Fairness in Graph Mining: Metrics and Algorithms

Yushun Dong
University of Virginia
Supervisor: Prof. Jundong Li

Outline

- Background Information
- Fairness Notions and Metrics
- Methodologies to Mitigate Bias
- Existing Problems & Future Works

Outline

Background Information

Fairness Notions and Metrics

Methodologies to Mitigate Bias

Existing Problems & Future Works

Background Information

- What are graph mining algorithms?

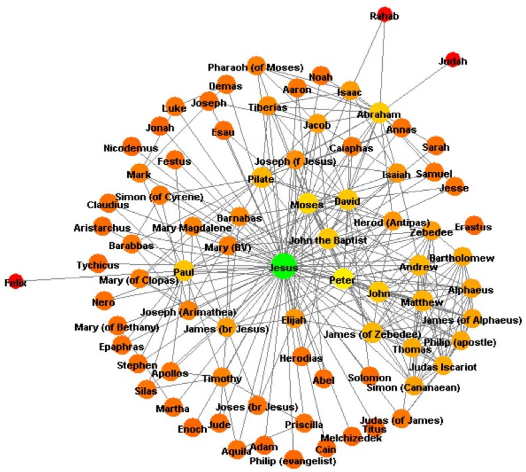

Social networks

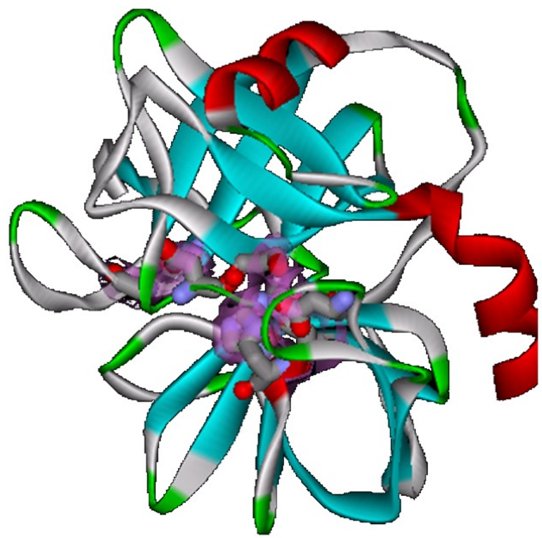

Protein Structures

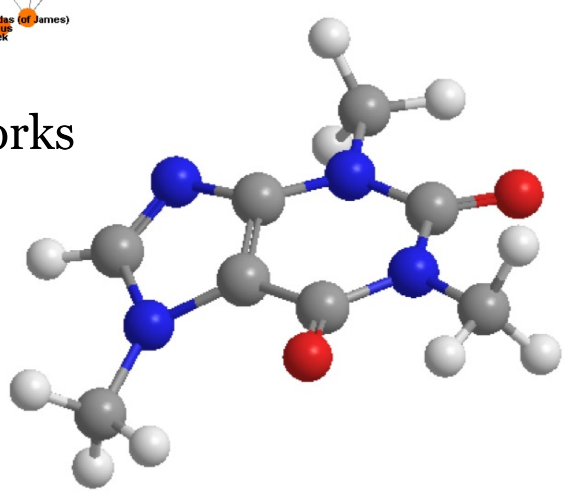

Chemical Compounds

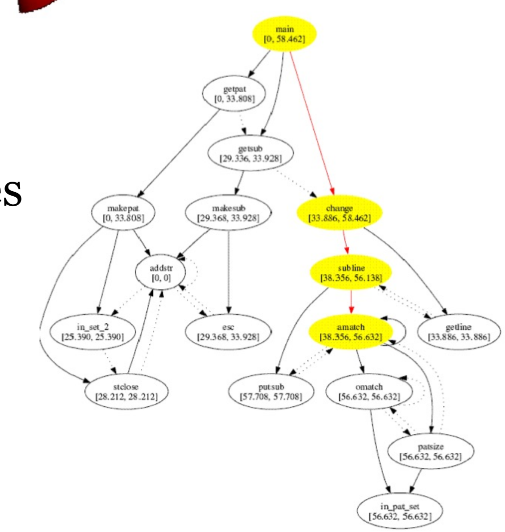

Program Flows

Background Information

- What are graph mining algorithms?

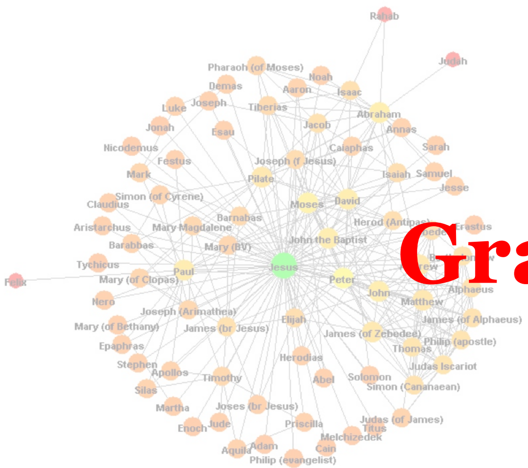

Social networks

Protein Structures

Graphs are everywhere

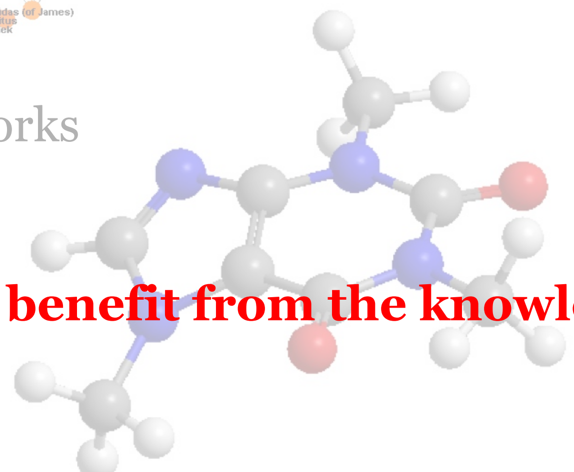

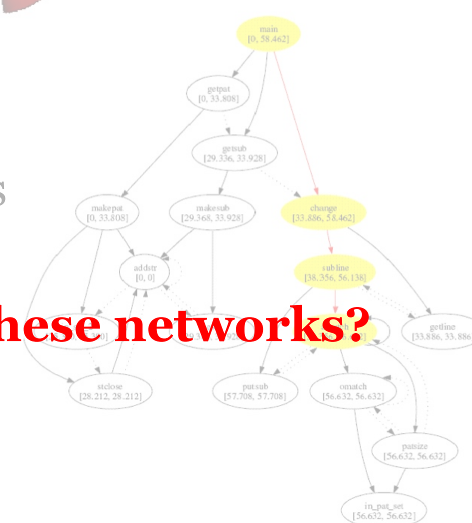

But how to benefit from the knowledge encoded in these networks?

Chemical Compounds

Program Flows

Background Information

- What are graph mining algorithms?

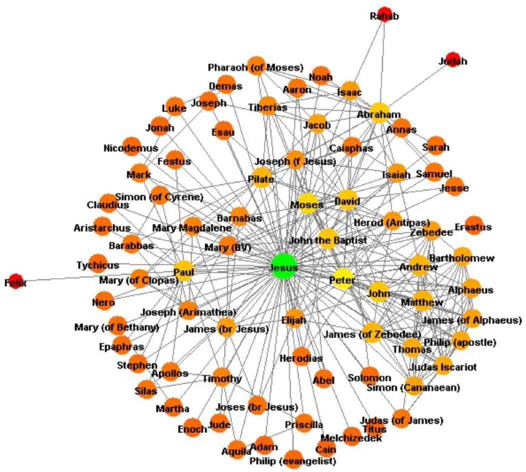

Social networks

Background Information

- What are graph mining algorithms?

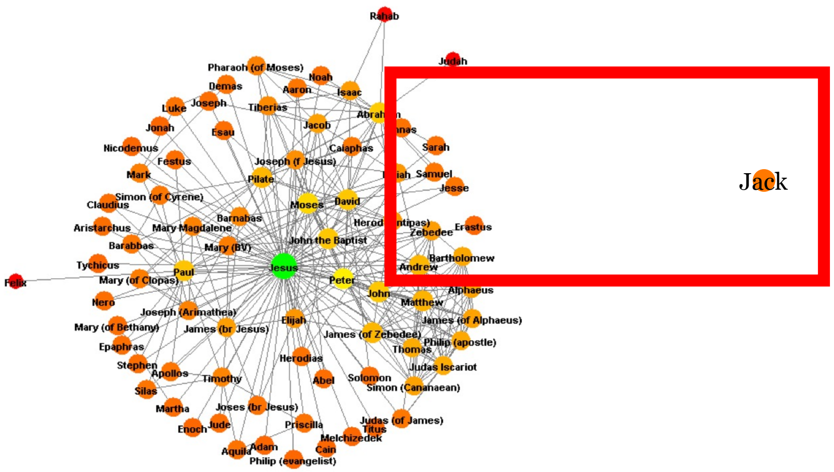

Social networks

Jack becomes a **new user** of the social networking platform.

Background Information

- What are graph mining algorithms?

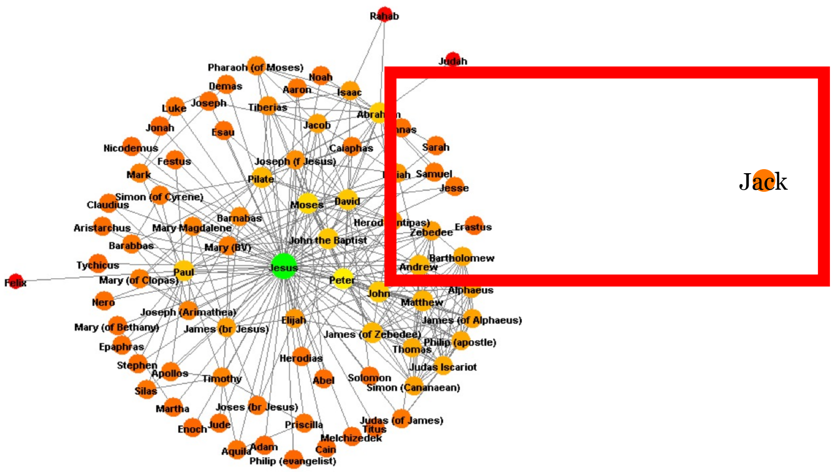

Jack becomes a **new user** of the social networking platform.

Who should be recommended to him?

Social networks

Background Information

- What are graph mining algorithms?

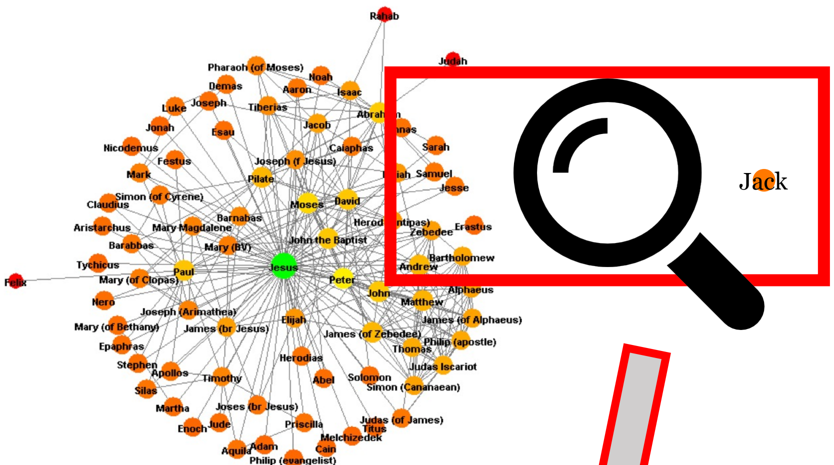

Social networks

Jack becomes a **new user** of the social networking platform.

Who should be recommended to him?

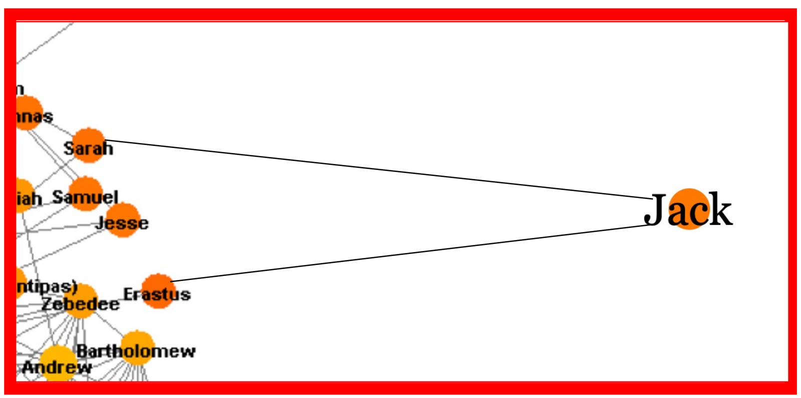

Jack chose to become **a friend** of Sarah and Erastus in the first week.

Background Information

- What are graph mining algorithms?

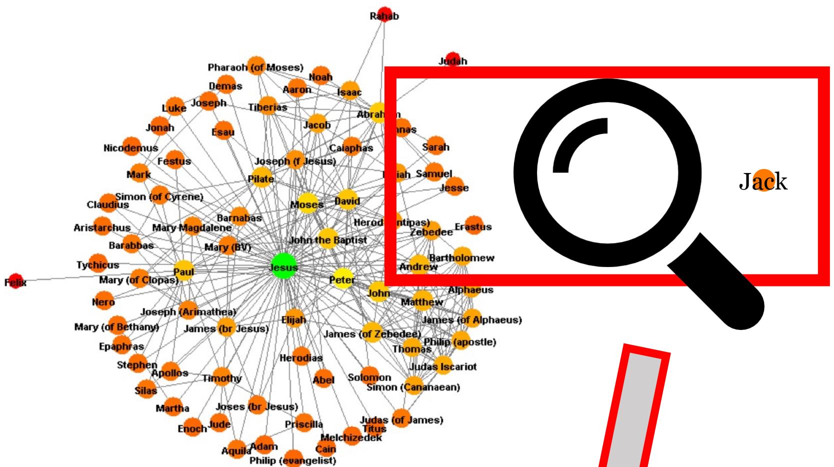

Social networks

Jack becomes a **new user** of the social networking platform.

Who should be recommended to him?

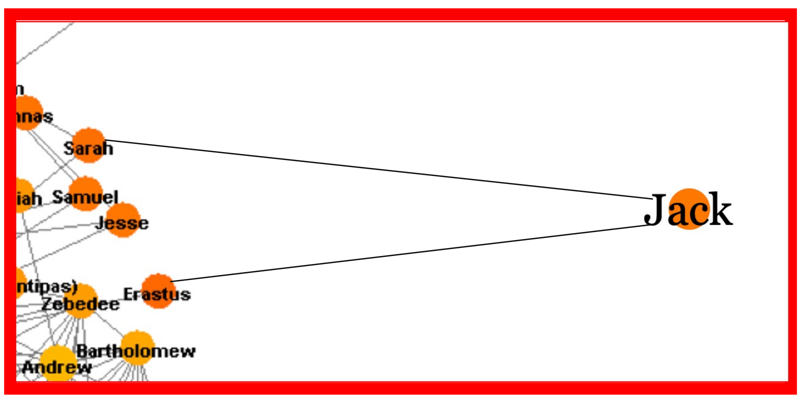

Jack chose to become **a friend** of Sarah and Erastus in the first week.

Who else should be recommended to him?

Background Information

- What are graph mining algorithms?

Graph mining algorithms are algorithms that **extracts information encoded in the graph data** to facilitate our understanding (on these graphs) and gain benefit on various predictive tasks.

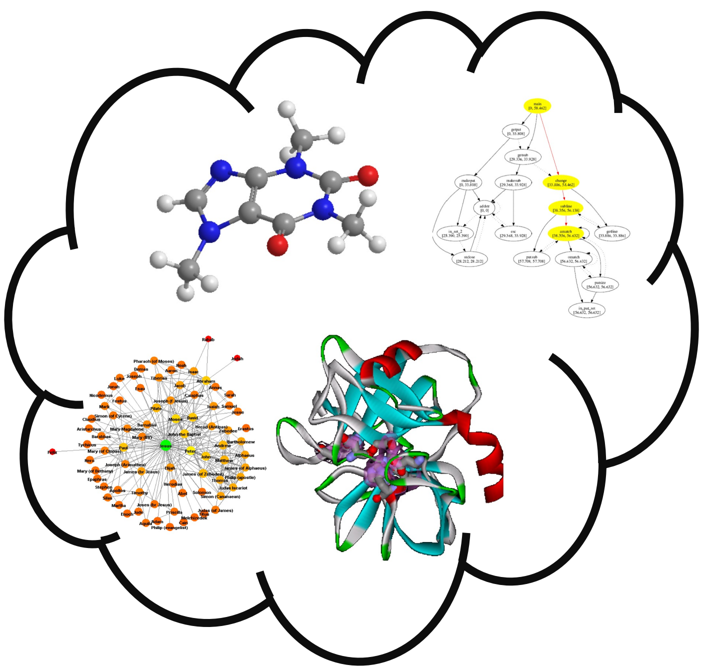

Graph Mining Algorithms

- Who are potential friends?
- What will the properties of this chemical compound be like?
- Which part of the protein accounts for its interactions with other compounds?

…

Background Information

- What is fairness?

Background Information

- What is fairness?

  Fairness is a vague notion: there is **no specific criterion** defines how to determine whether fairness has been fulfilled [1].

Background Information

- What is fairness?

Fairness is a vague notion: there is **no specific criterion** defines how to determine whether fairness has been fulfilled [1].

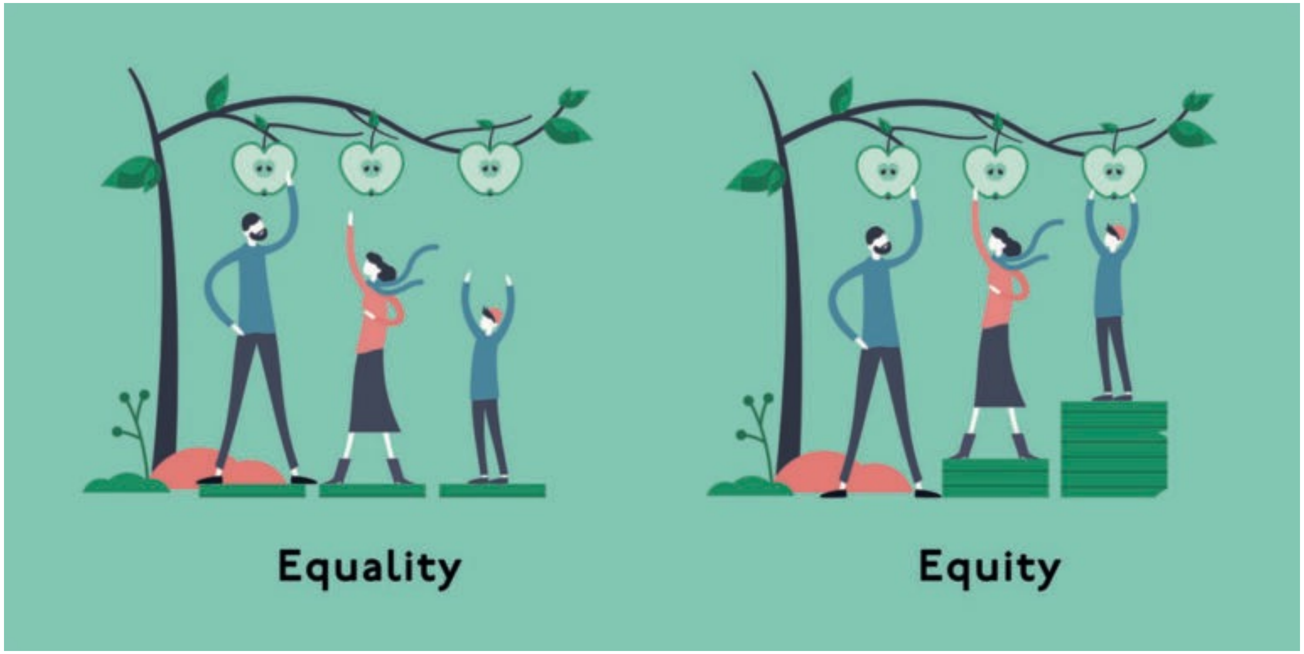

For example, it **depends on the specific application scenario** to determine whether "Equality" or "Equity" should be considered as fairness.

Background Information

- What is fairness?

Despite the lack of **specific criterion** for fairness, we could still study fairness in algorithms: there are **various existing fairness notions** based on people's awareness.

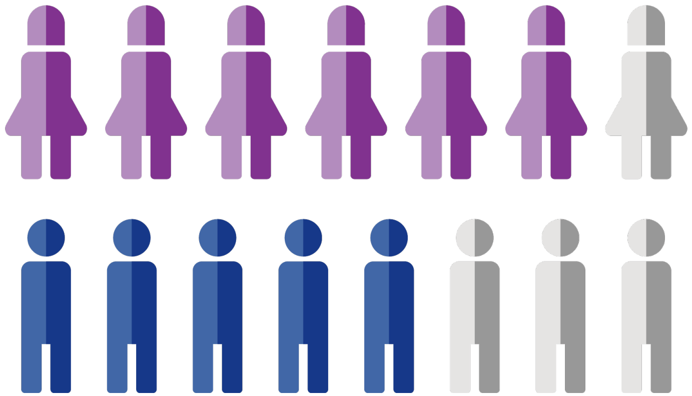

Background Information

- What is fairness?

Despite the lack of **specific criterion** for fairness, we could still study fairness in algorithms: there are **various existing fairness notions** based on people's awareness.

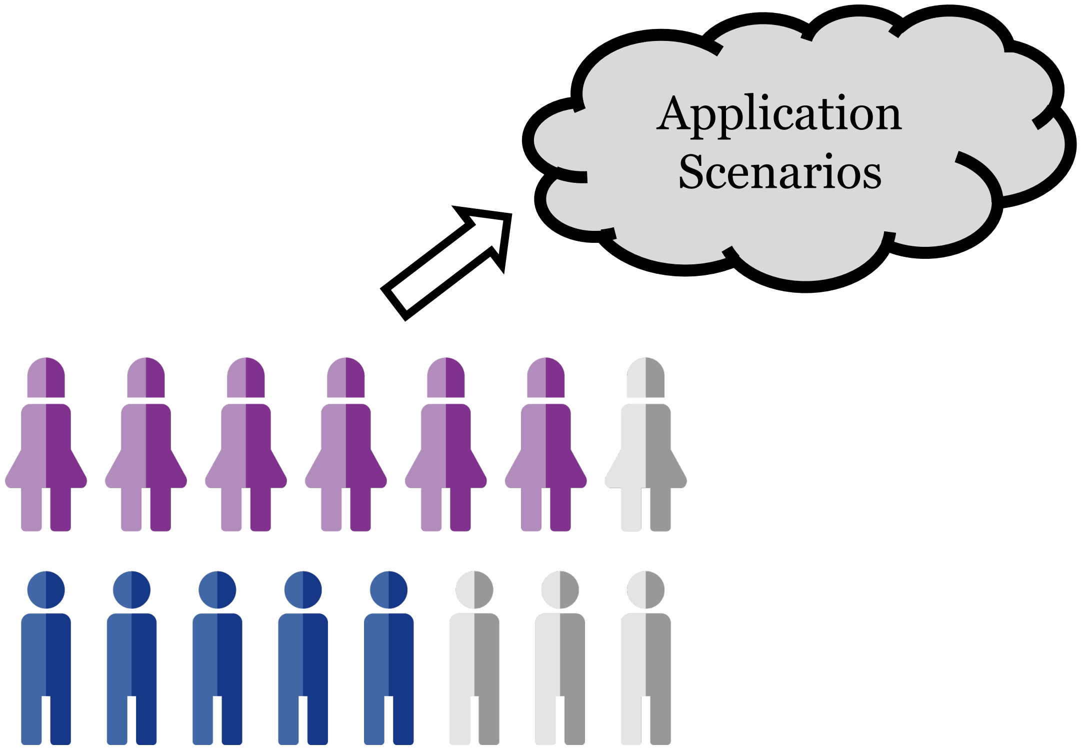

Background Information

- What is fairness?

Despite the lack of **specific criterion** for fairness, we could still study fairness in algorithms: there are **various existing fairness notions** based on people's awareness.

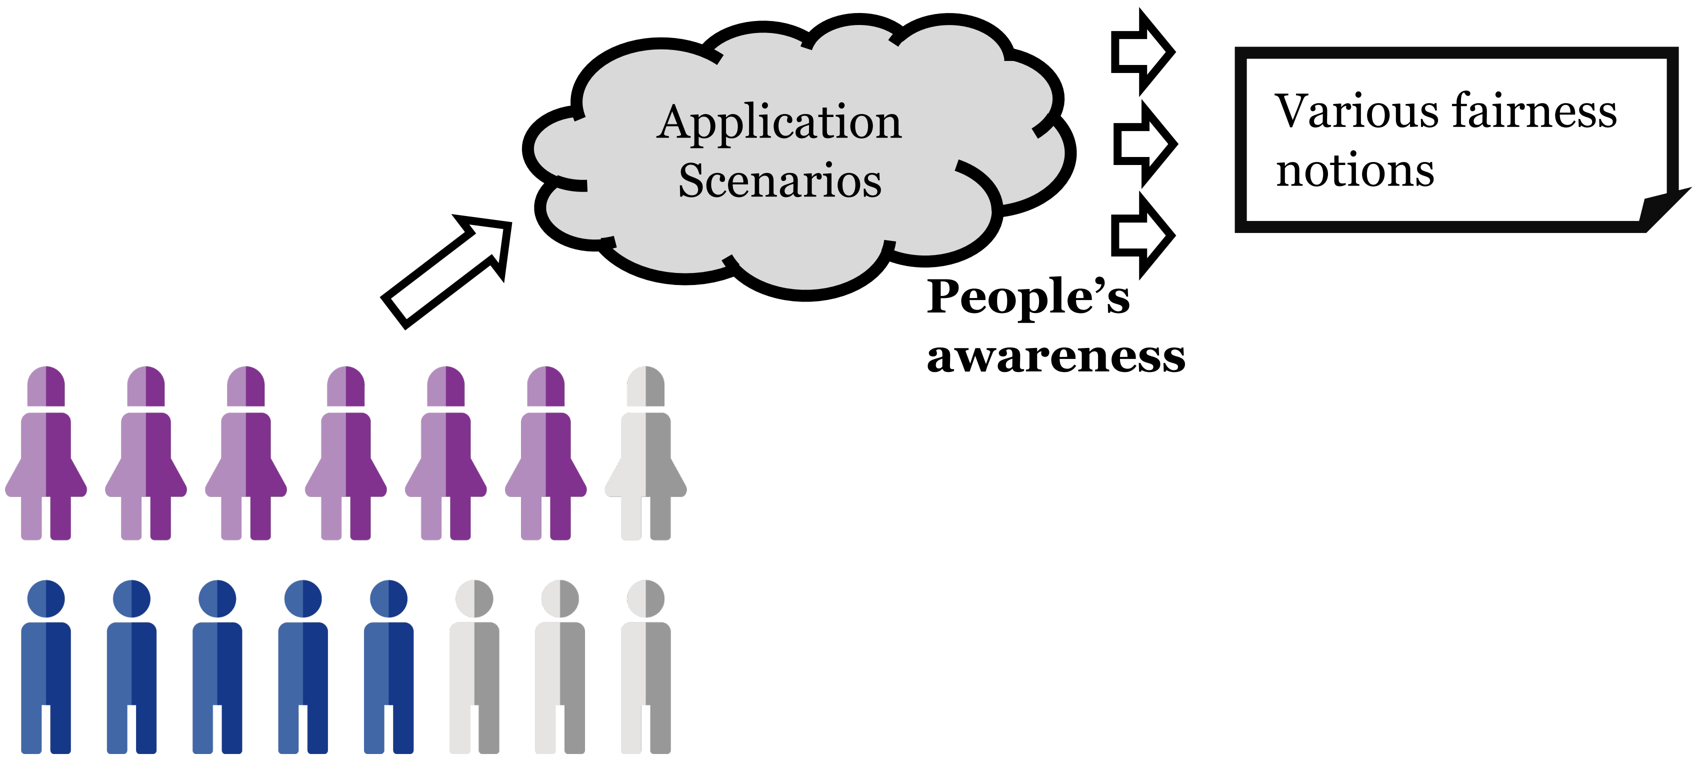

Background Information

- What is fairness?

Despite the lack of **specific criterion** for fairness, we could still study fairness in algorithms: there are **various existing fairness notions** based on people's awareness.

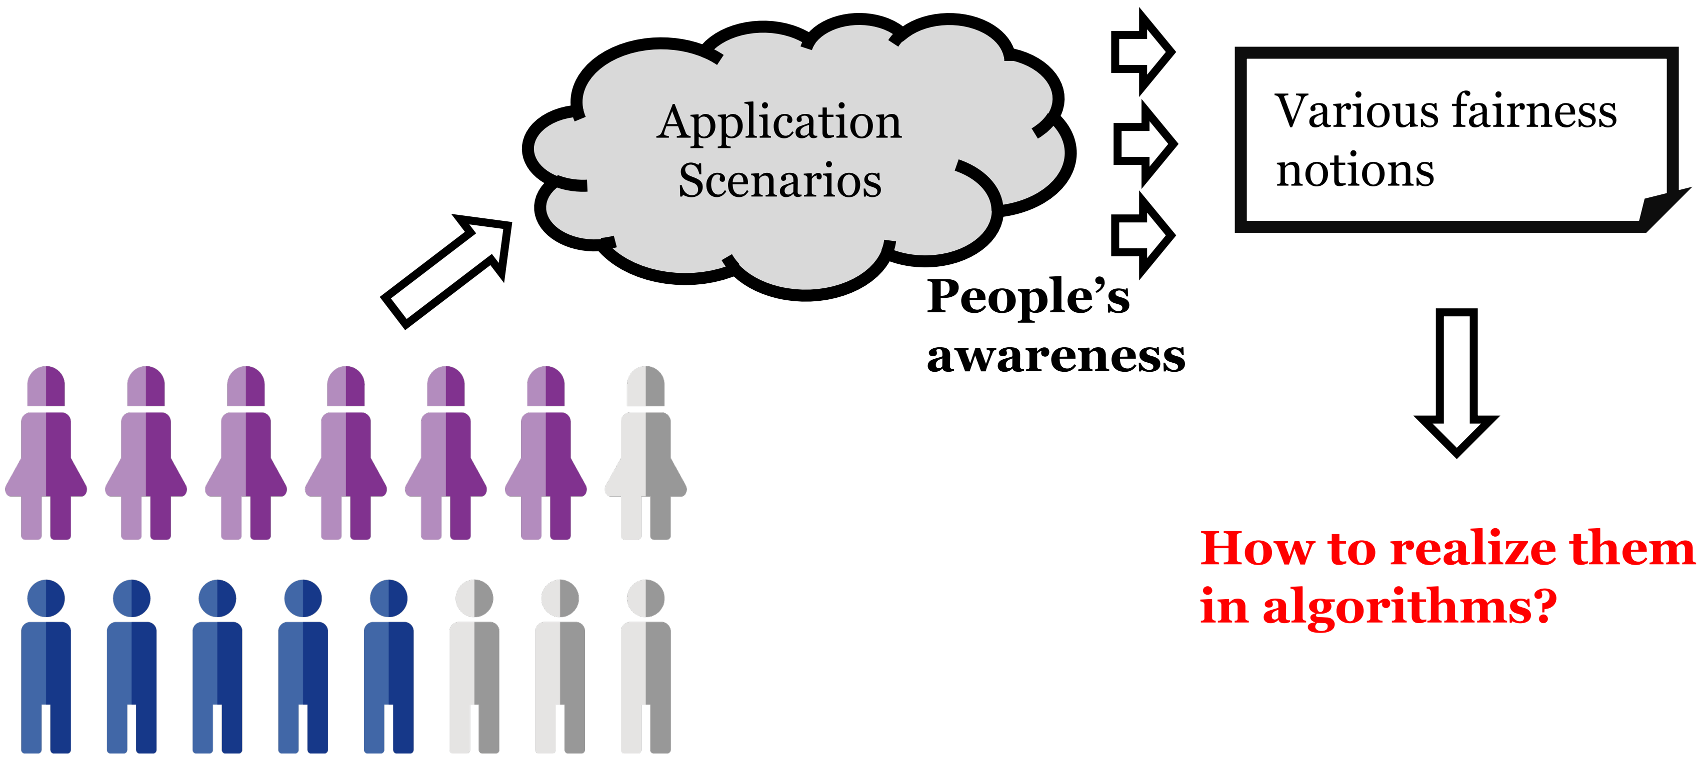

Background Information

- What is fairness?

In the realm of **graph mining…**

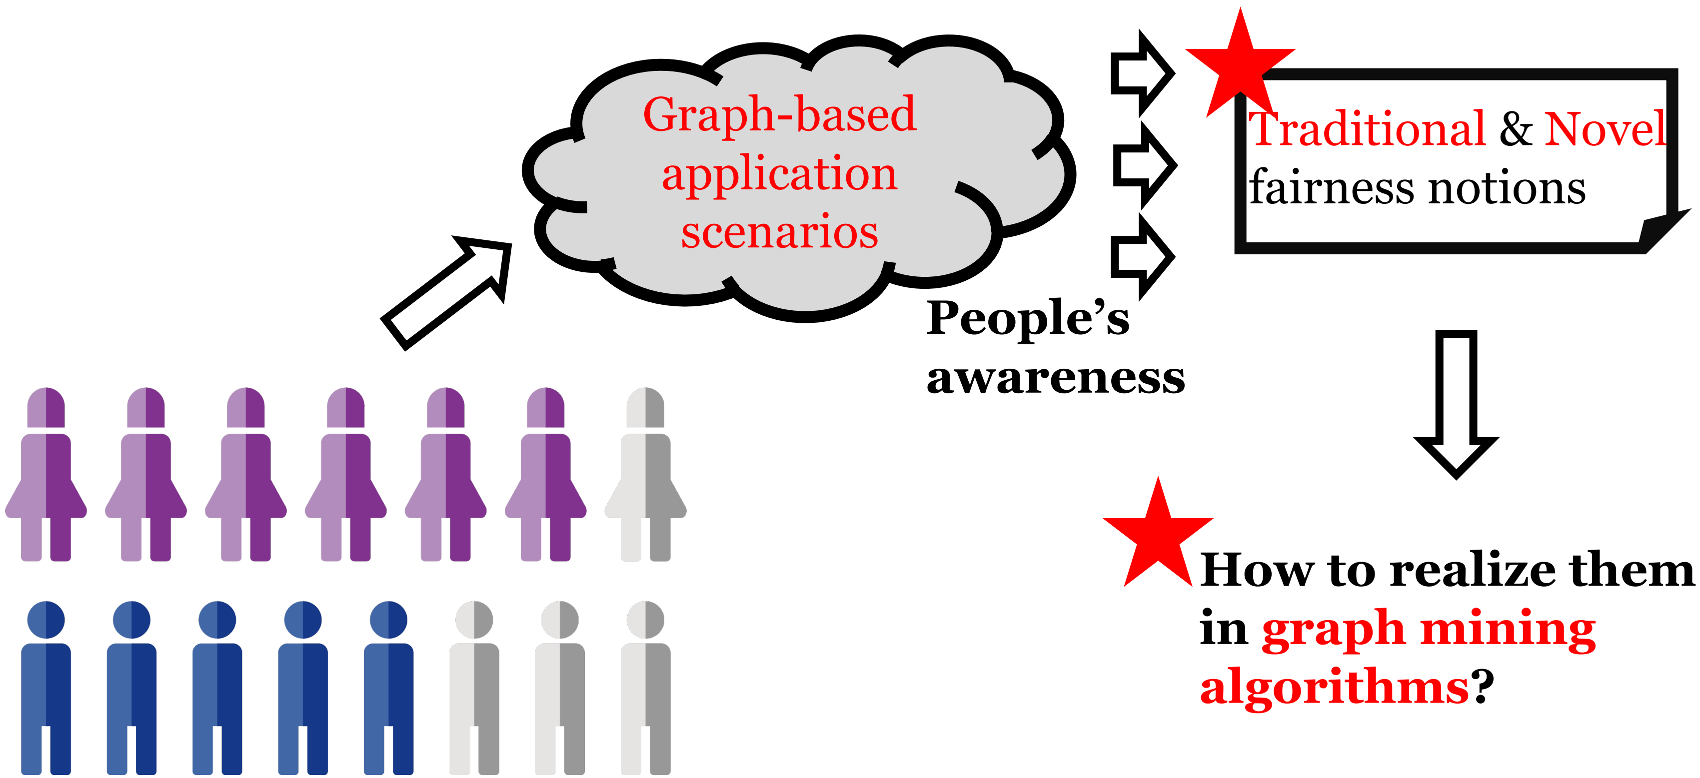

Outline

Background Information

Fairness Notions and Metrics

Methodologies to Mitigate Bias

Existing Problems & Future Works

Fairness Notions and Metrics

- A taxonomy of commonly used algorithmic fairness notions in graph mining algorithms [2].

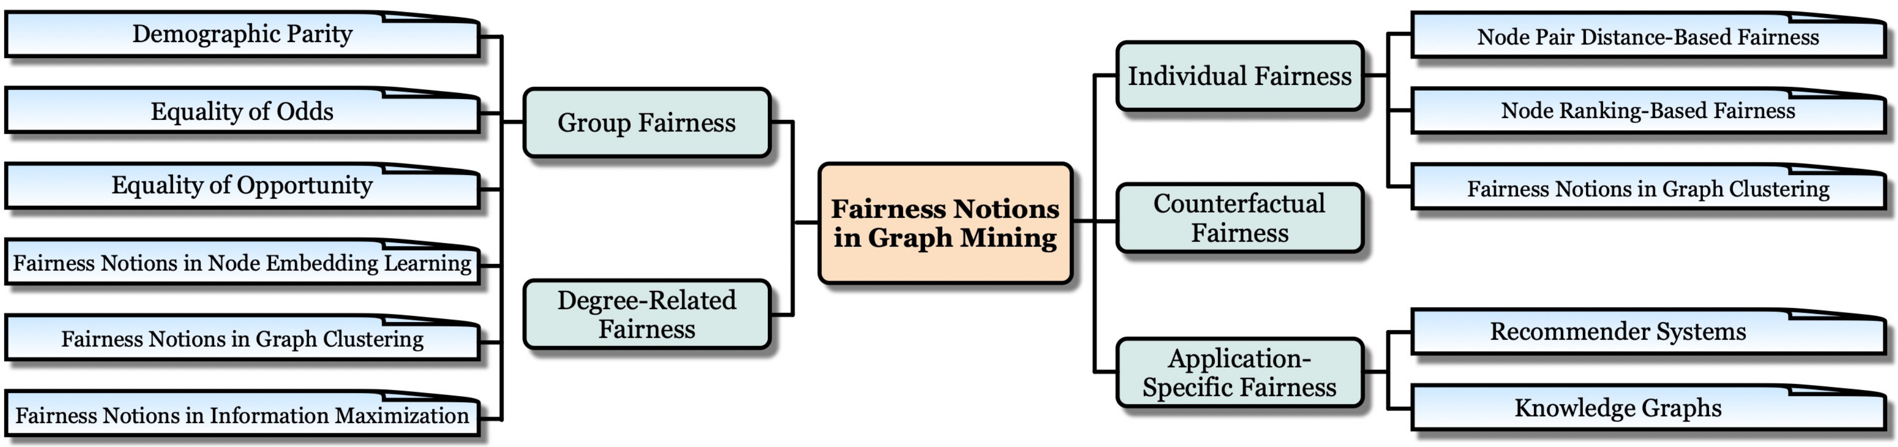

Fairness Notions and Metrics

- A taxonomy of commonly used algorithmic fairness notions in graph mining algorithms [2].

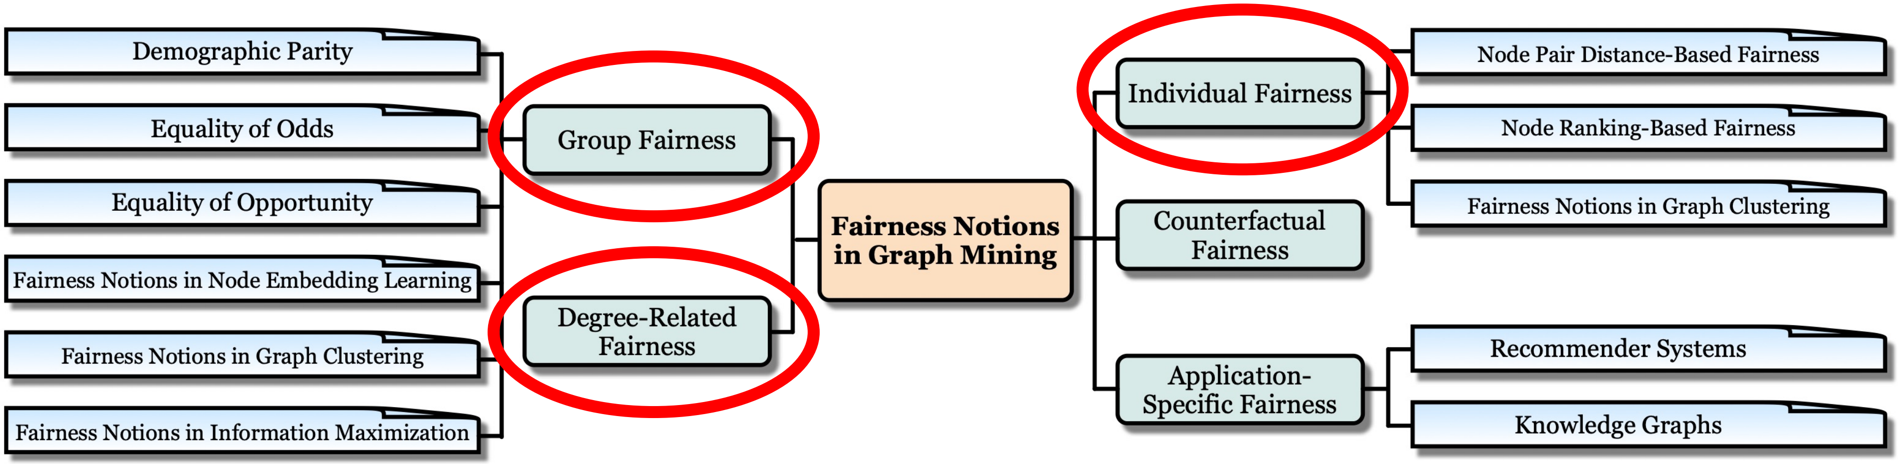

Fairness Notions and Metrics

- A taxonomy of commonly used algorithmic fairness notions in graph mining algorithms [2].

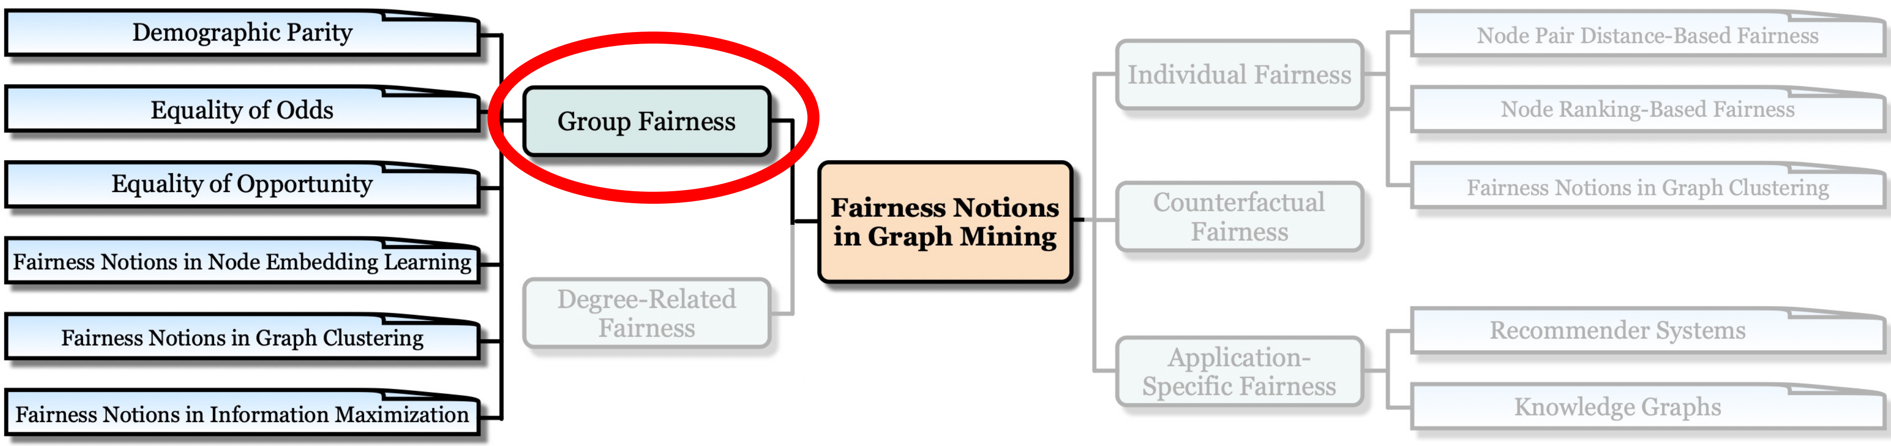

A general idea of group fairness: categorical **sensitive attributes** (e.g., gender, race) divide the whole population into different demographic subgroups, and each group should gain **their fair share of interest** [3].

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

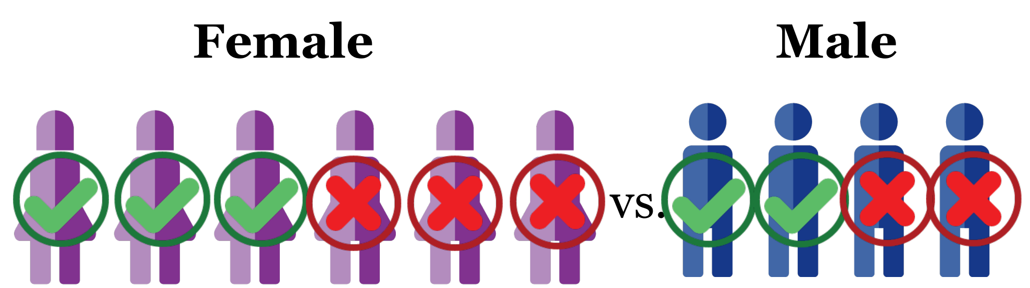

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

Fair in perspective of Demographic Parity.

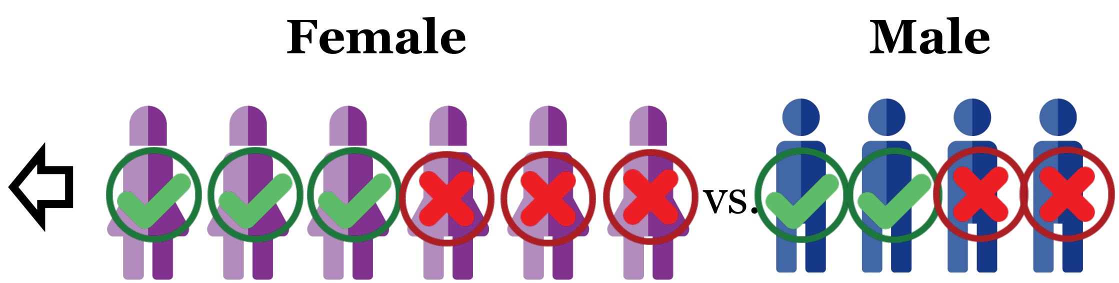

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

Fair in perspective of Demographic Parity.

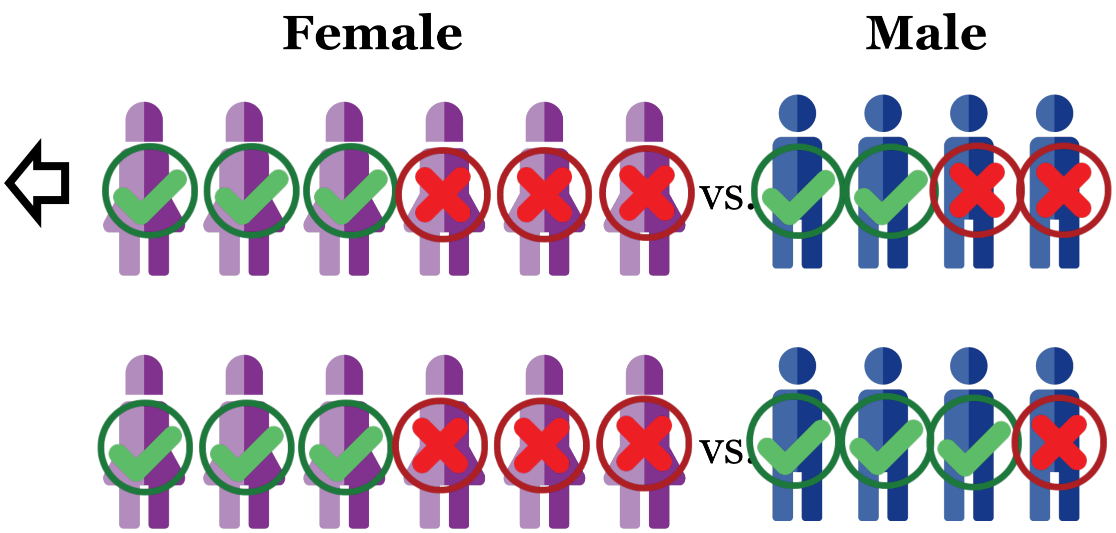

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

Fair in perspective of Demographic Parity.

Unfair in perspective of Demographic Parity.

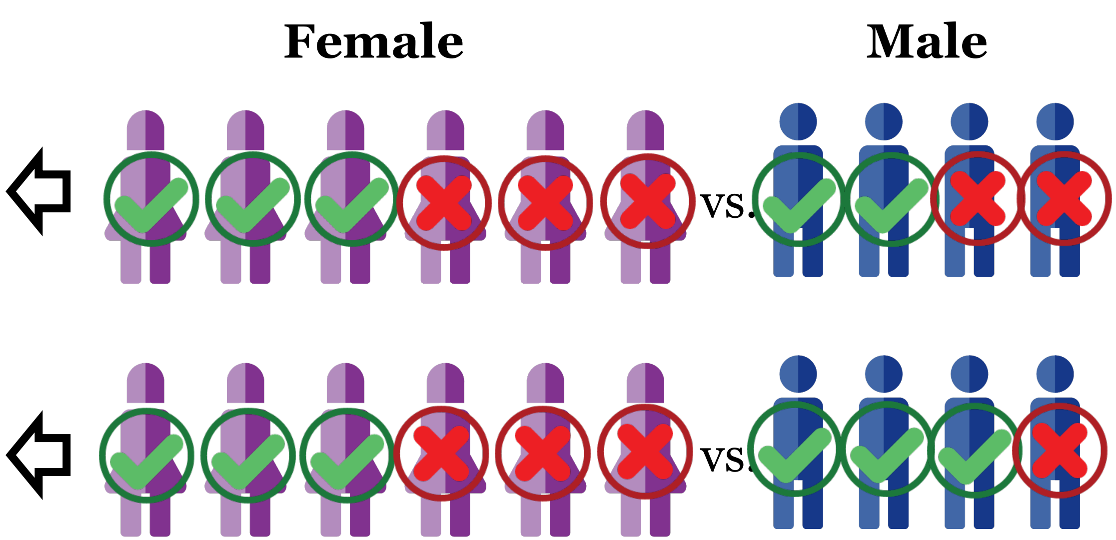

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Demographic Parity is first proposed in **binary classification task** for tabular data in [3].

Demographic Parity is considered as achieved if the model yields the **same positive rate** for individuals in both sensitive subgroups.

Criterion:

$$P(\hat{Y} = 1|S = 0) = P(\hat{Y} = 1|S = 1)$$

Metric:

$$\Delta_{DP} = |P(\hat{Y} = 1 \mid S = 0) - P(\hat{Y} = 1 \mid S = 1)|$$

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Interesting take-away points:

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Interesting take-away points:

(1) Demographic Parity is defined in **binary classification task** for tabular data [3], and there should be categorical sensitive feature(s);

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Interesting take-away points:

(1) Demographic Parity is defined in **binary classification task** for tabular data [3], and there should be categorical sensitive feature(s);

(2) In classification, it does not consider the ground truth labels. Consequently, enforcing demographic parity **may lead to sacrifice on utility** (e.g., classification accuracy) in practice;

Fairness Notions and Metrics

- Group Fairness: Demographic Parity

Interesting take-away points:

(1) Demographic Parity is defined in **binary classification task** for tabular data [3], and there should be categorical sensitive feature(s);

(2) In classification, it does not consider the ground truth labels. Consequently, enforcing demographic parity **may lead to sacrifice on utility** (e.g., classification accuracy) in practice;

(3) Recent works on fairness have **extended this notion to other settings**, including link prediction [4, 5] and scenarios with continuous sensitive feature(s) [6];

Fairness Notions and Metrics

- Group Fairness:
  Equality of Odds [7] vs. **Equality of Opportunity** [7]

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

Equality of Odds: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **ground truth class labels**.

$$P(\hat{Y} = 1 | S = 0, Y = y) = P(\hat{Y} = 1 | S = 1, Y = y)$$

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

Equality of Odds: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **ground truth class labels**.

$$P(\hat{Y}=1|S=0,Y=y)=P(\hat{Y}=1|S=1,Y=y)$$

Intuition: to enforce the true positive rate (right and beneficial results) and false positive rate (wrong but beneficial results) to be the same across groups;

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

Equality of Odds: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **ground truth class labels**.

$$P(\hat{Y} = 1|S = 0, Y = y) = P(\hat{Y} = 1|S = 1, Y = y)$$

Intuition: to enforce the true positive rate (right and beneficial results) and false positive rate (wrong but beneficial results) to be the same across groups;

Equality of Opportunity: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **positive ground truth class labels**.

$$P(\hat{Y} = 1|S = 0, Y = 1) = P(\hat{Y} = 1|S = 1, Y = 1)$$

Fairness Notions and Metrics

- Group Fairness:

Equality of Odds [7] vs. **Equality of Opportunity** [7]

Equality of Odds: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **ground truth class labels**.

Intuition: to enforce the true positive rate (right and beneficial results) and false positive rate (wrong but beneficial results) to be the same across groups;

$$P(\hat{Y}=1|S=0,Y=y)=P(\hat{Y}=1|S=1,Y=y)$$

Equality of Opportunity: the **positive rate** are enforced to be the same between demographic subgroups conditional on the **positive ground truth class labels**.

Intuition: to enforce the true positive rate (right and beneficial results) to be the same across groups;

$$P(\hat{Y}=1|S=0,Y=1)=P(\hat{Y}=1|S=1,Y=1)$$

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

  Equality of Opportunity is more widely studied in graph mining tasks, e.g., node classification.

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

  Equality of Opportunity is more widely studied in graph mining tasks, e.g., node classification.

  Criterion:

$$P(\hat{Y} = 1 | S = 0, Y = 1) = P(\hat{Y} = 1 | S = 1, Y = 1)$$

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

  Equality of Opportunity is more widely studied in graph mining tasks, e.g., node classification.

  Criterion:

$$P(\hat{Y} = 1 | S = 0, Y = 1) = P(\hat{Y} = 1 | S = 1, Y = 1)$$

  Metric:

$$\Delta_{EO} = |P(\hat{Y} = 1 \mid Y = 1, S = 0) - P(\hat{Y} = 1 \mid Y = 1, S = 1)|$$

Fairness Notions and Metrics

- Group Fairness:

  Equality of Odds [7] vs. **Equality of Opportunity** [7]

  Equality of Opportunity is more widely studied in graph mining tasks, e.g., node classification.

  Criterion:

  $$P(\hat{Y}=1|S=0,Y=1)=P(\hat{Y}=1|S=1,Y=1)$$

  Metric:

  $$\Delta_{EO}=|P(\hat{Y}=1 \mid Y=1,S=0) - P(\hat{Y}=1 \mid Y=1,S=1)|$$

  Equality of Opportunity has also been extended to tasks other than node classification, e.g., edge prediction [4, 8].

Fairness Notions and Metrics

- Another critical fairness notion in graph mining: Individual Fairness.

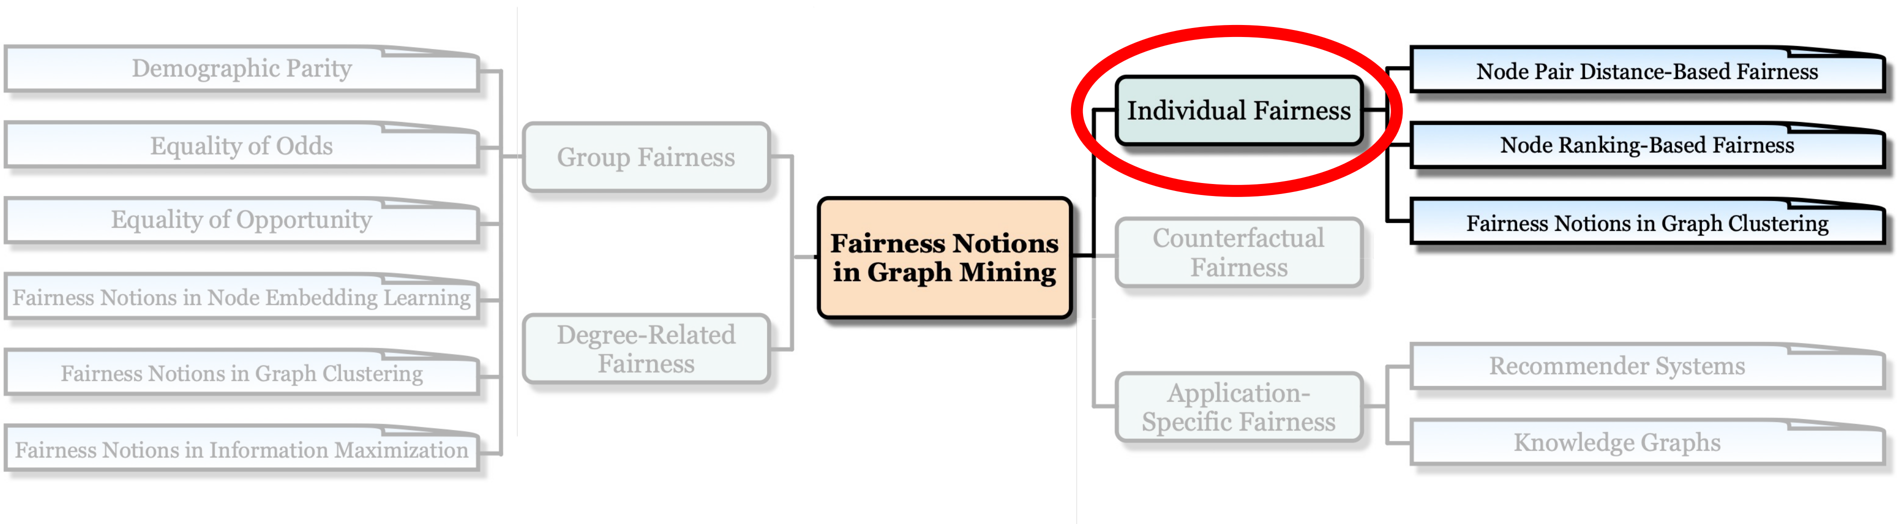

A general idea of individual fairness: **similar individual should receive similar output** from the graph mining algorithms [9].

Fairness Notions and Metrics

- Individual Fairness:

(1) Node Pair Distance-Based Fairness [10]

For any pair of node, this fairness notion enforces **the output distance to be smaller than a scaled input distance** - which is consistent with the general idea of "similar individual should receive similar output".

Fairness Notions and Metrics

- Individual Fairness:

(1) Node Pair Distance-Based Fairness [10]

For any pair of node, this fairness notion enforces **the output distance to be smaller than a scaled input distance** - which is consistent with the general idea of "similar individual should receive similar output".

Mathematically, we have

$$\underbrace{D_1(f(x), f(y))}_{\text{Output distance}} \leq L\underbrace{D_2(x, y)}_{\text{Input distance}} \quad \forall\ (x, y)$$

L: Lipschitz Constant

In practice, we enforce the following inequality

$$\|\mathbf{Y}[i,:] - \mathbf{Y}[j,:]\|_F^2 \leq \frac{\epsilon}{\mathbf{S}[i,j]}\ \forall i, j = 1, \ldots, n$$

Fairness Notions and Metrics

- Individual Fairness:

(2) Node Ranking-Based Fairness [11]

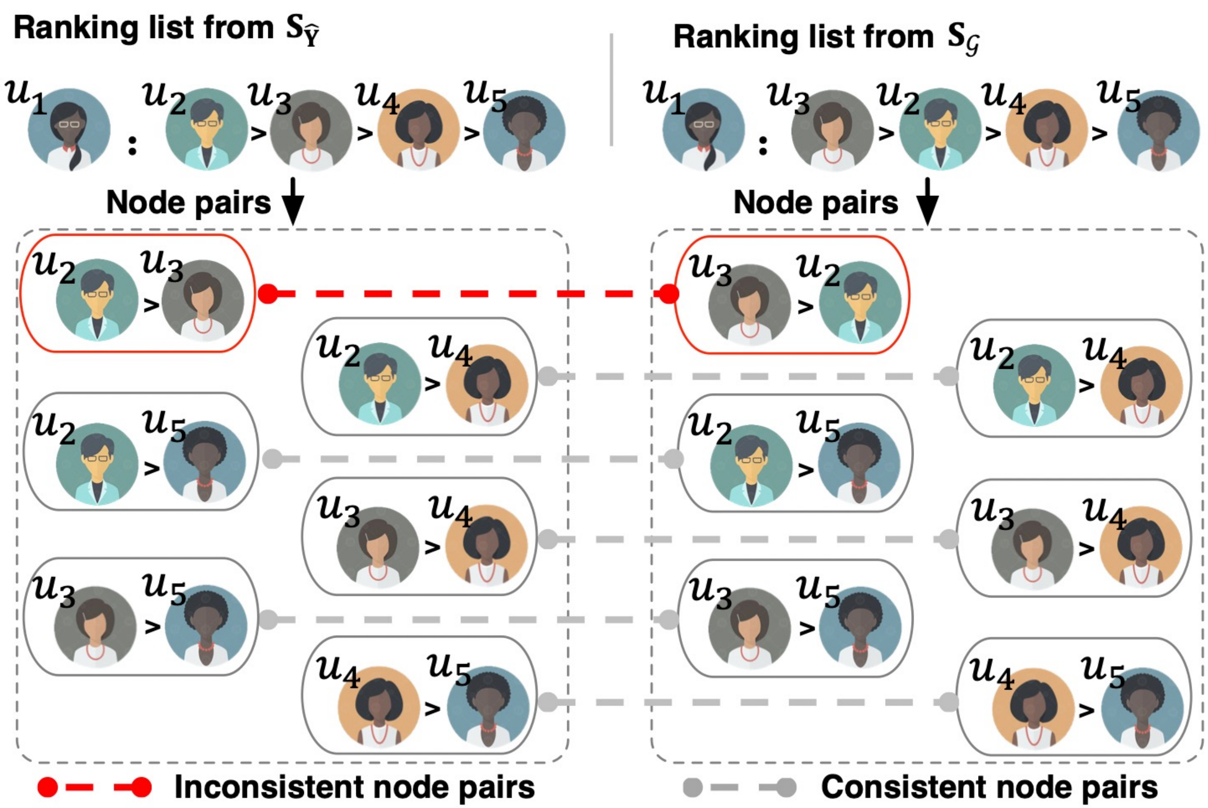

Fairness Notions and Metrics

- Individual Fairness:

(2) Node Ranking-Based Fairness [11]

Criterion: Given a graph $\mathcal{G}$ with five nodes, suppose the ranking list that encodes the similarity between node u_1 and other nodes from $\mathbf{S}_{\mathcal{G}}$ is $\{u_4, u_3, u_2, u_5\}$, we say the predictions are are individually fair for node u_1 if the ranking list that encodes the similarity between u_1 and other nodes from $\mathbf{S}_{\hat{\mathbf{Y}}}$ is also $\{u_4, u_3, u_2, u_5\}$.

Metrics: average ranking similarity across all individuals, e.g., average NDCG@k [12].

Fairness Notions and Metrics

- A fairness notion **tailored with graph structure**: Degree-Related Fairness.

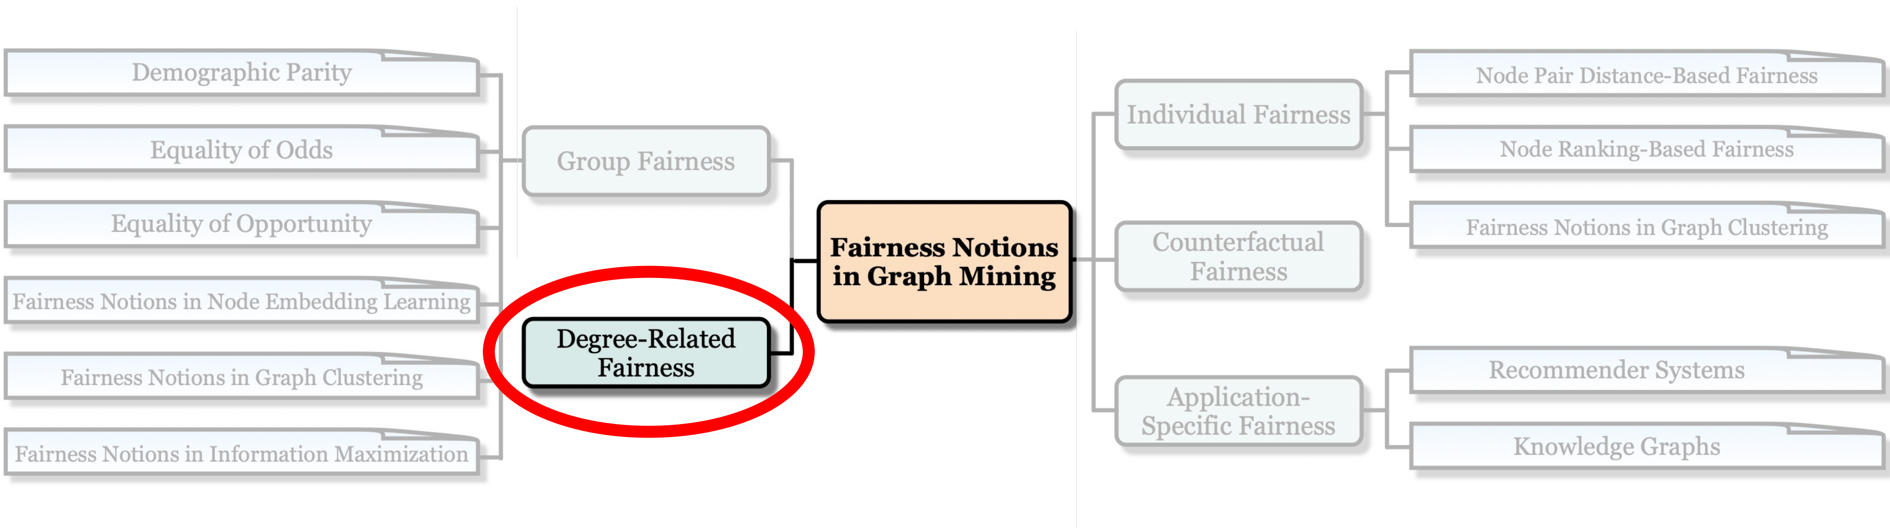

A general idea of degree-related fairness: the degree of nodes should be independent from the quality of their corresponding predictions [13, 14].

Fairness Notions and Metrics

- Degree-Related Fairness:

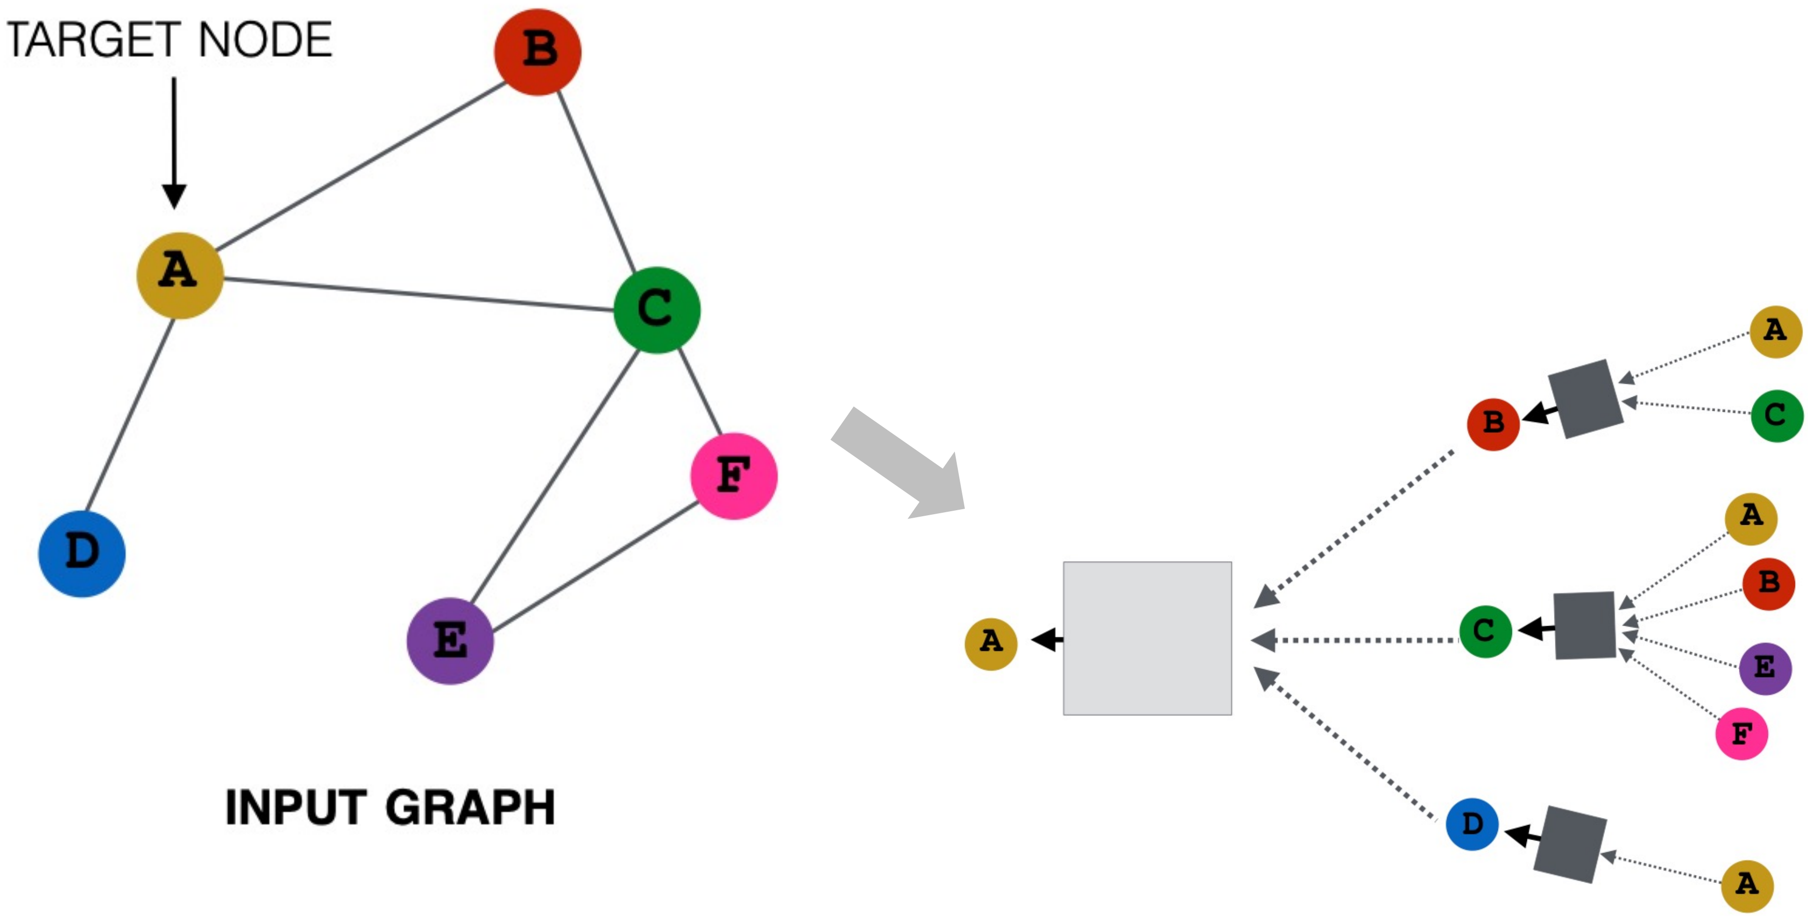

Fairness Notions and Metrics

- Degree-Related Fairness:

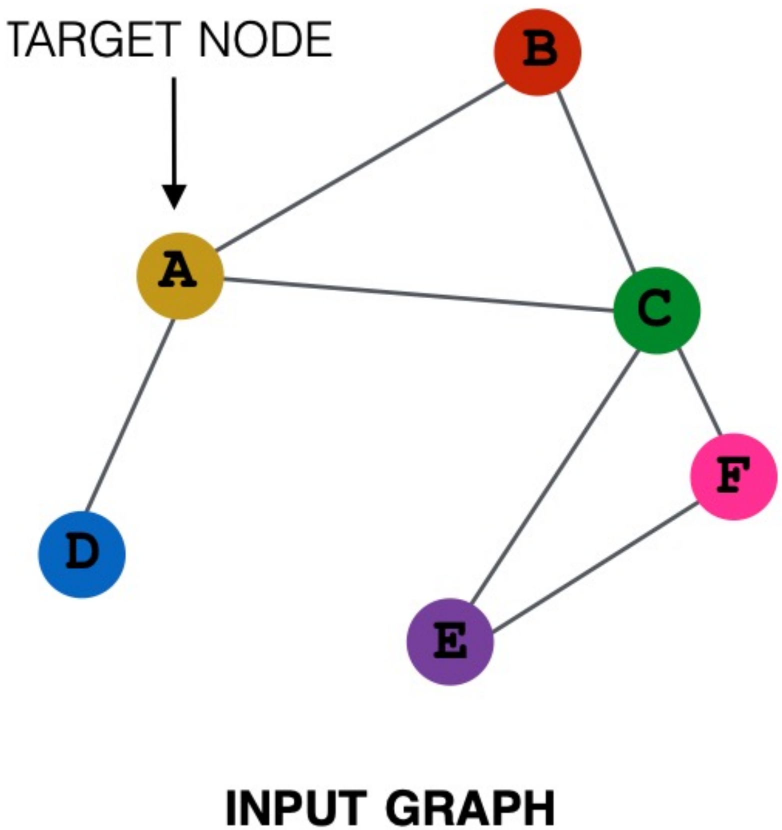

In graph data, a critical source of information is **the complementary information between neighbors.**

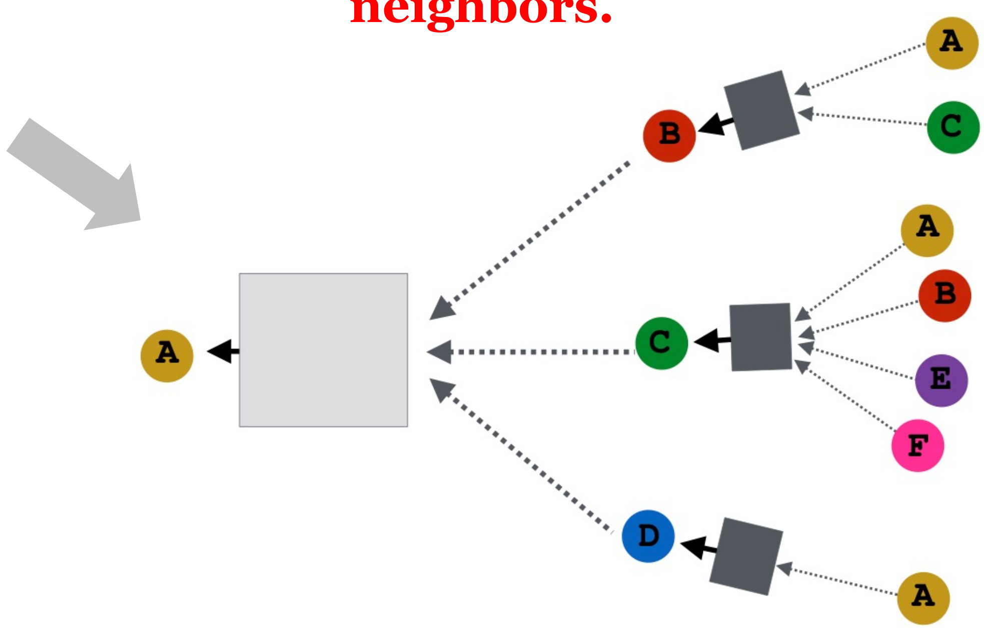

Fairness Notions and Metrics

- Degree-Related Fairness:

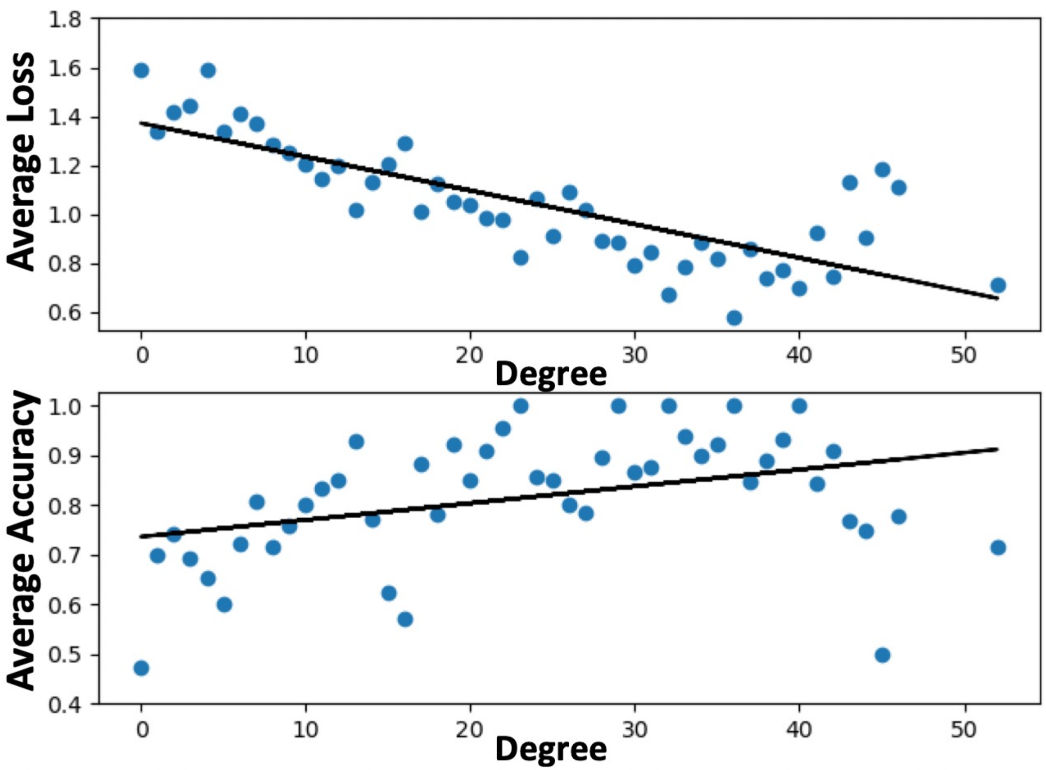

In graph data, a critical source of information is **the complementary information between neighbors.**

Fairness Notions and Metrics

- Degree-Related Fairness:

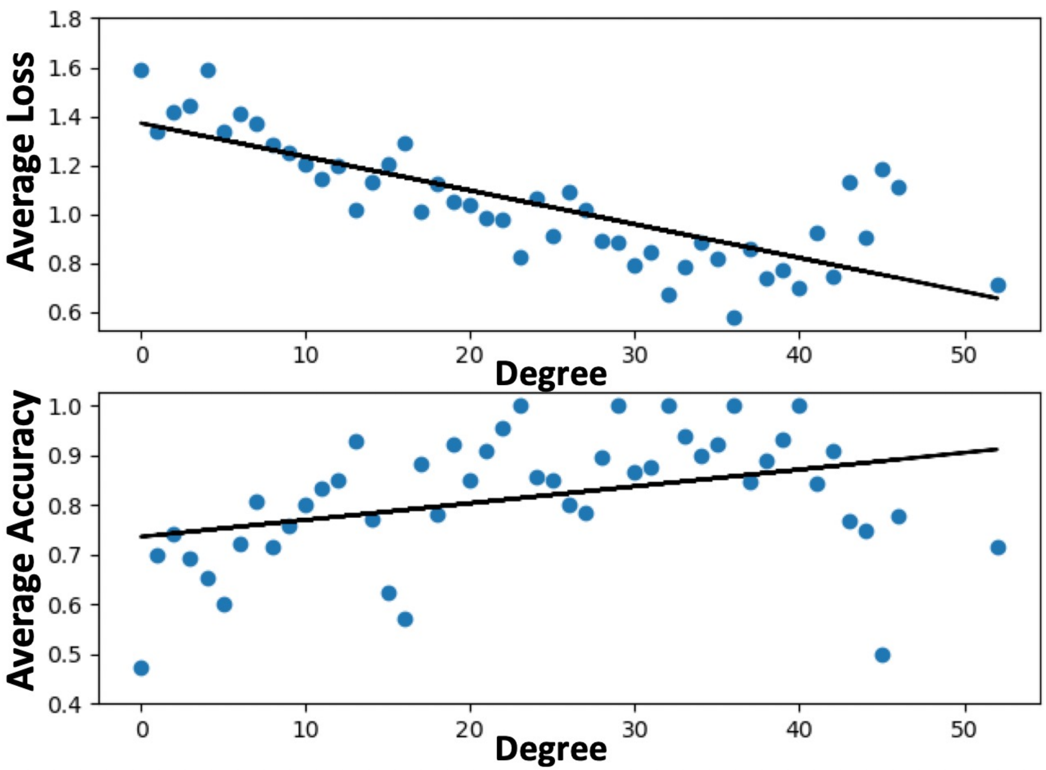

In graph data, a critical source of information is **the complementary information between neighbors.**

However, graph mining algorithms rely on such information tend to yield predictions with **much worse quality** for low-degree nodes, as they have **fewer neighbors**.

Fairness Notions and Metrics

- Degree-Related Fairness:

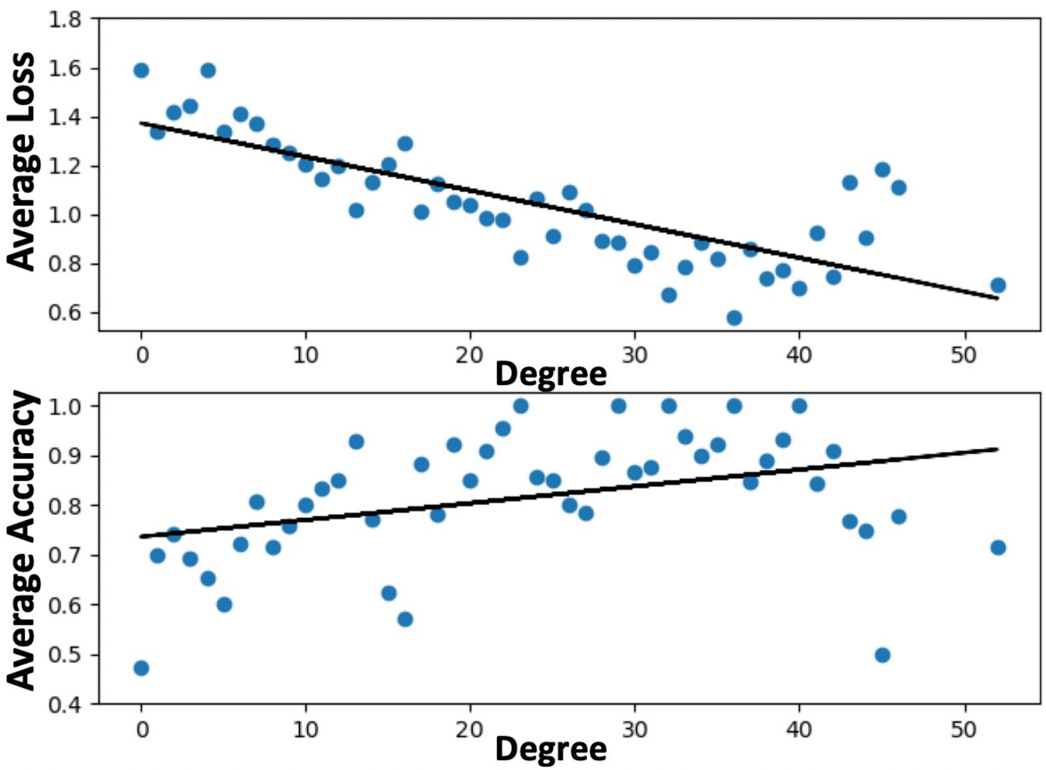

In graph data, a critical source of information is **the complementary information between neighbors.**

However, graph mining algorithms rely on such information tend to yield predictions with **much worse quality** for low-degree nodes, as they have **fewer neighbors**.

Degree-Related Fairness requires that nodes should bear similar utility (e.g., node classification accuracy) in the graph mining algorithms **regardless of their degrees**.

Outline

Background Information

Fairness Notions and Metrics

Methodologies to Mitigate Bias

Existing Problems & Future Works

Methodologies to Mitigate Bias

- In general, there are six main categories of commonly used techniques to improve fairness in graph mining.

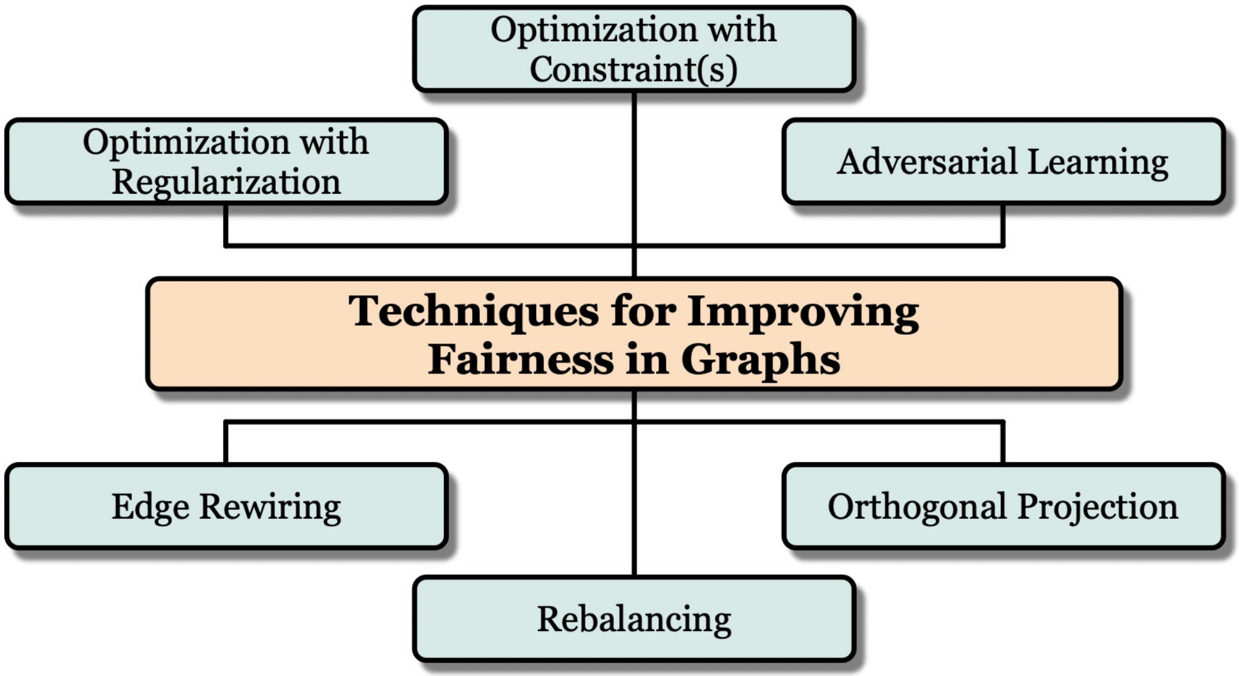

Methodologies to Mitigate Bias

- Optimization with Regularization.

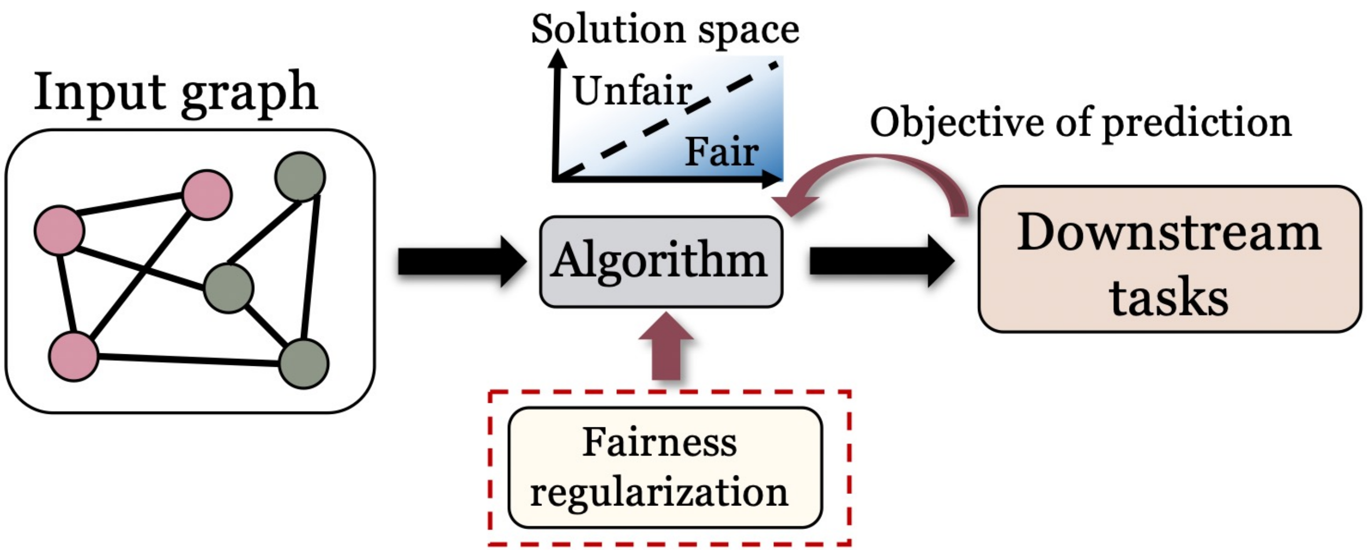

Methodologies to Mitigate Bias

- Optimization with Regularization.

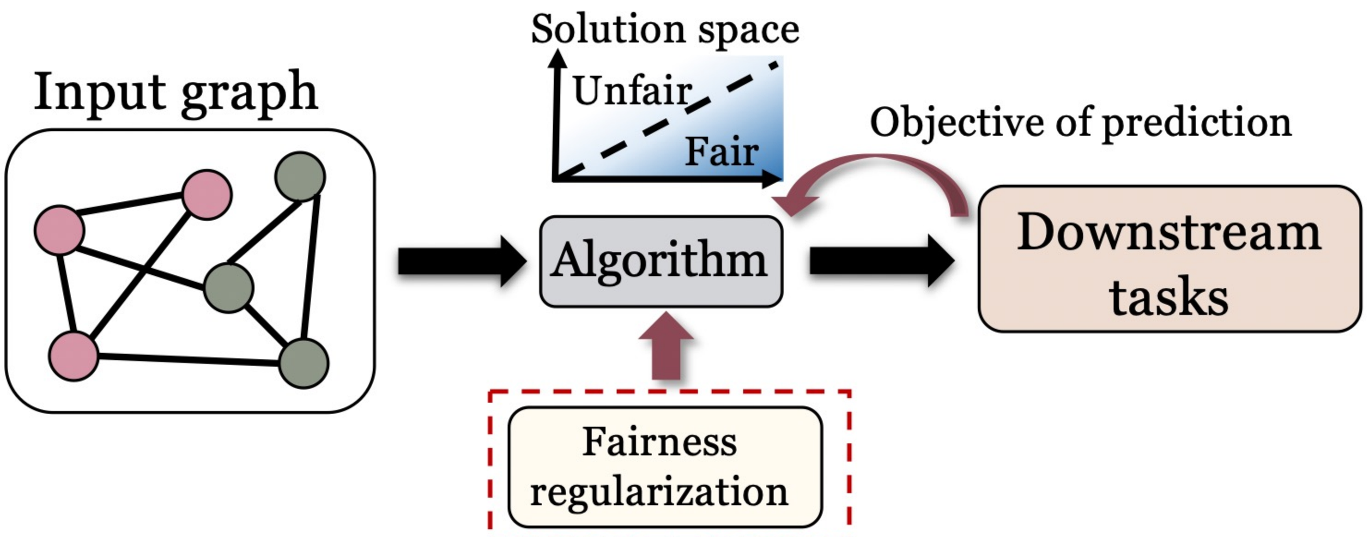

Adding a fairness-aware loss term on the total objective function.

$$\mathscr{L} = \mathscr{L}_{\text{utility}} + \lambda \mathscr{L}_{\text{fair}}$$

- Output logits-based regularization [15].
- Network topology-based regularization [16].
- Node Embedding-Based Regularization [17].

Methodologies to Mitigate Bias

- Optimization with Constraint(s).

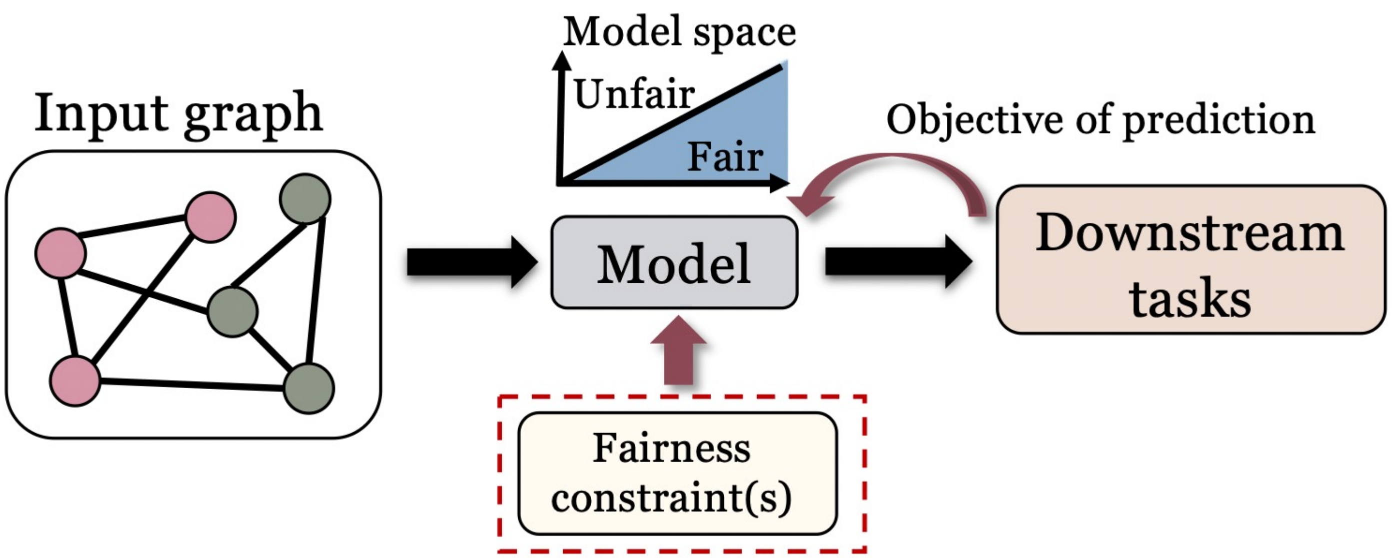

Methodologies to Mitigate Bias

- Optimization with Constraint(s).

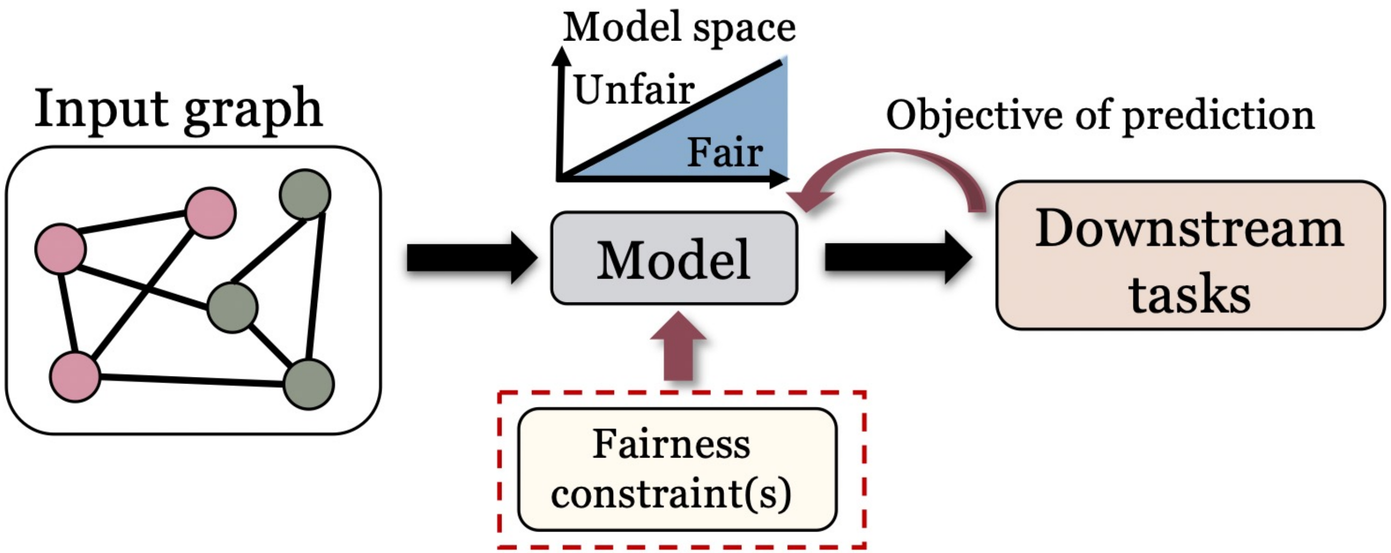

Adding a fairness-aware constraint on the optimization problem [18].

$$\begin{aligned} \min \quad & \mathscr{L}_{\text{utility}}, \\ \text{subject to} \quad & \text{certain fairness constraint(s)} \end{aligned}$$

Most existing works formulate such a constraint with the **performance difference** on different demographic subgroups.

Methodologies to Mitigate Bias

- Adversarial Learning.

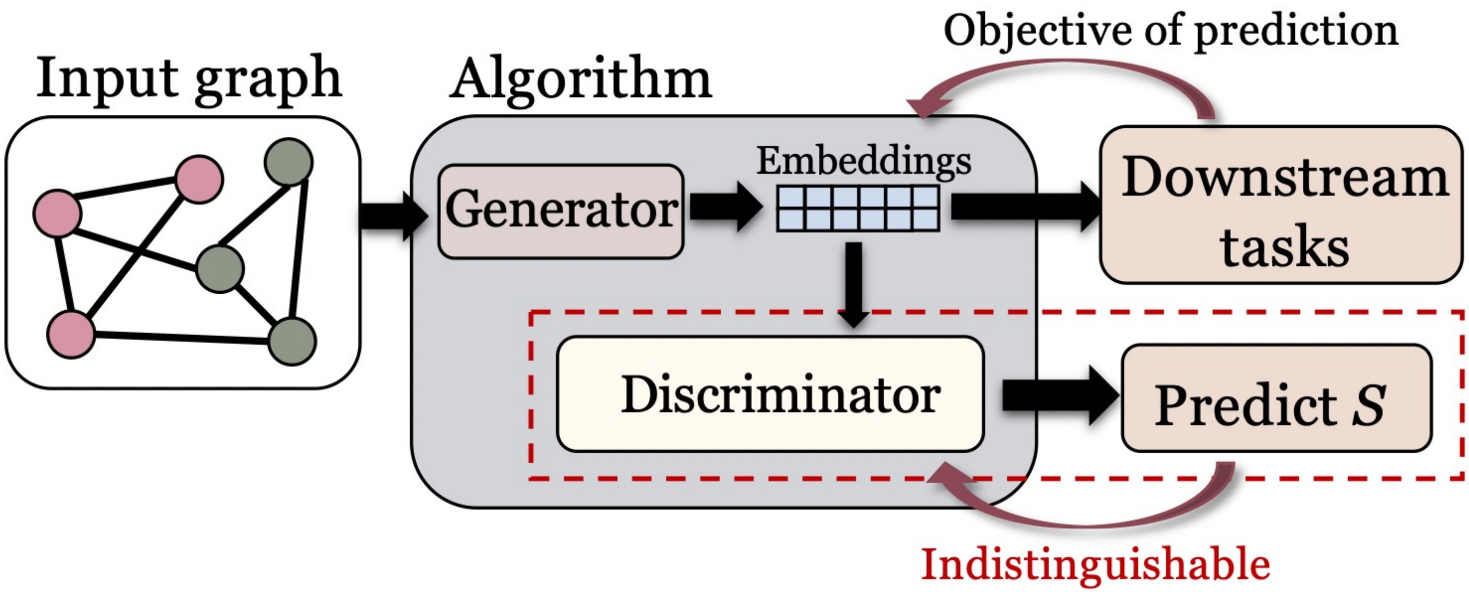

A general formulation of fulfilling fairness with adversarial learning includes a **generator** and a **discriminator** [19]:

Methodologies to Mitigate Bias

- Adversarial Learning.

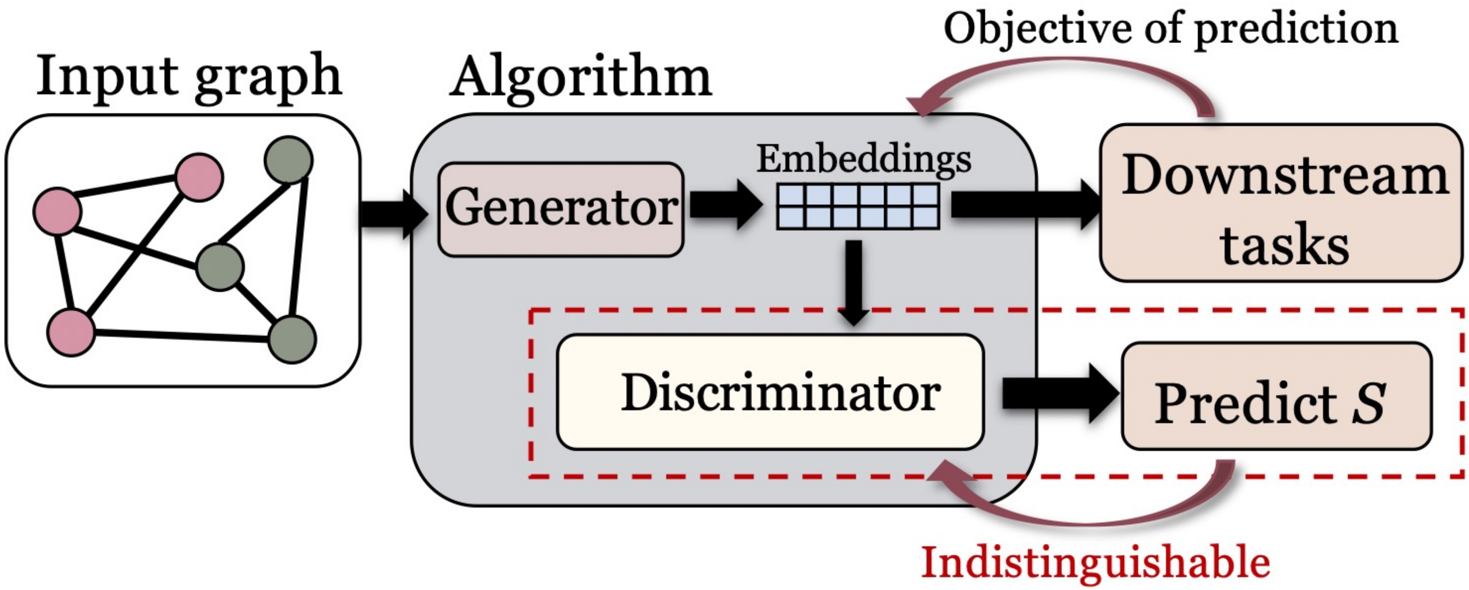

A general formulation of fulfilling fairness with adversarial learning includes a **generator** and a **discriminator** [19]:

Generator: generate node embeddings for downstream tasks;

Discriminator: distinguish the embeddings between demographic subgroups;

Methodologies to Mitigate Bias

- Edge Rewiring.

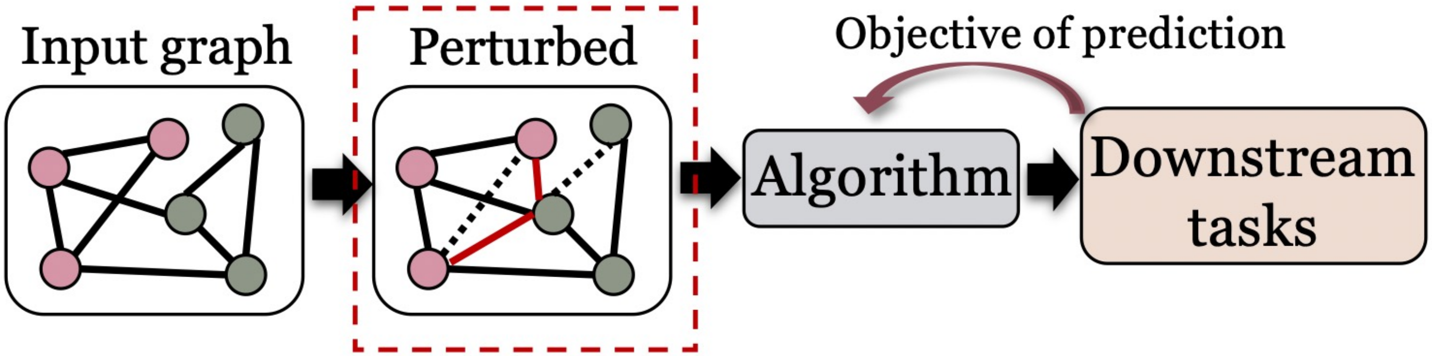

Methodologies to Mitigate Bias

- Edge Rewiring.

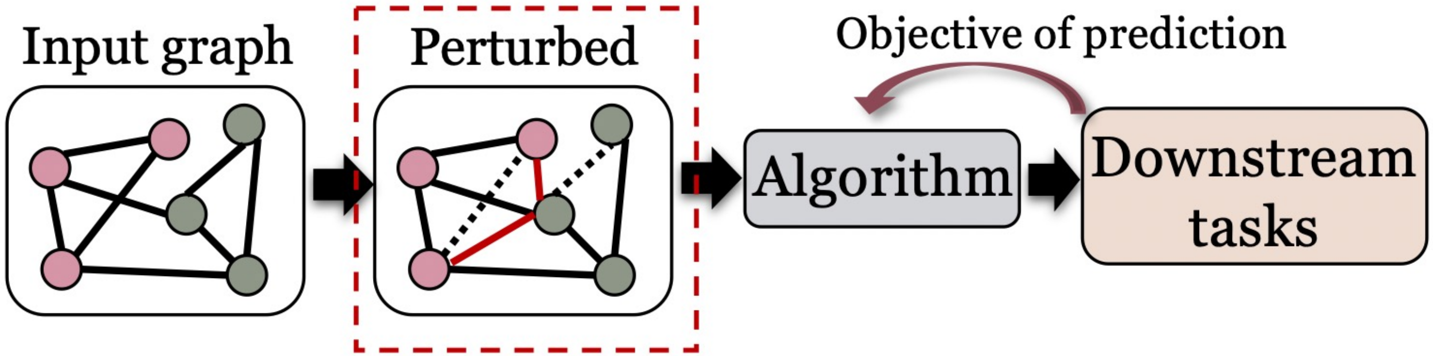

There could be **bias encoded in the network structure**, and edge rewiring aims to achieve a fairer structure for the graph mining algorithm.

Methodologies to Mitigate Bias

- Edge Rewiring.

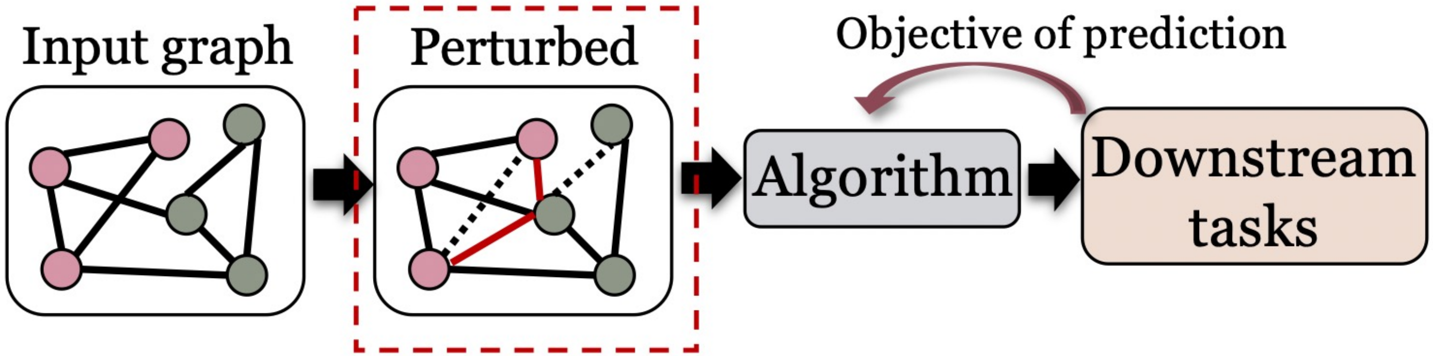

There could be **bias encoded in the network structure**, and edge rewiring aims to achieve a fairer structure for the graph mining algorithm.

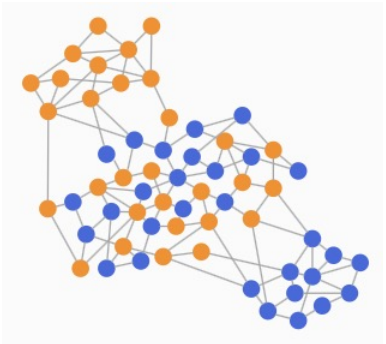

An example of **biased graph structure**: clear **community structure** between two groups of nodes, where the membership is dependent on sensitive feature(s) [20].

Methodologies to Mitigate Bias

- Rebalancing.

  Rebalancing could be achieved in different ways **depending on the characteristics of the graph mining algorithms & tasks**.

Methodologies to Mitigate Bias

- Rebalancing.

  Rebalancing could be achieved in different ways **depending on the characteristics of the graph mining algorithms & tasks**.

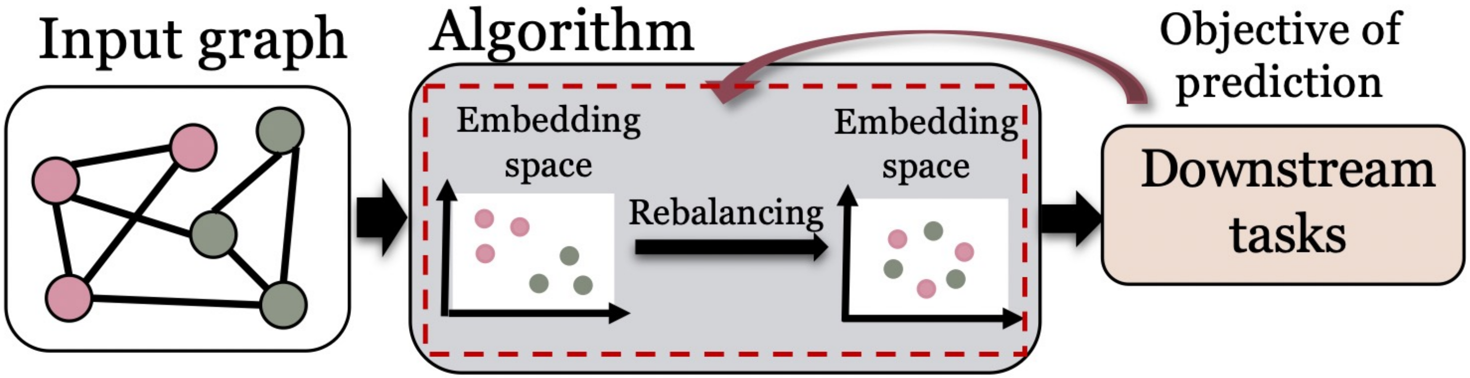

Methodologies to Mitigate Bias

- Rebalancing.

  Rebalancing could be achieved in different ways **depending on the characteristics of the graph mining algorithms & tasks**.

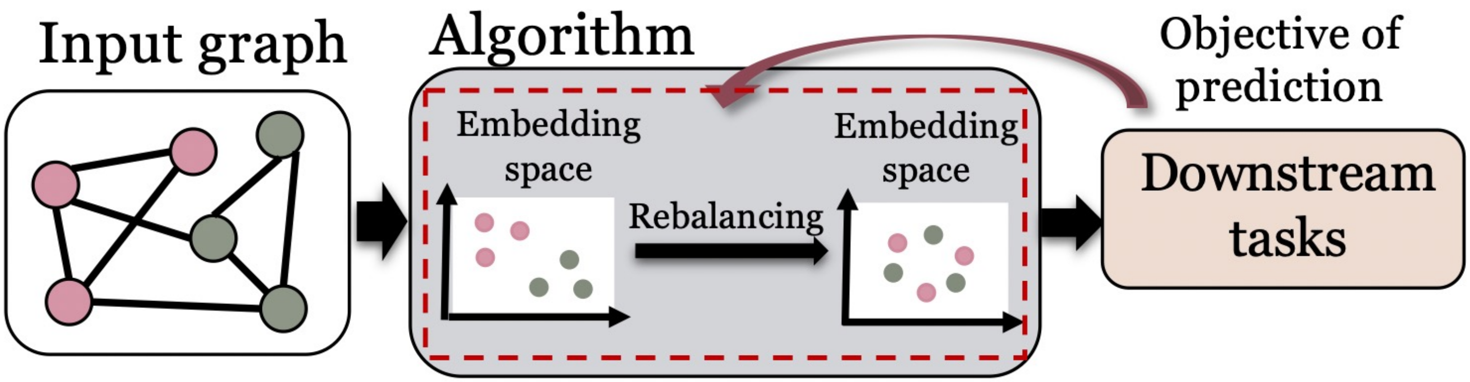

An example of rebalancing in **random walk** on graphs [21]:

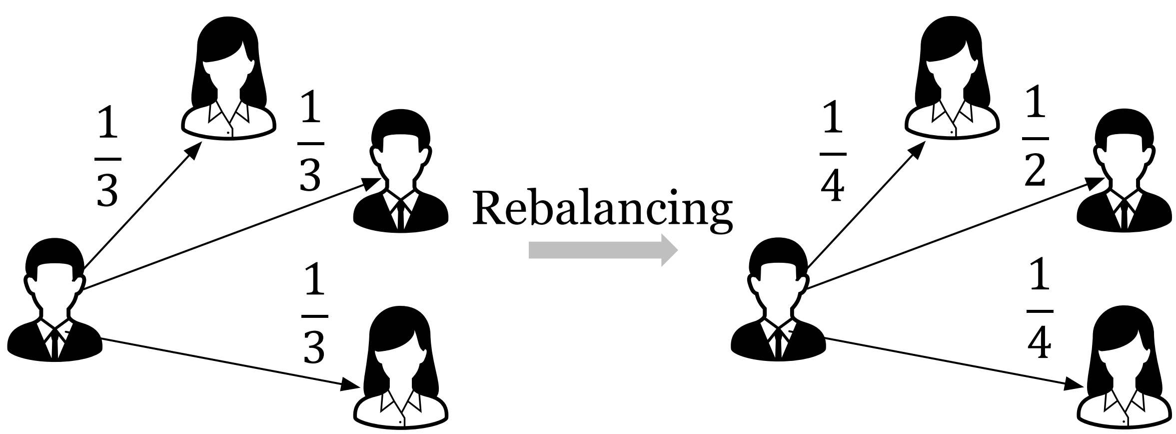

Methodologies to Mitigate Bias

- Orthogonal Projection.

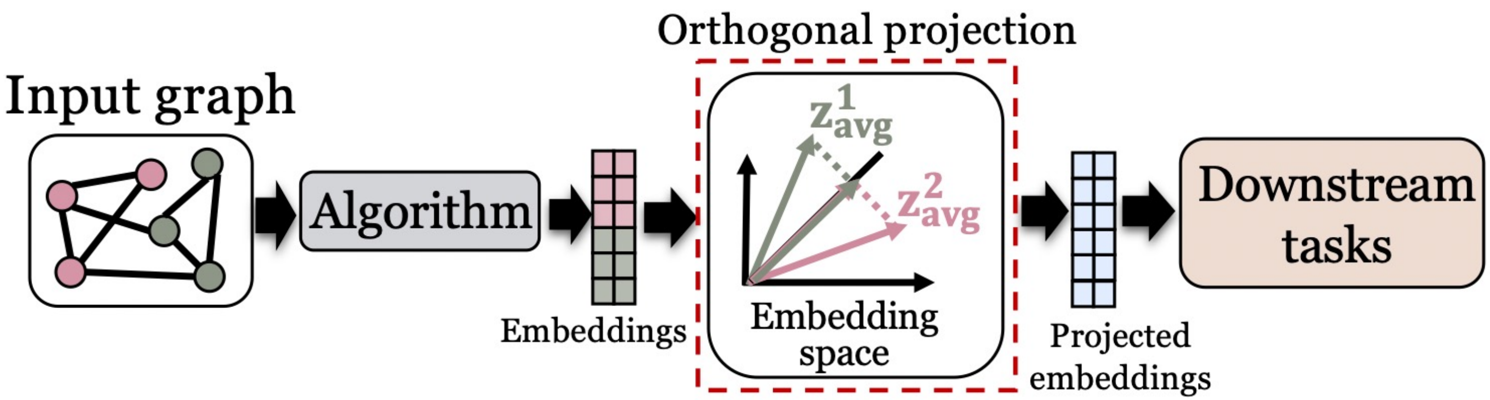

Methodologies to Mitigate Bias

- Orthogonal Projection.

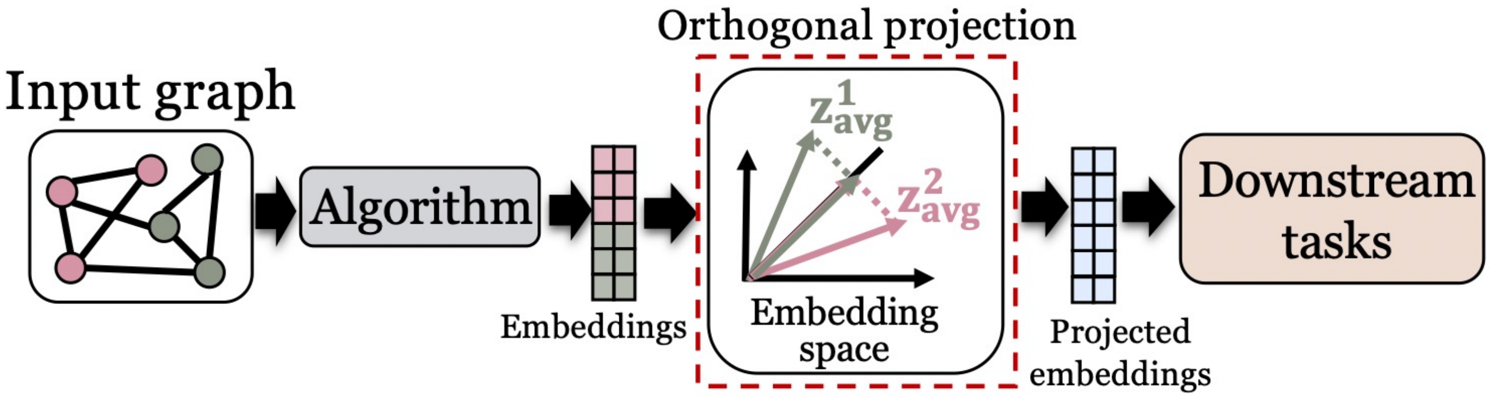

Intuition: if the node embeddings are projected onto the same hyper-plane, then there will be **no correlation** between node embeddings and bias (usually sensitive features).

Methodologies to Mitigate Bias

- Orthogonal Projection.

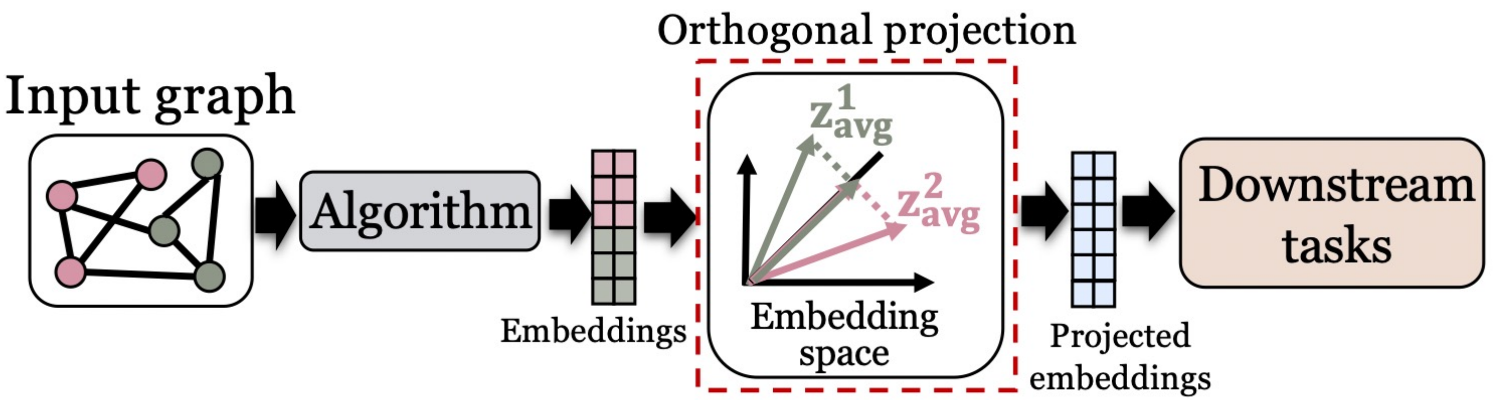

Intuition: if the node embeddings are projected onto the same hyper-plane, then there will be **no correlation** between node embeddings and bias (usually sensitive features).

An example [22]:

Find bias direction.

Project embeddings onto a hyper-plane orthogonal to the bias direction.

$$\mathbf{z}^i_{\text{avg}} = \frac{\mathbf{z}_1 + \mathbf{z}_2 + \ldots + \mathbf{z}_{|\mathcal{V}_i|}}{\|\mathbf{z}_1 + \mathbf{z}_2 + \ldots + \mathbf{z}_{|\mathcal{V}_i|}\|_2} \rightarrow \mathbf{z}_{\text{bias}} = \frac{\mathbf{z}^1_{\text{avg}} - \mathbf{z}^2_{\text{avg}}}{\|\mathbf{z}^1_{\text{avg}} - \mathbf{z}^2_{\text{avg}}\|_2} \rightarrow \mathbf{z}'_j = \mathbf{z}_j - < \mathbf{z}_j, \mathbf{z}_{\text{bias}} > \mathbf{z}_{\text{bias}}$$

Outline

Background Information

Fairness Notions and Metrics

Methodologies to Mitigate Bias

Existing Problems & Future Works

Existing Problems & Future Works

- (1) Lack of Fairness Notions.

Can existing fairness notions help to avoid all cases where people may feel unfair?

Existing Problems & Future Works

- (1) Lack of Fairness Notions.

- (2) Fulfilling Multiple Types of Fairness.

How to achieve multiple types of fairness?

Are some of the existing fairness notions in conflict with each other?

If we could achieve multiple types of fairness, will people get a stronger sense of fairness? If not, what will be beneficial for social good?

Existing Problems & Future Works

- (1) Lack of Fairness Notions.
- (2) Fulfilling Multiple Types of Fairness.
- (3) Balancing Model Utility and Algorithmic Fairness.

How to achieve fairness at low or no cost of utility?

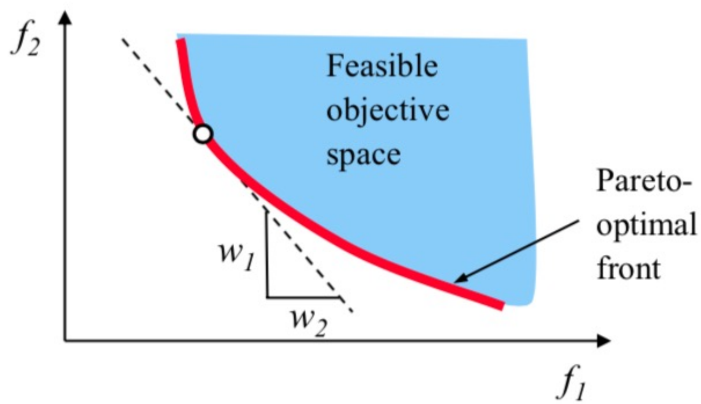

Existing Problems & Future Works

- (1) Lack of Fairness Notions.
- (2) Fulfilling Multiple Types of Fairness.
- (3) Balancing Model Utility and Algorithmic Fairness.
- (4) Explaining How Unfairness Arises.

  How to interpret why unfairness arises in graph mining algorithms?

  Is the graph data biased?

  Is the model biased naturally?

Existing Problems & Future Works

- (1) Lack of Fairness Notions.

- (2) Fulfilling Multiple Types of Fairness.

- (3) Balancing Model Utility and Algorithmic Fairness.

- (4) Explaining How Unfairness Arises.

- (5) Enhancing Robustness of Algorithms on Fairness.

  How would existing graph mining algorithms perform in perspective of fairness under malicious attack?

  How to achieve better robustness in perspective of fairness?

References

[1] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. Fairness in deep learning: A computational perspective. IEEE Intelligent Systems, 2020.

[2] Dong, Y., Ma, J., Chen, C., & Li, J. (2022). Fairness in Graph Mining: A Survey. arXiv preprint arXiv:2204.09888.

[3] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In Innovations in Theoretical Computer Science, 2012.

[4] Maarten Buyl and Tijl De Bie. Debayes: a bayesian method for debiasing network embeddings. In ICML, 2020.

[5] Charlotte Laclau, Ievgen Redko, Manvi Choudhary, and Chris- tine Largeron. All of the fairness for edge prediction with optimal transport. In AISTATS, 2021.

[6] Zhimeng Jiang, Xiaotian Han, Chao Fan, Fan Yang, Ali Mostafavi, and Xia Hu. Generalized demographic parity for group fairness. In ICLR, 2021.

[7] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In NeurIPS, 2016.

[8] Blake W., Suriya G., Mesrob I O., and Nathan S. Learning nondiscriminatory predictors. In COLT, 2017.

[9] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In Innovations in Theoretical Computer Science, 2012.

[10] Jian Kang, Jingrui He, Ross Maciejewski, and Hanghang Tong. Inform: Individual fairness on graph mining. In SIGKDD, 2020.

[11] Yushun Dong, Jian Kang, Hanghang Tong, and Jundong Li. Individual fairness for graph neural networks: A ranking based approach. In SIGKDD, 2021.

[12] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. ACM Transactions on Information Systems 20, 2002.

[13] Xianfeng Tang, Huaxiu Yao, Yiwei Sun, Yiqi Wang, Jiliang Tang, Charu Aggarwal, Prasenjit Mitra, and Suhang Wang. Investigating and mitigating degree-related biases in graph convoltuional networks. In CIKM, 2020

[14] Jian Kang, Yan Zhu, Jiebo Luo, Yinglong Xia, and Hanghang Tong. Rawlsgcn: Towards rawlsian difference principle on graph convolutional network. In WWW, 2022.

[15] Ziqian Zeng, Rashidul Islam, et al. Fair representation learning for heterogeneous information networks. In AAAI, 2021.

[16] Zhimeng Jiang, Xiaotian Han, Chao Fan, Zirui Liu, Na Zou, Ali Mostafavi, and Xia Hu. Fmp: Toward fair graph message passing against topology bias. arXiv preprint arXiv:2202.04187, 2022.

[17] Preethi Lahoti, Krishna P. Gummadi, and Gerhard Weikum. Operationalizing individual fairness with pairwise fair represen- tations. VLDB, 2019.

[18] Mattha ̈us Kleindessner, Samira Samadi, Pranjal Awasthi, and Jamie Morgenstern. Guarantees for spectral clustering with fairness constraints. In ICML, 2019.

[19] Enyan Dai and Suhang Wang. Say no to the discrimination: Learning fair graph neural networks with limited sensitive at- tribute information. In WSDM, 2021.

[20] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. EDITS: modeling and mitigating data bias for graph neural networks. In WWW, 2022.

[21] Tahleen Rahman, Bartlomiej Surma, Michael Backes, and Yang Zhang. Fairwalk: Towards fair graph embedding. In IICAI, 2019.

[22] John Palowitch and Bryan Perozzi. MONET: debiasing graph embeddings via the metadata-orthogonal training unit. ASONAM, 2020.

The End

VIRGINIA DATA SCIENCE AND ANALYTICS (VAST) LAB

Thanks for listening!