On the Role of Edge Dependency in Graph Generative Models

Sudhanshu Chanpuriya, Cameron Musco,
Konstantinos Sotiropoulos, Charalampos Tsourakakis

Graph Generative Models

Given a graph, we want to create a model that generates novel but similar graphs.

Why novel graphs?

- Test scalability & robustness of models and algorithms.
- Study networks under privacy.
- Drug discovery.

How similar?

- Match the input graph's structure and statistics as well as possible. . .
- But not to the point of reproducing the same edges!

How can we analyze the inherent trade-off between expressiveness and variety?

Graph Generative Models

Given a graph, we want to create a model that generates novel but similar graphs.

Why novel graphs?

- Test scalability & robustness of models and algorithms.
- Study networks under privacy.
- Drug discovery.

How similar?

- Match the input graph's structure and statistics as well as possible...
- But not to the point of reproducing the same edges!

How can we analyze the inherent trade-off between expressiveness and variety?

Graph Generative Models

Given a graph, we want to create a model that generates novel but similar graphs.

Why novel graphs?

- Test scalability & robustness of models and algorithms.
- Study networks under privacy.
- Drug discovery.

How similar?

- Match the input graph's structure and statistics as well as possible...
- But not to the point of reproducing the same edges!

How can we analyze the inherent trade-off between expressiveness and variety?

Graph Generative Models

Given a graph, we want to create a model that generates novel but similar graphs.

Why novel graphs?

- Test scalability & robustness of models and algorithms.
- Study networks under privacy.
- Drug discovery.

How similar?

- Match the input graph's structure and statistics as well as possible...
- But not to the point of reproducing the same edges!

How can we analyze the inherent trade-off between expressiveness and variety?

Quantifying Graph Variety

To concretize *variety*, prior work defines the **overlap** of a random graph model $\mathcal{G}$:

$$\mathrm{Ov}(\mathcal{G}) := \frac{\mathbb{E}_{G_1,G_2\sim\mathcal{G}}|E(G_1)\cap E(G_2)|}{\mathbb{E}_{G\sim\mathcal{G}}|E(G)|}.$$

Quantifying Graph Variety

To concretize *variety*, prior work defines the **overlap** of a random graph model $\mathcal{G}$:

$$\mathrm{Ov}(\mathcal{G}) := \frac{\mathbb{E}_{G_1, G_2 \sim \mathcal{G}} |E(G_1) \cap E(G_2)|}{\mathbb{E}_{G \sim \mathcal{G}} |E(G)|}.$$

Consider common **edge-independent** graph models: some probability matrix $\boldsymbol{P} \in [0,1]^{n \times n}$ and the distribution $\mathcal{G}$ on graphs where each edge is added independently with probability $\boldsymbol{P}_{ij}$.

It is shown that edge-independent models with low overlap are inherently limited in representing dense subgraph structure: the max expected triangle count **shrinks cubically in overlap**!

Quantifying Graph Variety

To concretize *variety*, prior work defines the **overlap** of a random graph model $\mathcal{G}$:

$$\mathrm{Ov}(\mathcal{G}) := \frac{\mathbb{E}_{G_1, G_2 \sim \mathcal{G}} |E(G_1) \cap E(G_2)|}{\mathbb{E}_{G \sim \mathcal{G}} |E(G)|}.$$

Consider common **edge-independent** graph models: some probability matrix $\boldsymbol{P} \in [0,1]^{n \times n}$ and the distribution $\mathcal{G}$ on graphs where each edge is added independently with probability $\boldsymbol{P}_{ij}$.

It is shown that edge-independent models with low overlap are inherently limited in representing dense subgraph structure: the max expected triangle count **shrinks cubically in overlap**!

This is very constraining. In this work, we explore what happens beyond edge independence.

Beyond Edge Independence

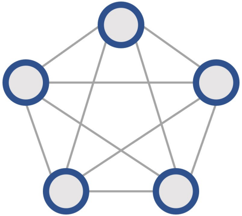

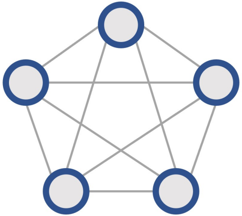

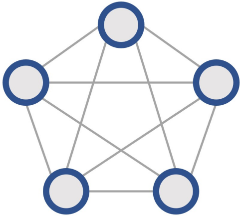

Beyond Edge Independence

Edge Activation

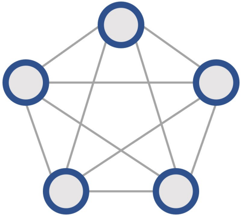

Node Activation

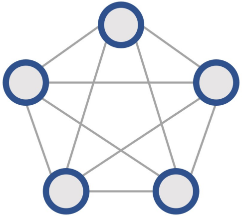

Graph Activation

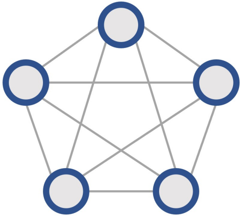

Beyond Edge Independence

Edge Activation

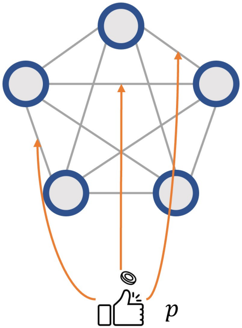

Node Activation

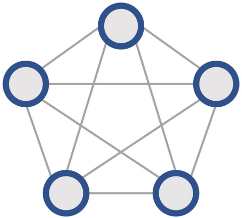

Graph Activation

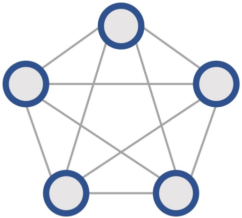

Beyond Edge Independence

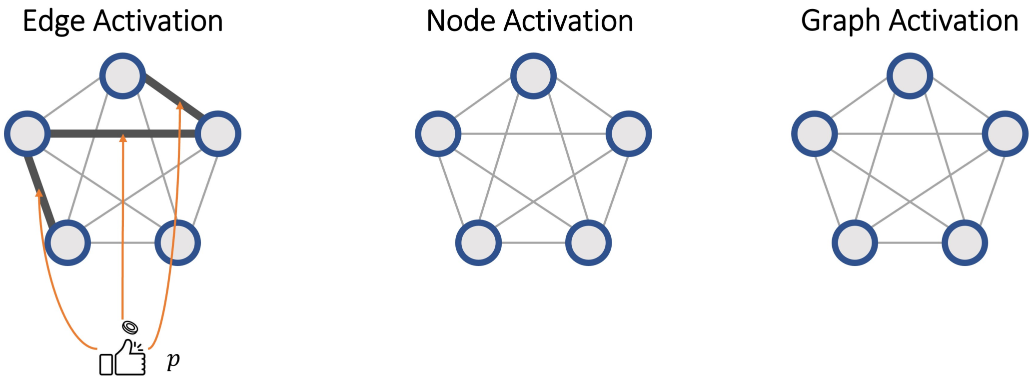

Beyond Edge Independence

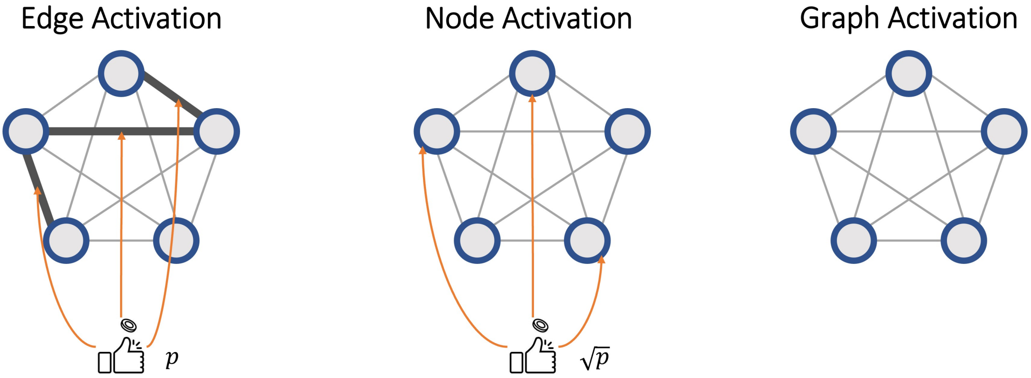

Beyond Edge Independence

Edge Activation

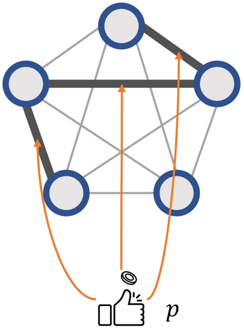

Node Activation

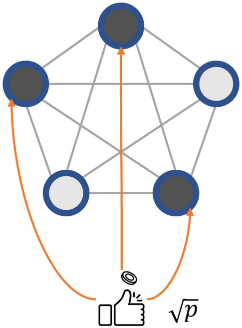

Graph Activation

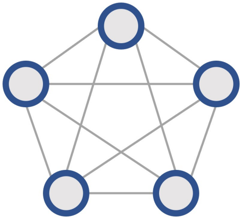

Beyond Edge Independence

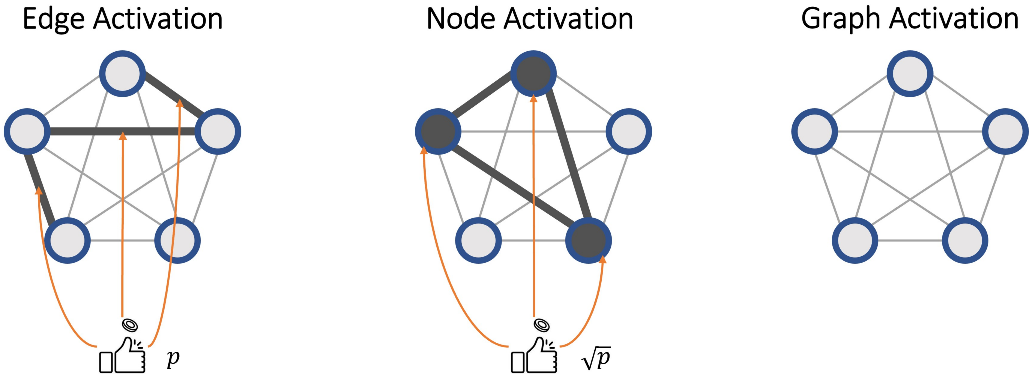

Beyond Edge Independence

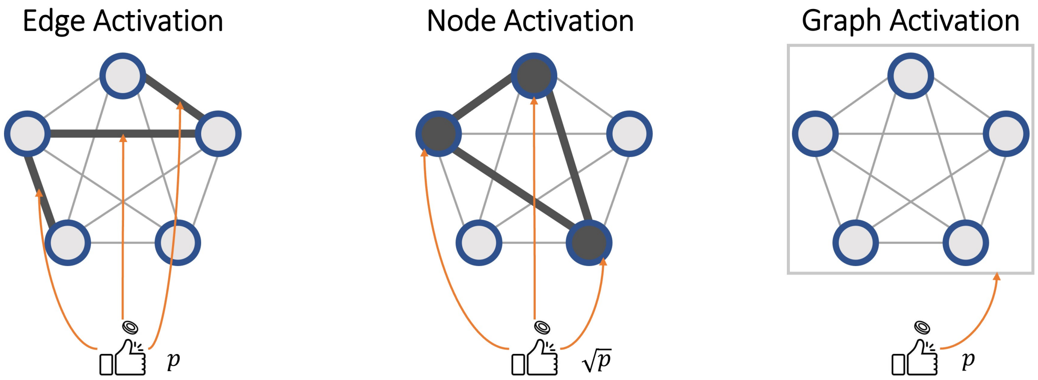

Beyond Edge Independence

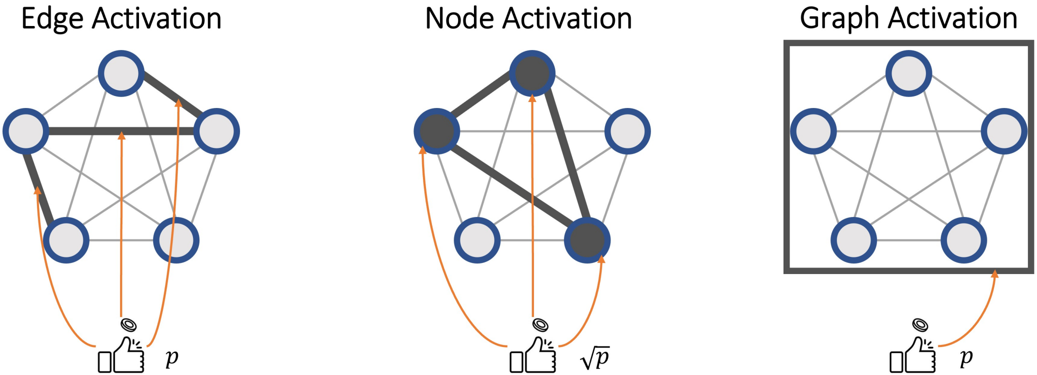

Beyond Edge Independence

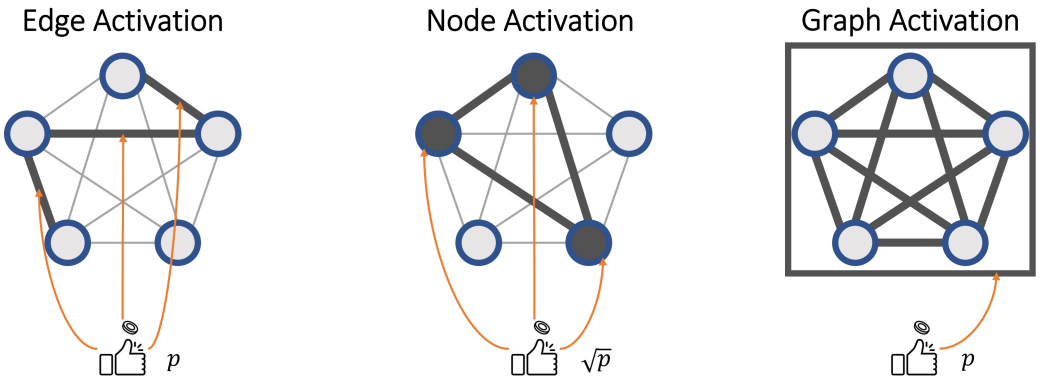

Beyond Edge Independence

Edge Activation

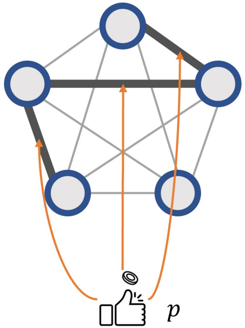

$\Pr(—) = p$

Node Activation

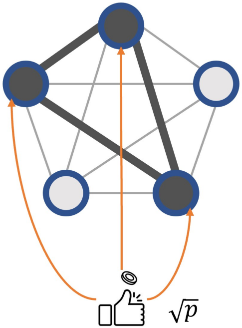

$\Pr(—) = p$

Graph Activation

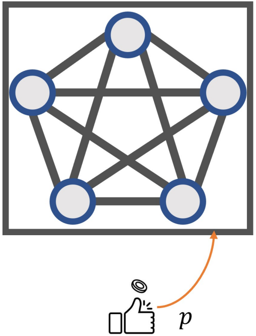

$\Pr(—) = p$

Beyond Edge Independence

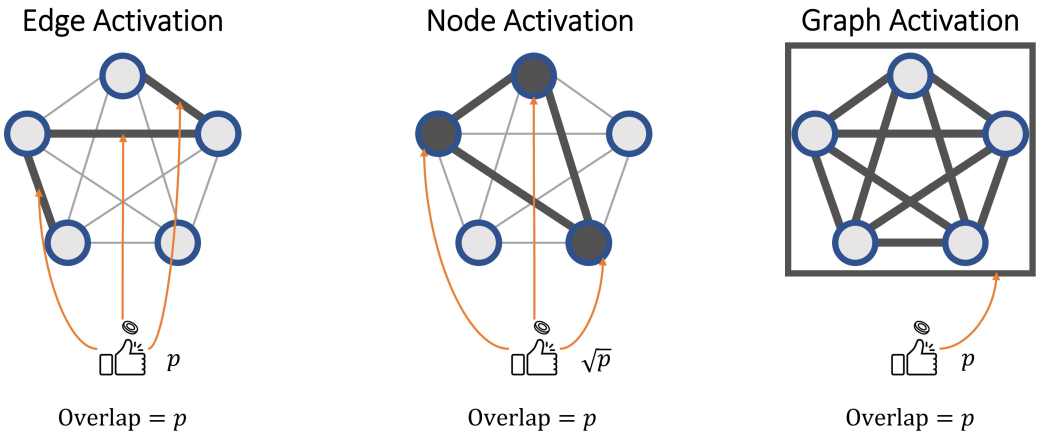

Beyond Edge Independence

Edge Activation

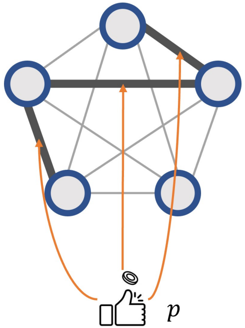

Overlap $= p$

$\Pr(\Delta) = p^3$

Node Activation

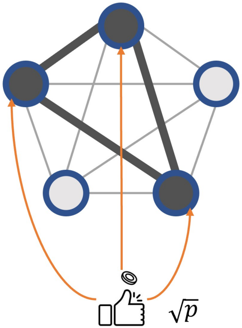

Overlap $= p$

$\Pr(\Delta) = p^{1.5}$

Graph Activation

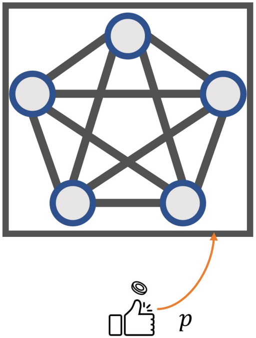

Overlap $= p$

$\Pr(\Delta) = p^1$

A Nested Hierarchy of Edge Dependence

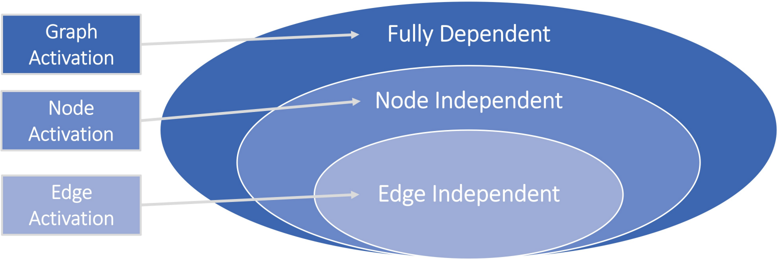

We introduce the *node independent* model: Each node is given an independent distribution over representation vectors.

Definition (Node Independent Graph Model)

Distribution over $\boldsymbol{A} = \sigma(\boldsymbol{X}\boldsymbol{Y}^\top)$ for some random node representation matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$, where $[\boldsymbol{X}; \boldsymbol{Y}]$'s rows are mutually independent.

A Nested Hierarchy of Edge Dependence

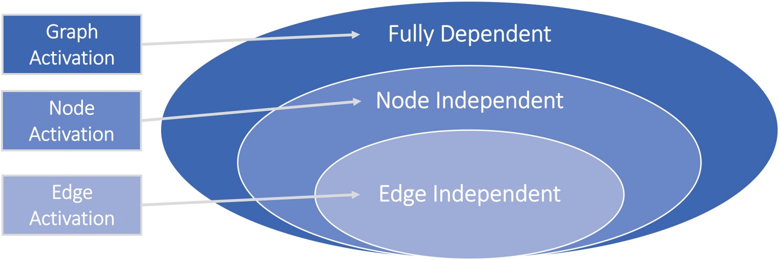

We introduce the *node independent* model: Each node is given an independent distribution over representation vectors.

Definition (Node Independent Graph Model)

Distribution over $\boldsymbol{A} = \sigma(\boldsymbol{X}\boldsymbol{Y}^\top)$ for some random node representation matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$, where $[\boldsymbol{X}; \boldsymbol{Y}]$'s rows are mutually independent.

A Nested Hierarchy of Edge Dependence

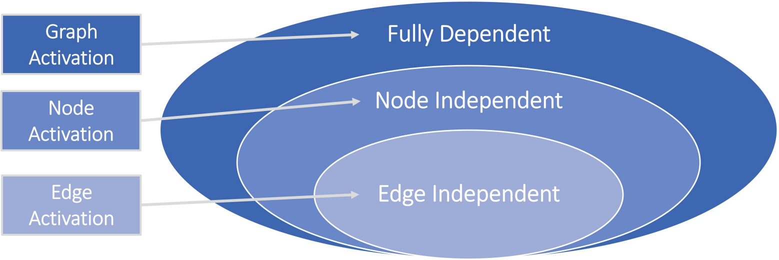

We introduce the *node independent* model: Each node is given an independent distribution over representation vectors.

Definition (Node Independent Graph Model)

Distribution over $\boldsymbol{A} = \sigma(\boldsymbol{X}\boldsymbol{Y}^\top)$ for some random node representation matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$, where $[\boldsymbol{X}; \boldsymbol{Y}]$'s rows are mutually independent.

A Nested Hierarchy of Edge Dependence

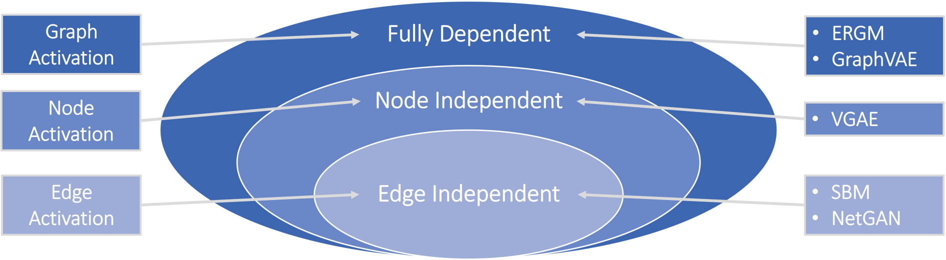

We introduce the *node independent* model: Each node is given an independent distribution over representation vectors.

Definition (Node Independent Graph Model)

Distribution over $\boldsymbol{A} = \sigma(\boldsymbol{X}\boldsymbol{Y}^\top)$ for some random node representation matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$, where $[\boldsymbol{X}; \boldsymbol{Y}]$'s rows are mutually independent.

Bounds for Edge Dependent Graph Models

We can prove bounds on the capabilities of models with edge dependency:

Model	Upper Bound on Δ/n^3	Examples
Edge Independent	$\mathrm{Ov}(\mathcal{A})^{3.0}$	SBM, NetGAN
Node Independent	$\mathrm{Ov}(\mathcal{A})^{1.5}$	VGAE
Fully Dependent	$\mathrm{Ov}(\mathcal{A})^{1.0}$	GraphVAE, ERGM

Note that these bounds match the {edge/node/graph} activation random graph families.
This means the above bounds are tight!

Bounds for Edge Dependent Graph Models

We can prove bounds on the capabilities of models with edge dependency:

Model	Upper Bound on Δ/n^3	Examples
Edge Independent	$\mathrm{Ov}(\mathcal{A})^{3.0}$	SBM, NetGAN
Node Independent	$\mathrm{Ov}(\mathcal{A})^{1.5}$	VGAE
Fully Dependent	$\mathrm{Ov}(\mathcal{A})^{1.0}$	GraphVAE, ERGM

Note that these bounds match the {edge/node/graph} activation random graph families.
This means the above bounds are tight!

Bounds for Edge Dependent Graph Models

We can prove bounds on the capabilities of models with edge dependency:

Model	Upper Bound on Δ/n^3	Examples
Edge Independent	$\text{Ov}(\mathcal{A})^{3.0}$	SBM, NetGAN
Node Independent	$\text{Ov}(\mathcal{A})^{1.5}$	VGAE
Fully Dependent	$\text{Ov}(\mathcal{A})^{1.0}$	GraphVAE, ERGM

Note that these bounds match the {edge/node/graph} activation random graph families. This means the above bounds are tight!

Empirical Contributions

We further explore empirical implications of our theoretical work:

- We motivate evaluation of graph generative models using not only error on matching graph statistics, but also overlap.
- We introduce three simple baselines inspired by {edge/node/graph} activation.
- We compare our baselines against modern deep-learning based models, and we find they are often competitive at matching graph statistics at the same levels of overlap.

On the Role of Edge Dependency in Graph Generative Models

Sudhanshu Chanpuriya, Cameron Musco,
Konstantinos Sotiropoulos, Charalampos Tsourakakis