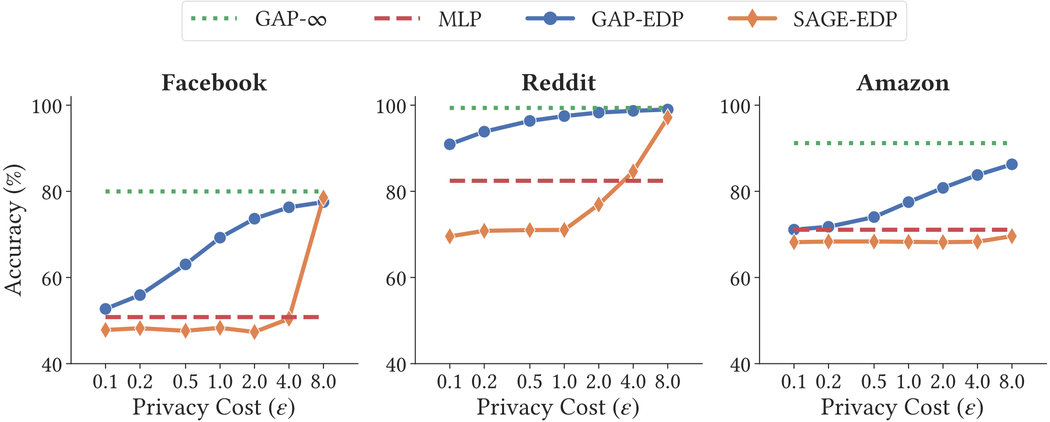

Node-Level DP Accuracy-Privacy Trade-Off

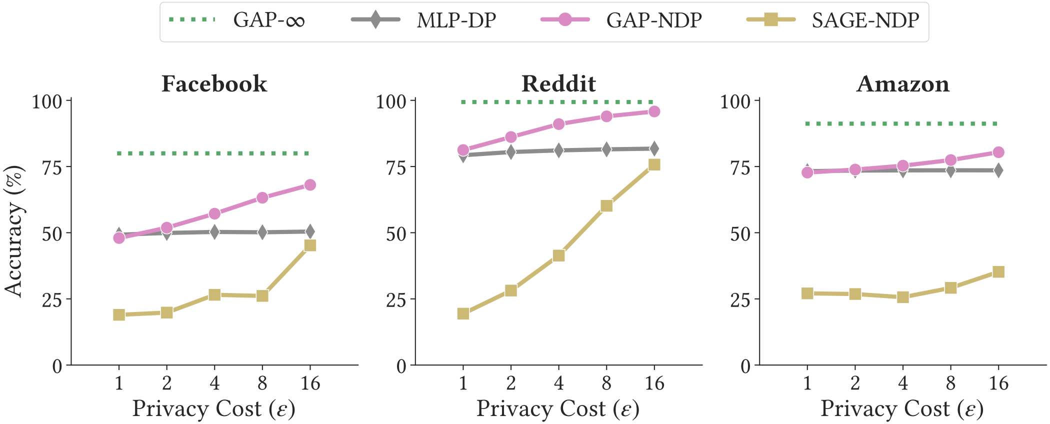

Resiliency to Membership Inference Attack

Mean AUC of node-level membership inference attack.

Dataset	Method	$\epsilon = 1$	$\epsilon = 2$	$\epsilon = 4$	$\epsilon = 8$	$\epsilon = 16$	$\epsilon = \infty$
Facebook	GAP-NDP	50.16	50.25	50.61	51.11	52.66	81.67
Reddit	GAP-NDP	50.04	50.39	51.20	52.23	52.54	54.97
Amazon	GAP-NDP	50.06	50.23	50.54	51.53	51.72	66.68

Conclusion

- GNNs leak private information
  - They are vulnerable to privacy attacks
- Implementing DP in GNNs is challenging
  - Exploding sensitivity
  - Inference privacy
- Our Differentially Private GNN: GAP
  - Ensures both edge-level and node-level DP
  - Supports multi-hop aggregations
  - Provides inference privacy

Conclusion

- GNNs leak private information
  - They are vulnerable to privacy attacks
- Implementing DP in GNNs is challenging
  - Exploding sensitivity
  - Inference privacy
- Our Differentially Private GNN: GAP
  - Ensures both edge-level and node-level DP
  - Supports multi-hop aggregations
  - Provides inference privacy

Thank You !

Questions?

References I

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006).
Calibrating noise to sensitivity in private data analysis.
In *Theory of cryptography conference*, pages 265–284. Springer.

He, X., Jia, J., Backes, M., Gong, N. Z., and Zhang, Y. (2021).
Stealing links from graph neural networks.
In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.

Olatunji, I. E., Nejdl, W., and Khosla, M. (2021).
Membership inference attack on graph neural networks.
arXiv preprint arXiv:2101.06570.

Perozzi, B., Al-Rfou, R., and Skiena, S. (2014).
Deepwalk: Online learning of social representations.
In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710.

Effect of the Number of Hops

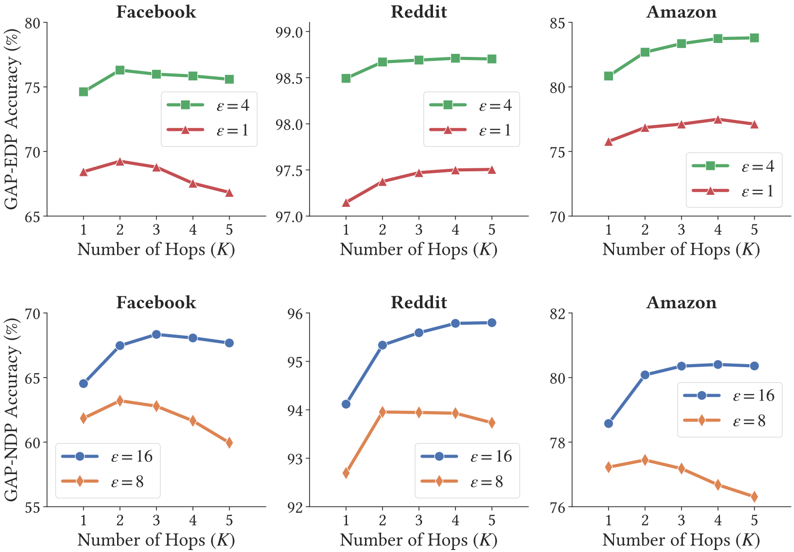

Effect of the Encoder Module

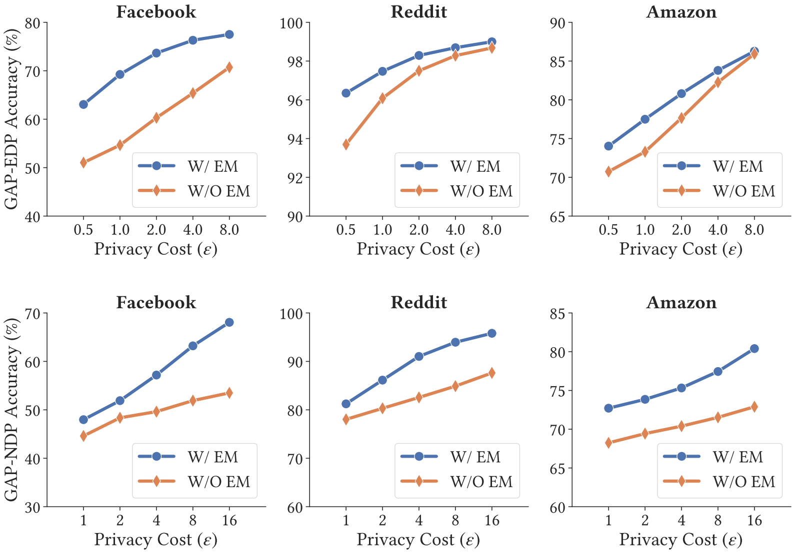

Effect of the Maximum Degree