On Estimating Link Prediction Uncertainty using Stochastic Centering

Puja Trivedi

Danai Koutra

Jay Thiagarajan

Link Prediction

- Graph Neural Networks (GNNs) are used for ***link prediction*** in high impact tasks.

Product Recommendation

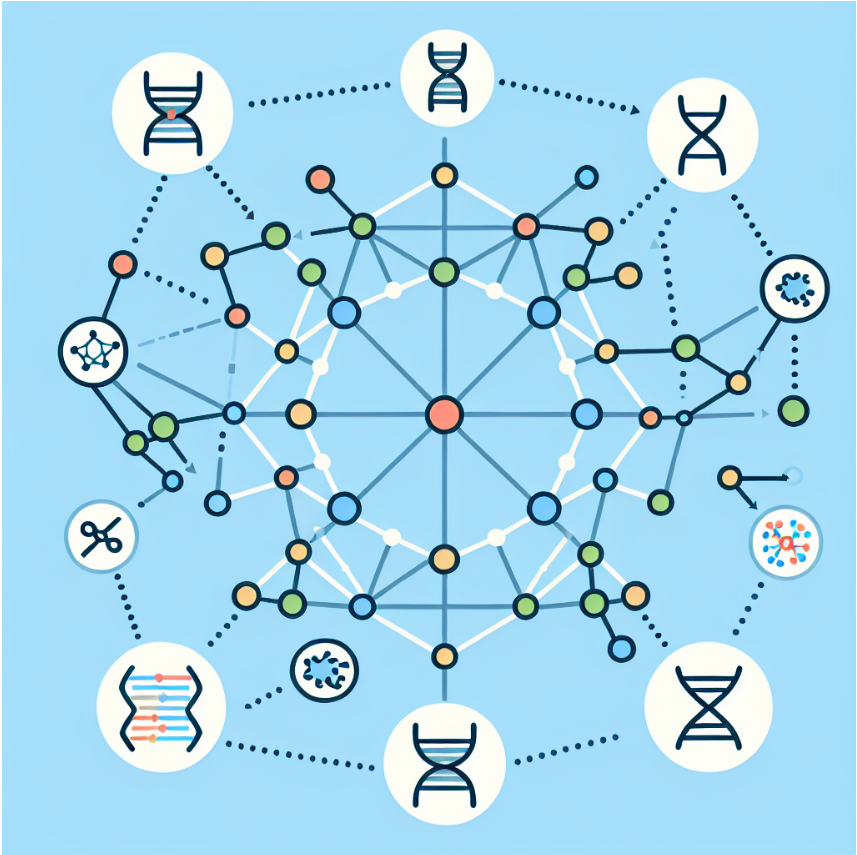

Gene-Gene Interaction

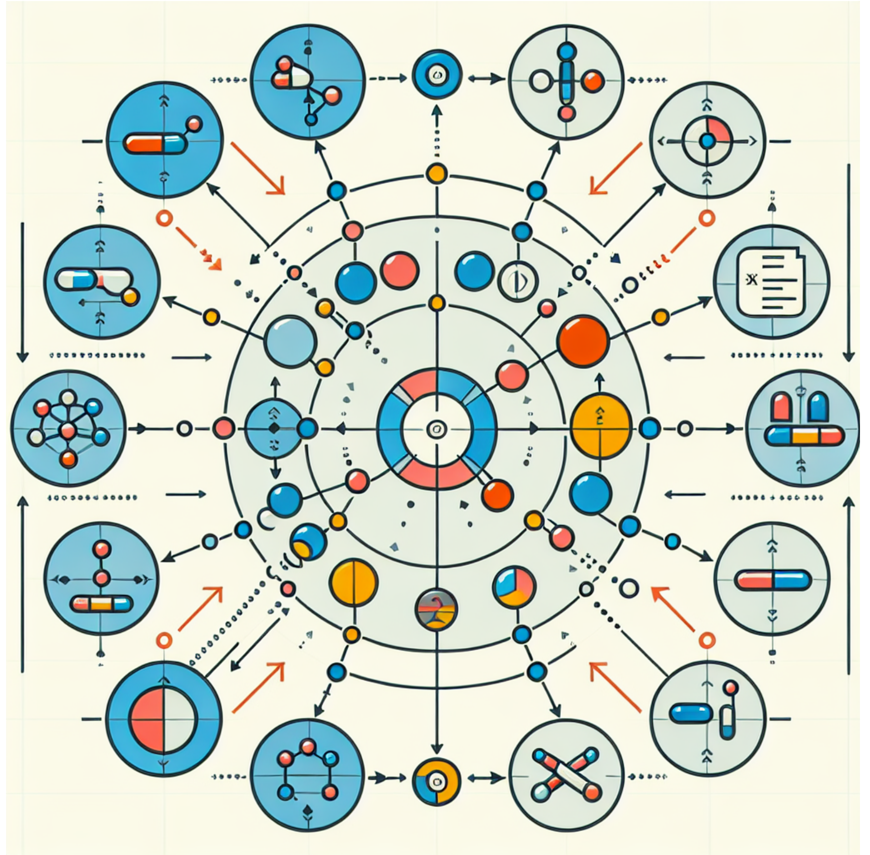

Drug-Drug Interaction

Link Prediction

- Graph Neural Networks (GNNs) are used for ***link prediction*** in high impact tasks.

Product Recommendation

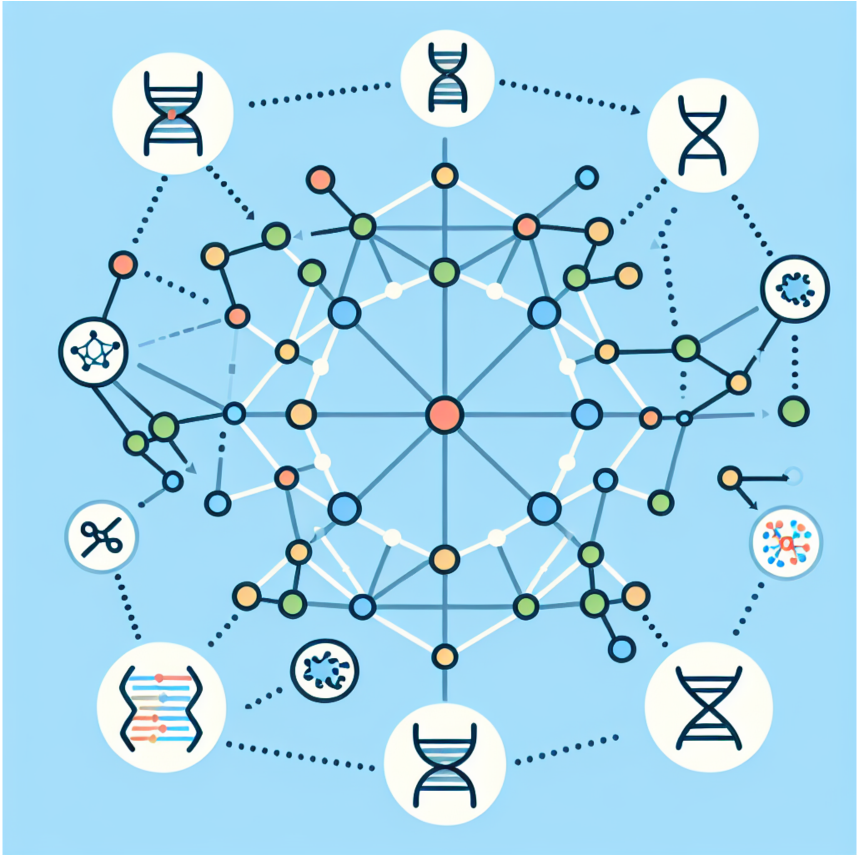

Gene-Gene Interaction

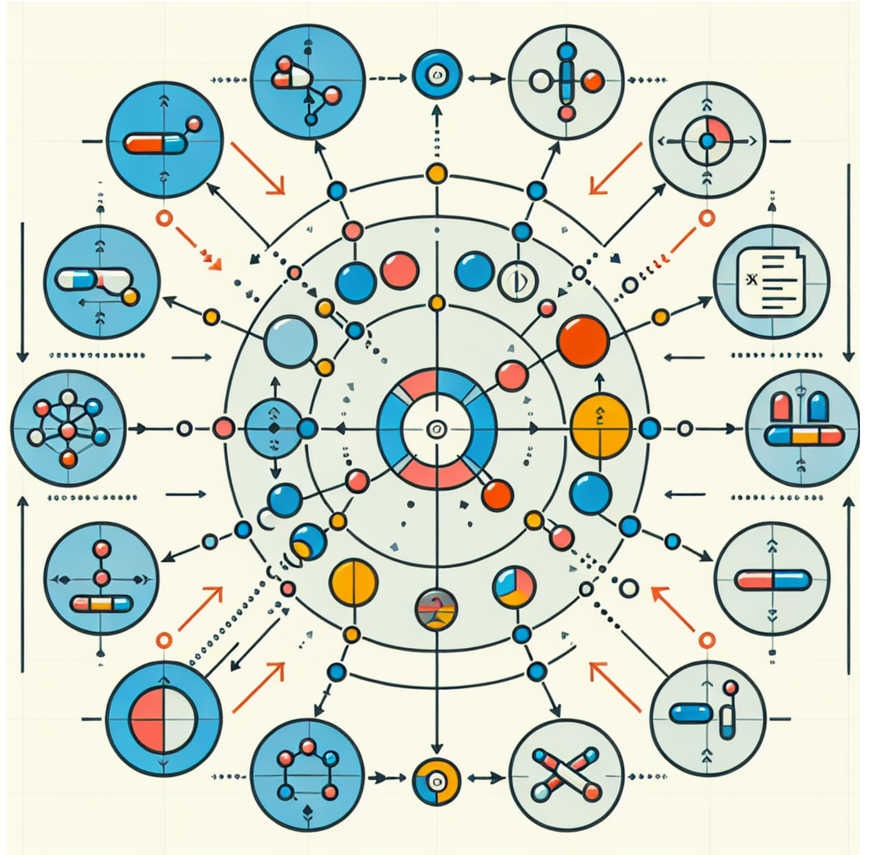

Drug-Drug Interaction

- These predictions are used to invoke expensive or time-consuming actions.
- Understanding prediction ***confidence*** is important for making informed decisions.

Calibration for Trust Worthy Predictions

- ***Calibration*** is the process of adjusting a model's ***output probabilities*** to ensure that they accurately reflect the ***true likelihood*** associated with a specific prediction.

$$\mathop{\mathbb{E}}_{\hat{P}}\left[\left|\mathbb{P}\left(\hat{Y}=Y \mid \hat{P}=p\right)-p\right|\right]$$

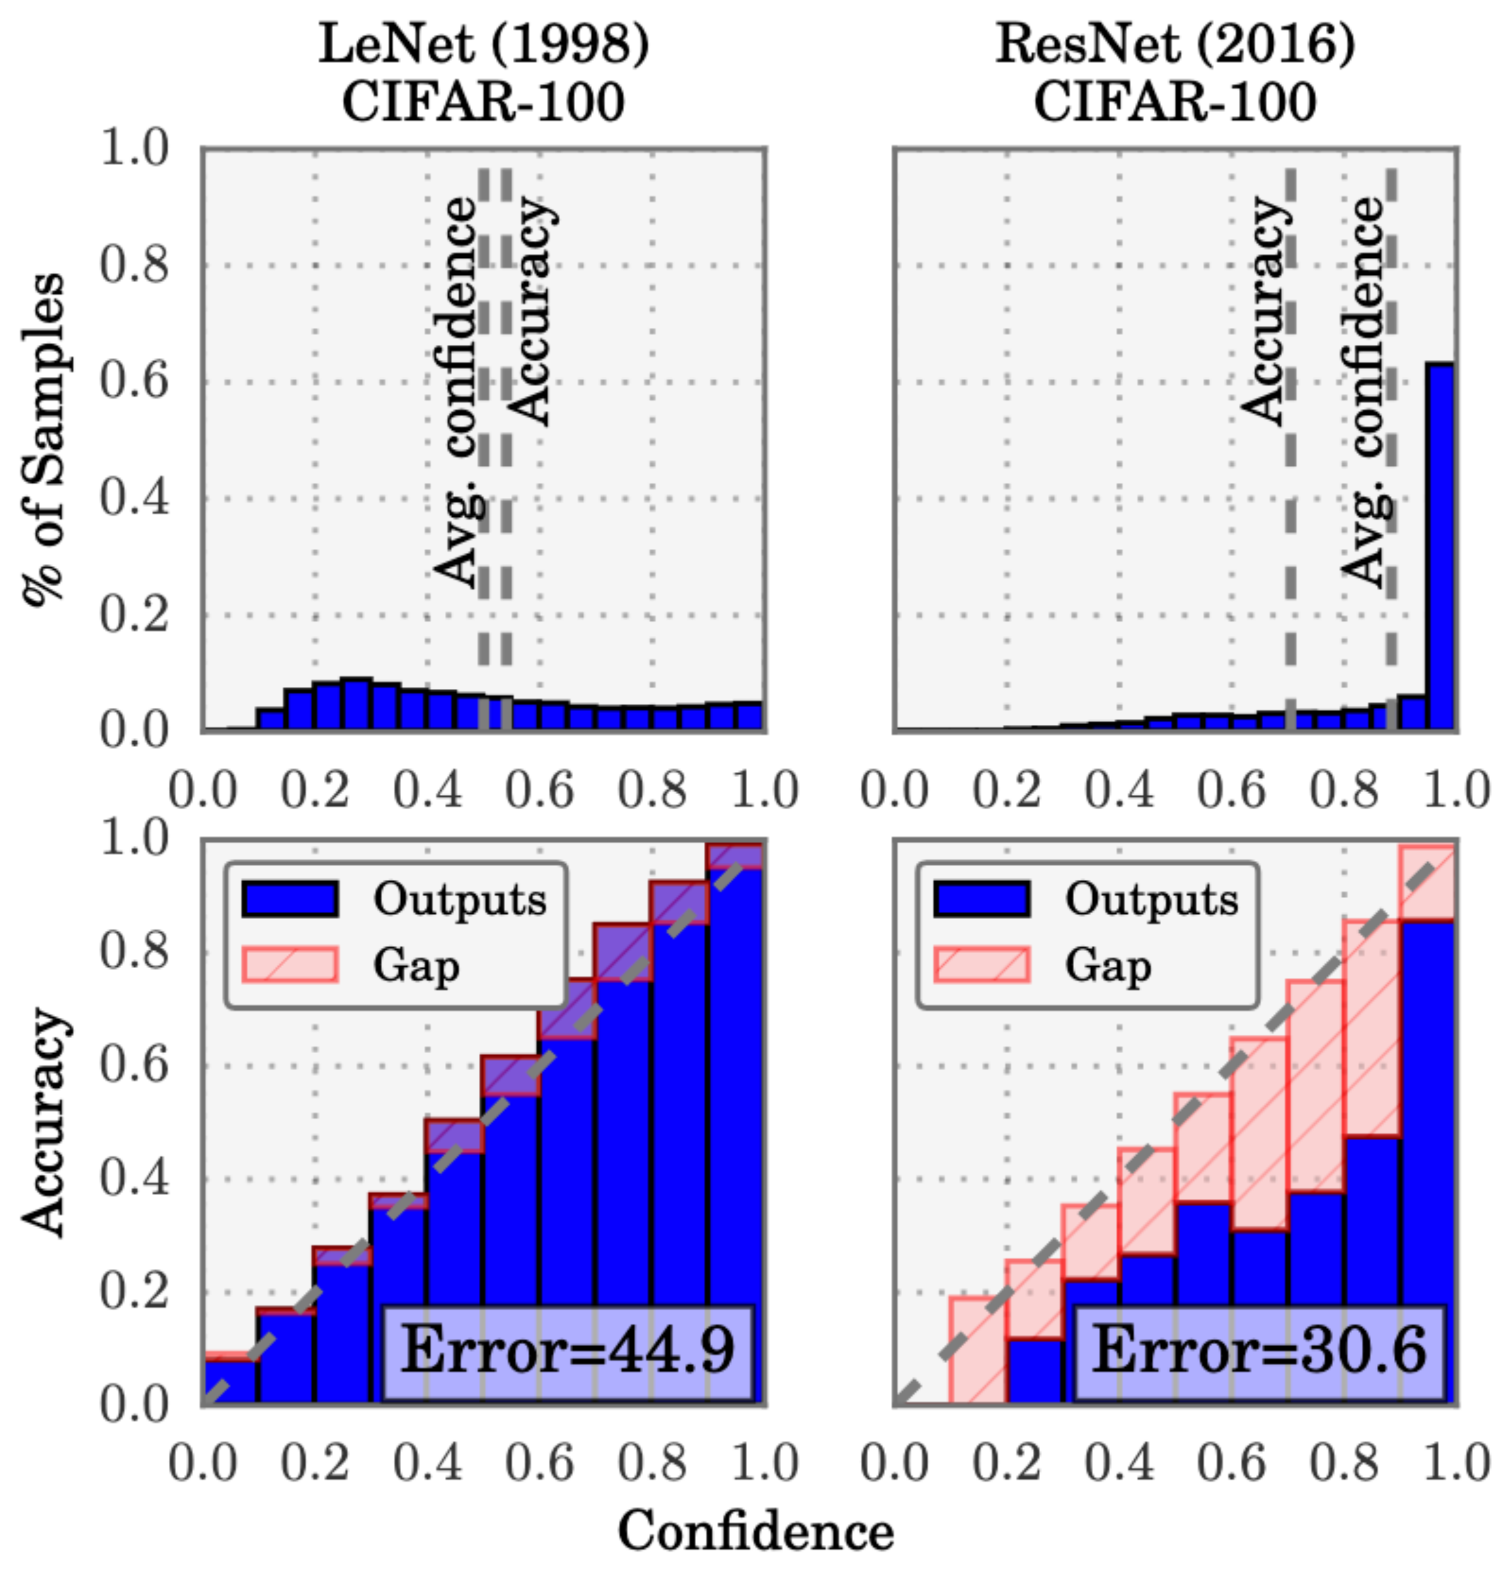

Calibration for Trust Worthy Predictions

- ***Calibration*** is the process of adjusting a model's ***output probabilities*** to ensure that they accurately reflect the ***true likelihood*** associated with a specific prediction.

$$\mathbb{E}_{\hat{P}}\left[\left|\mathbb{P}\left(\hat{Y}=Y \mid \hat{P}=p\right)-p\right|\right]$$

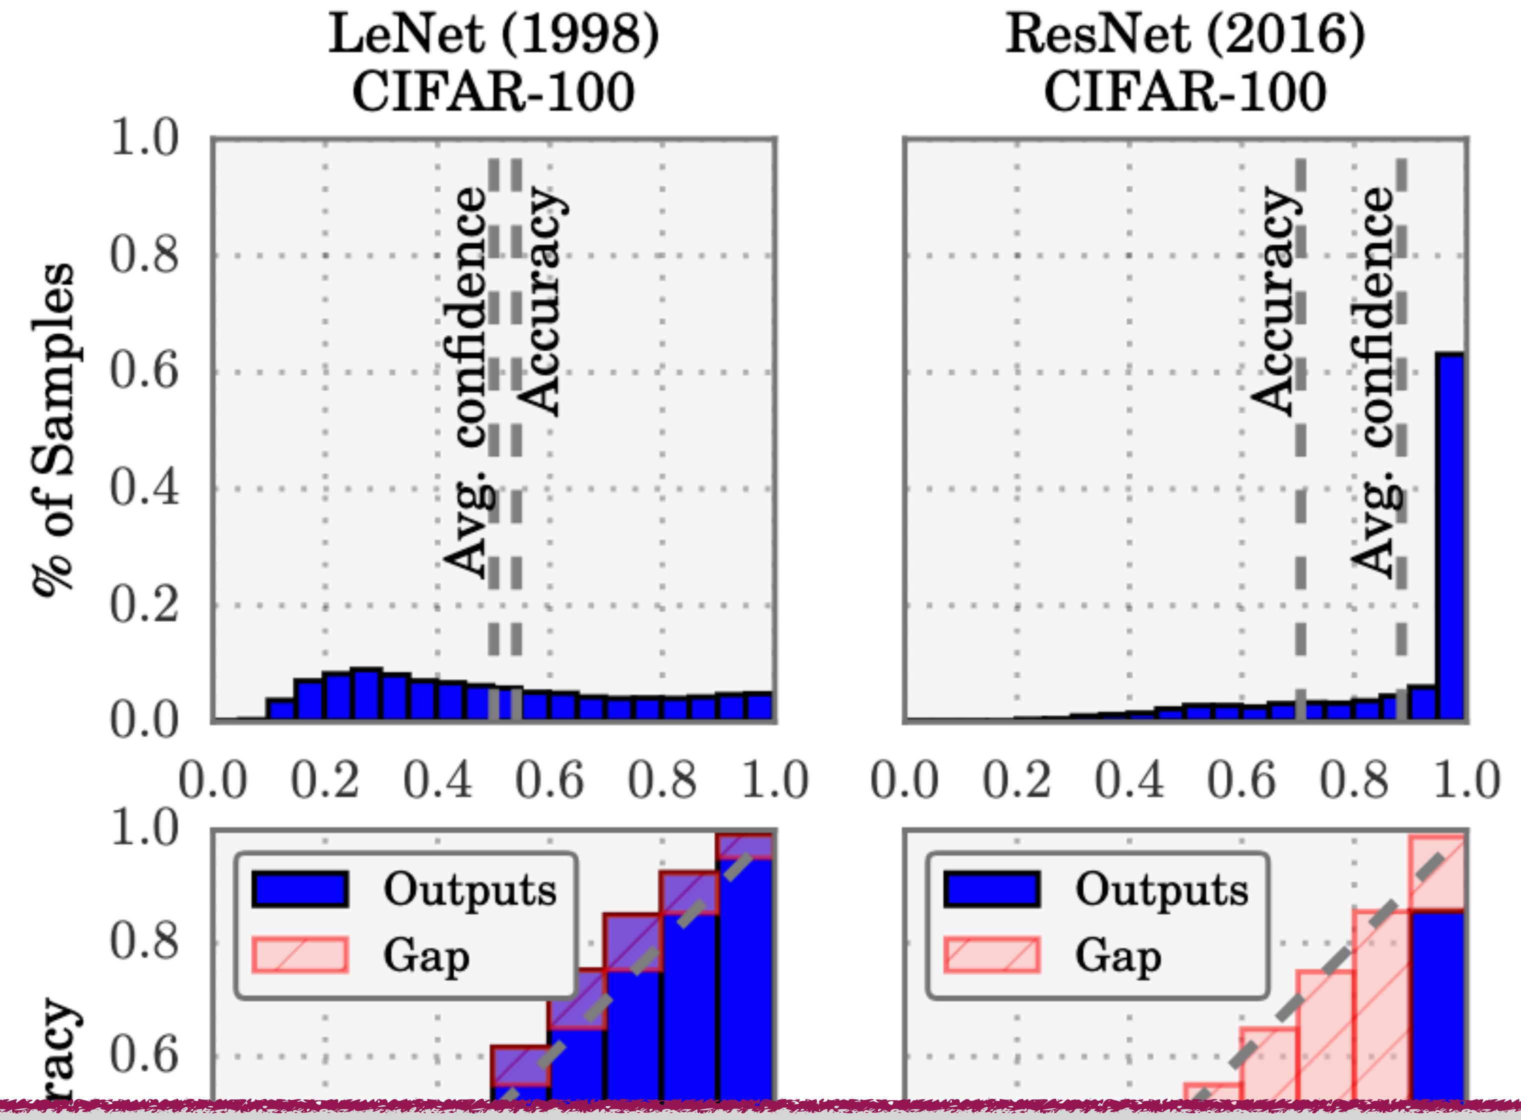

However, calibration for link prediction remains understudied!

Why is Calibration for Link Prediction Challenging?

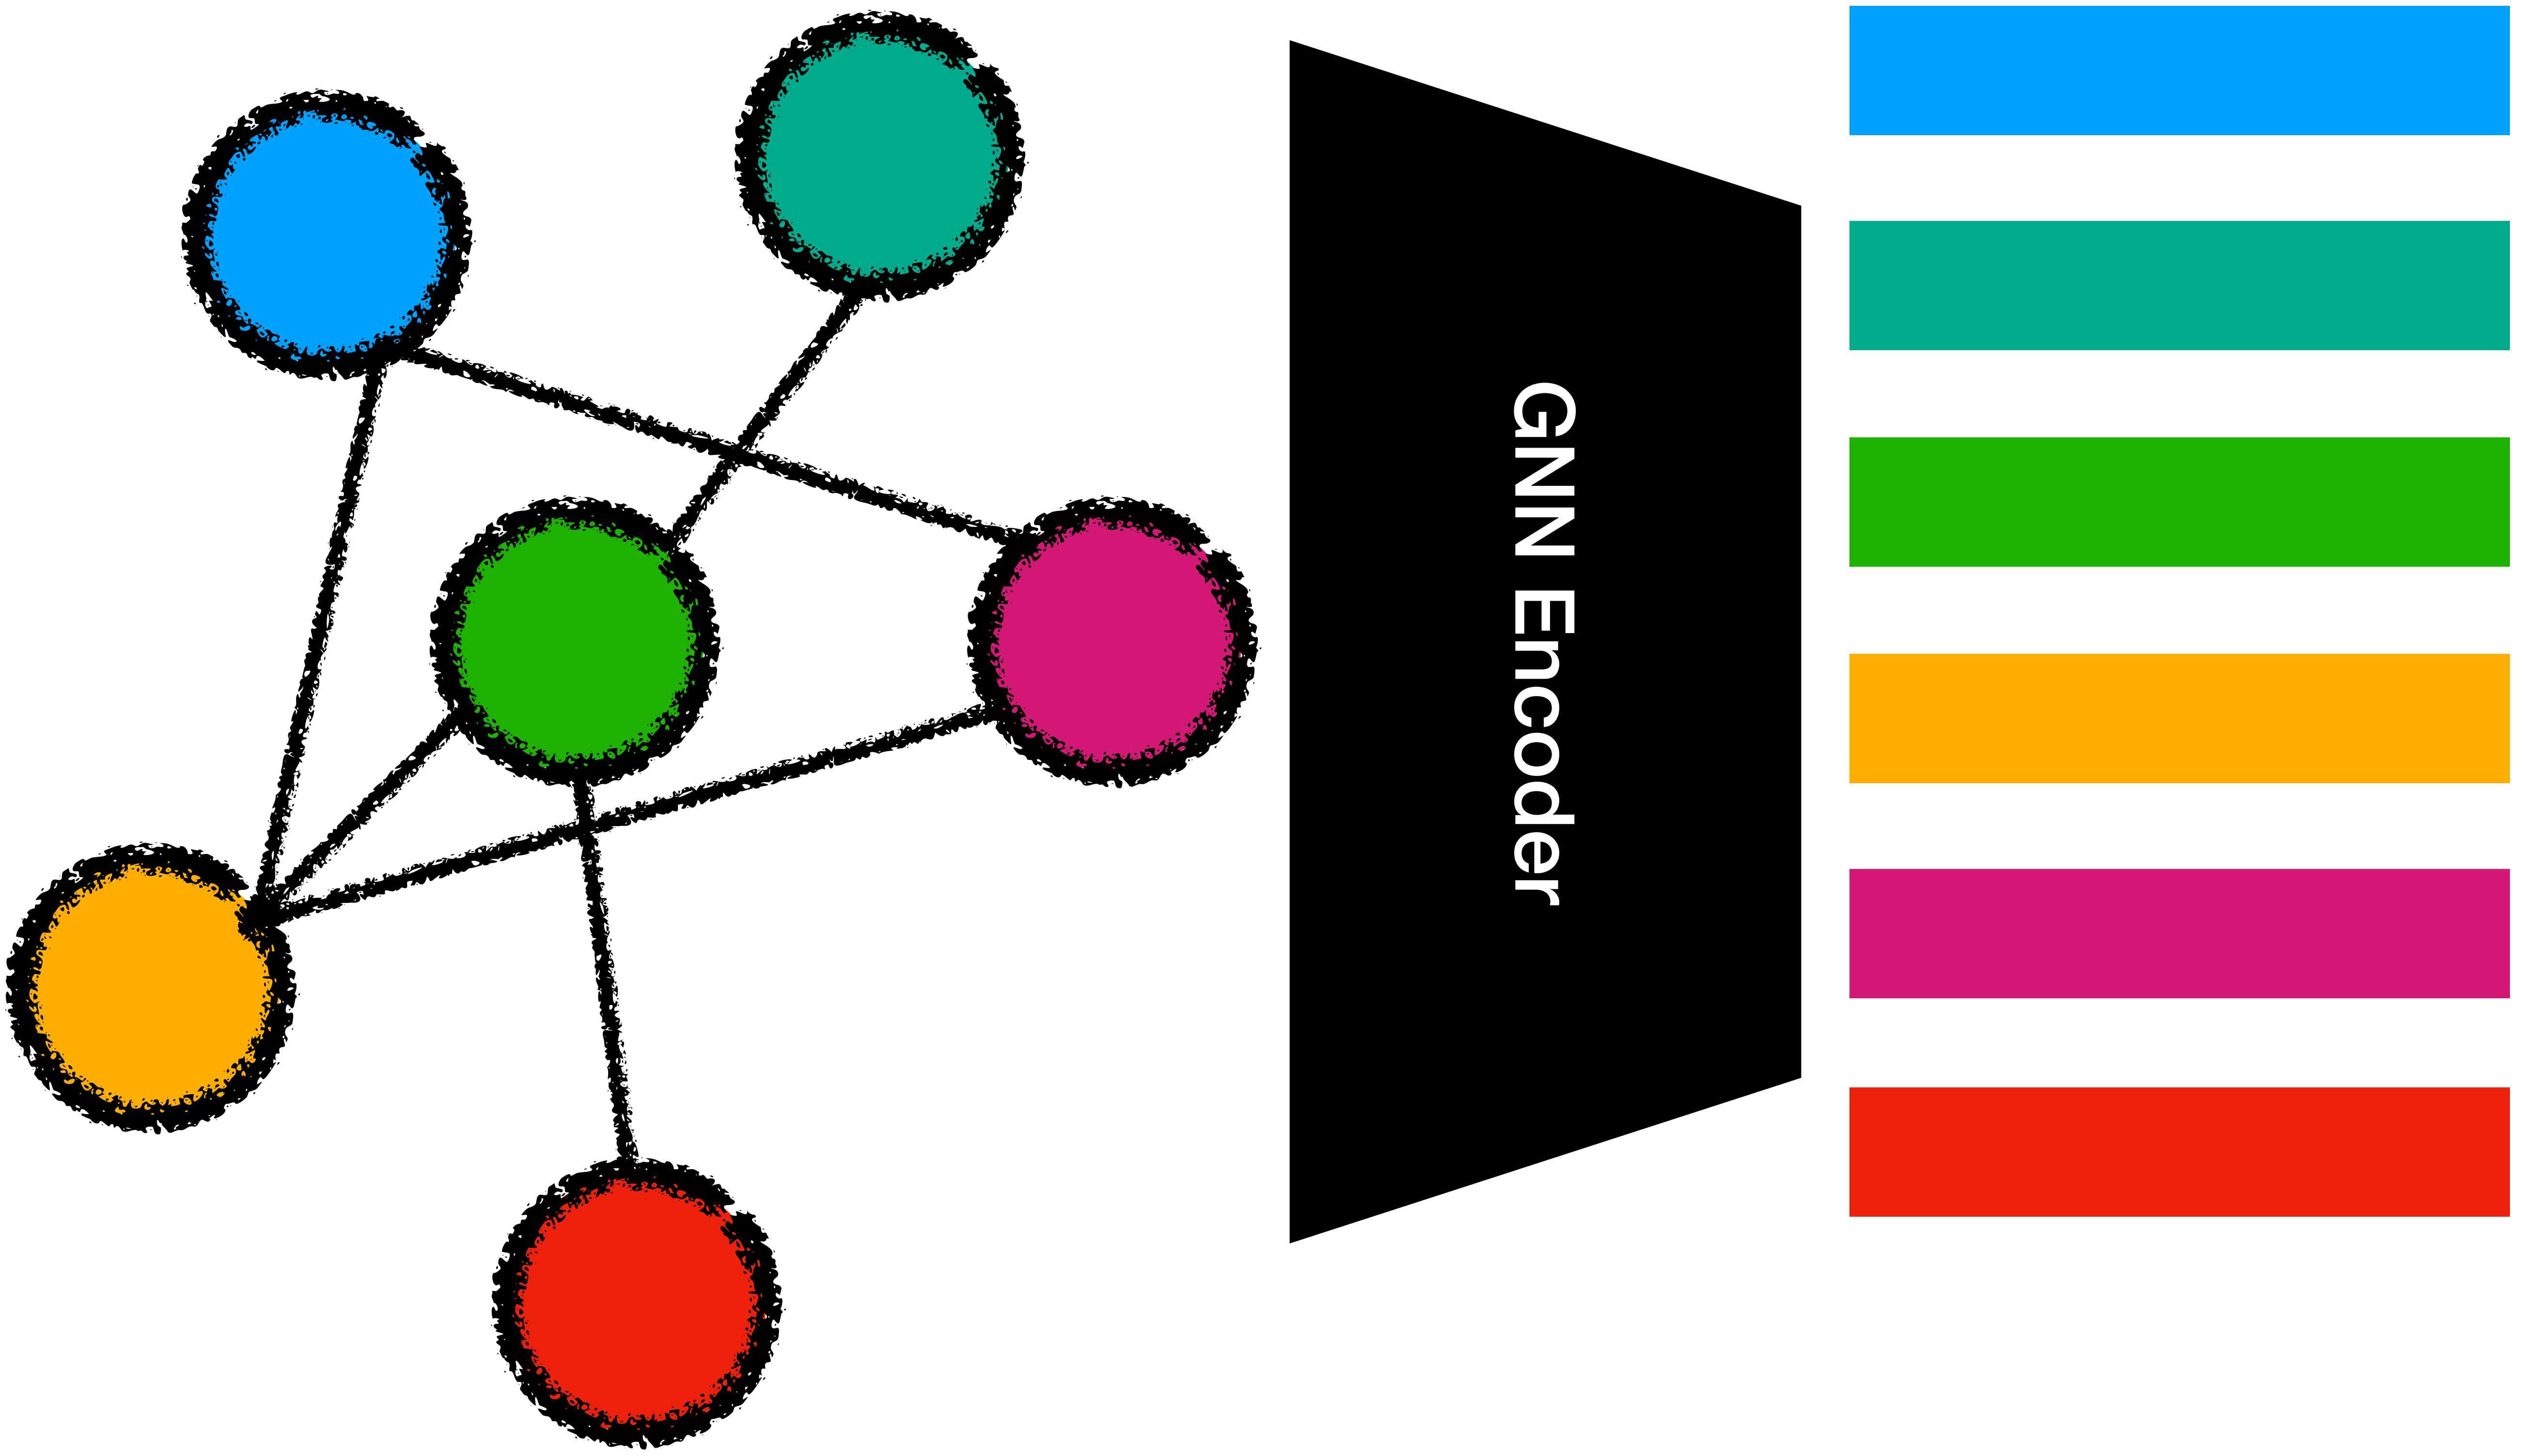

Why is Calibration for Link Prediction Challenging?

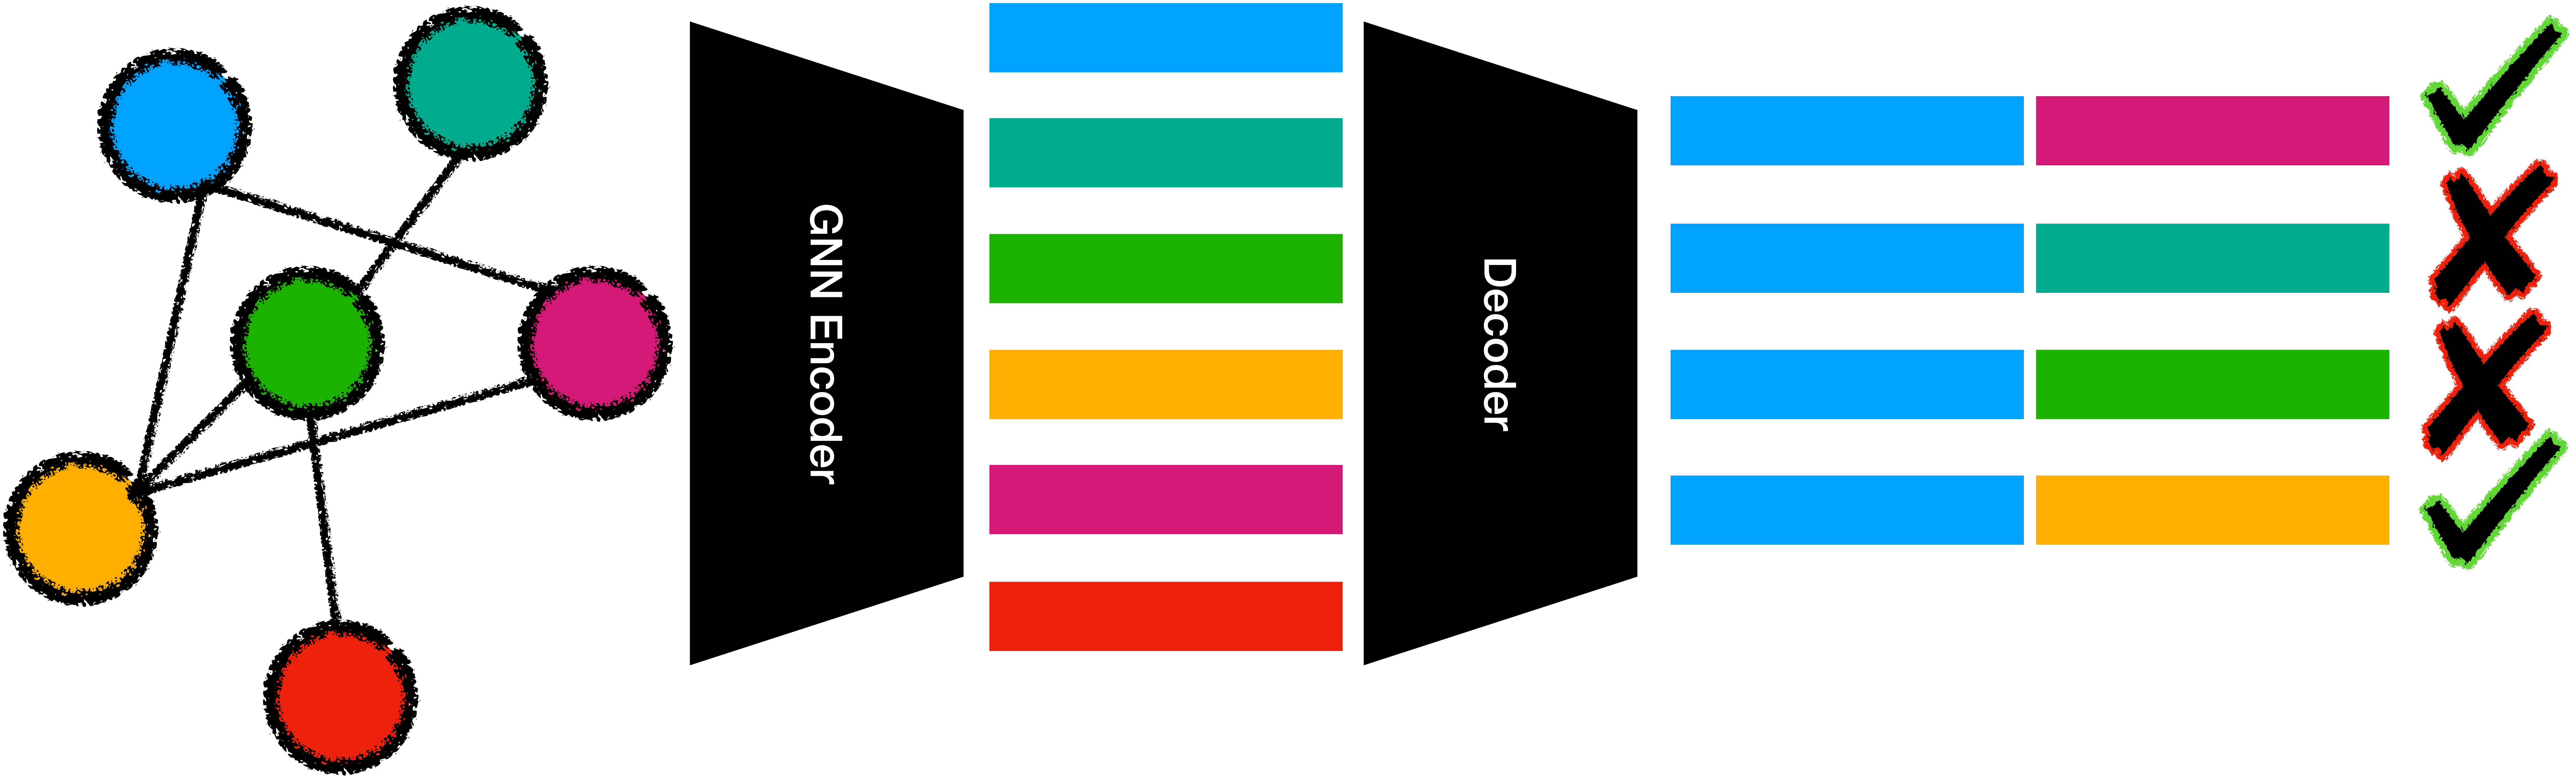

Why is Calibration for Link Prediction Challenging?

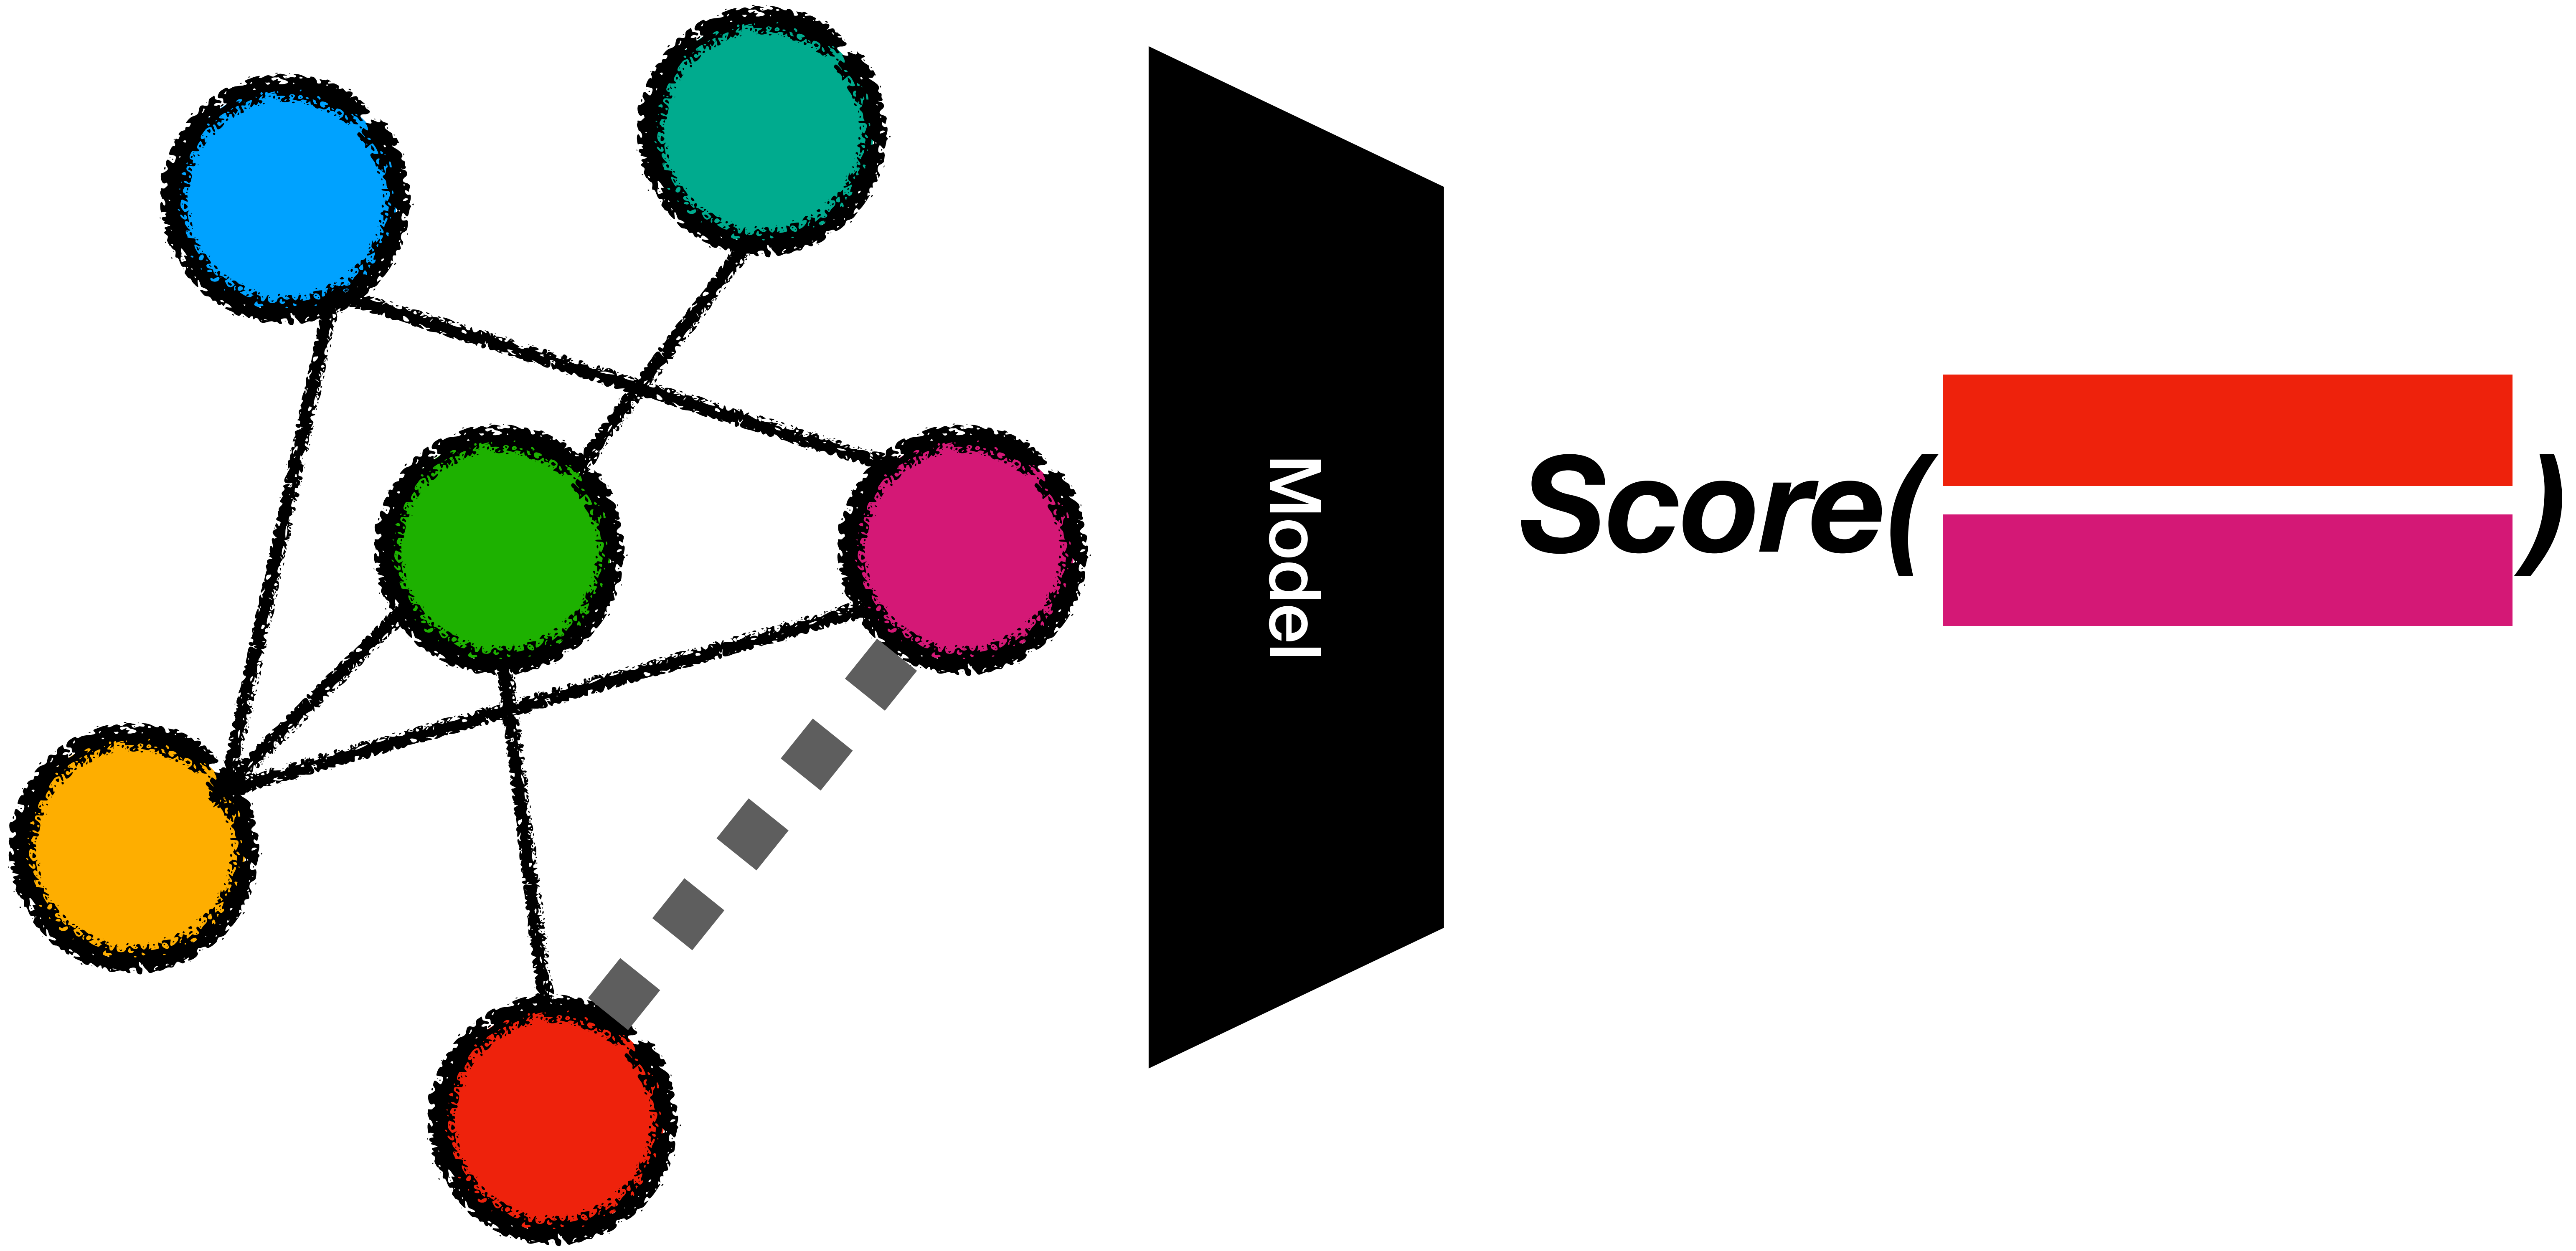

Why is Calibration for Link Prediction Challenging?

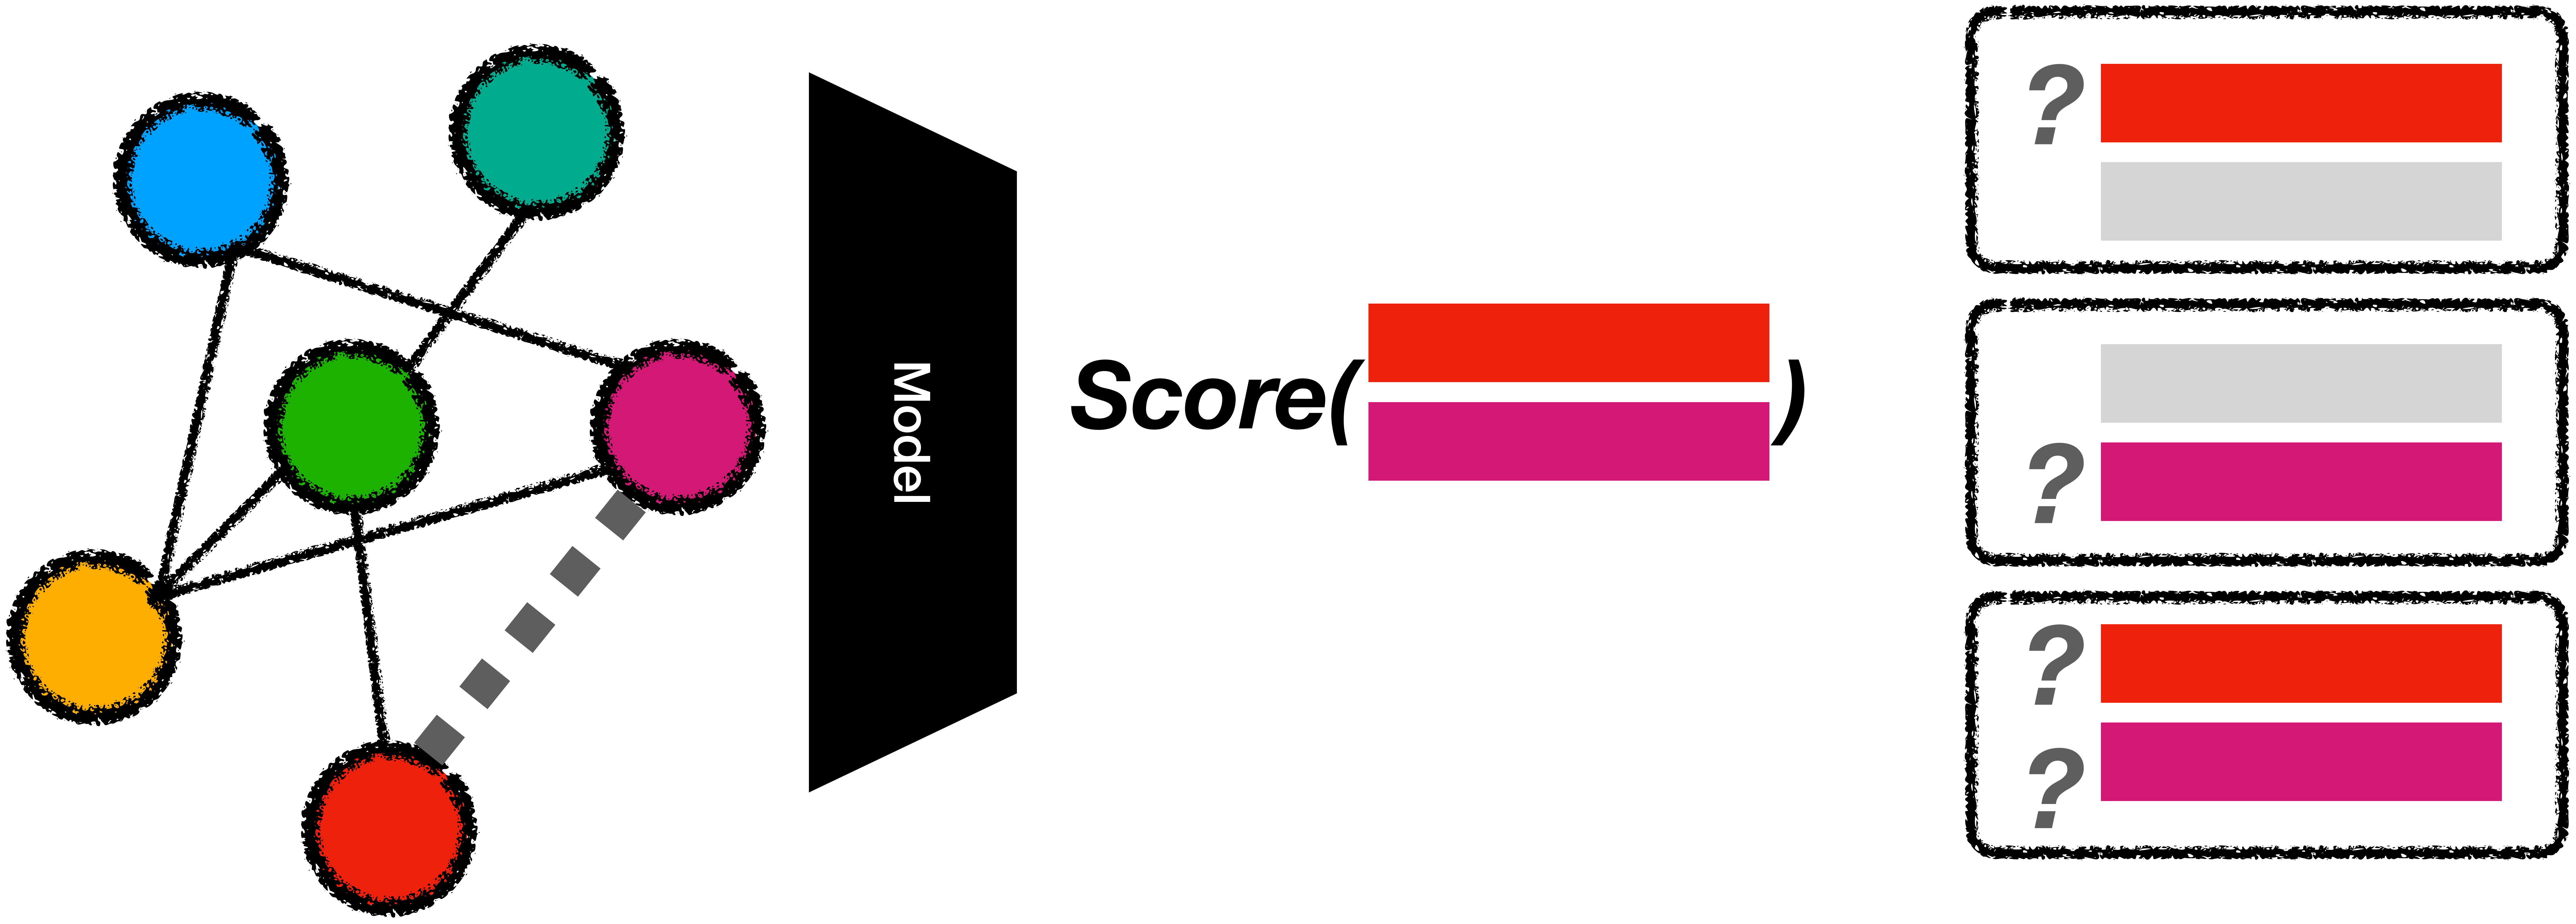

Why is Calibration for Link Prediction Challenging?

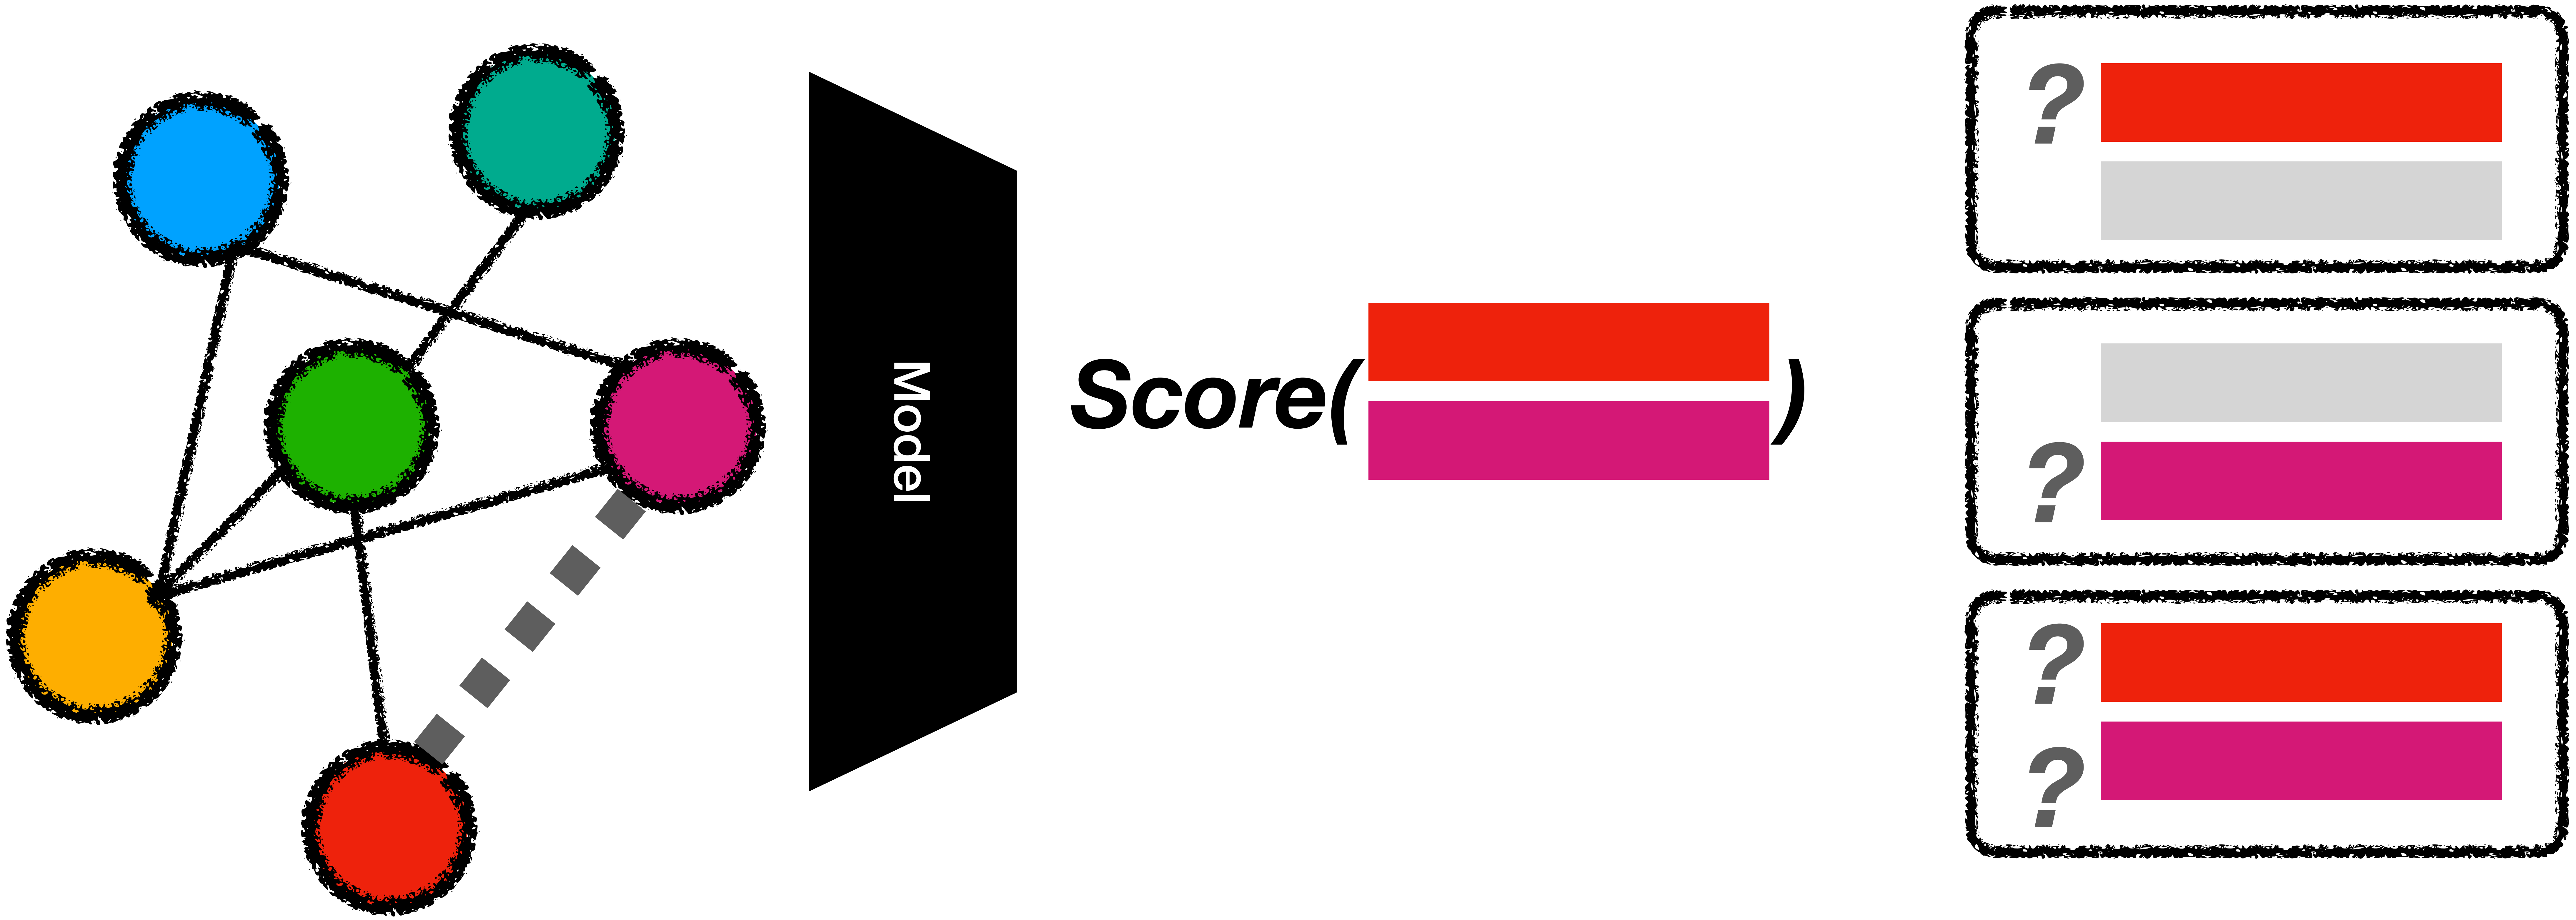

- *Model's uncertainty in its prediction is a function of its node-level uncertainties.*
- *Existing calibration methods do not take individual node's into account.*

Our Contributions

1. Extending Stochastic Centering to Edge-Level Uncertainty

2. Creating Meaningful Node-level uncertainties

3. Experimental Evaluation of Edge-ΔUQ

Our Contributions

1. Extending Stochastic Centering to Edge-Level Uncertainty

2. Creating Meaningful Node-level uncertainties.

3. Experimental Evaluation of Edge-ΔUQ

What is Stochastic Centering (ΔUQ)?

- Stochastic Centering uses ***anchoring*** to simulate the ***behavior of an ensemble*** using only a ***single trained model***.

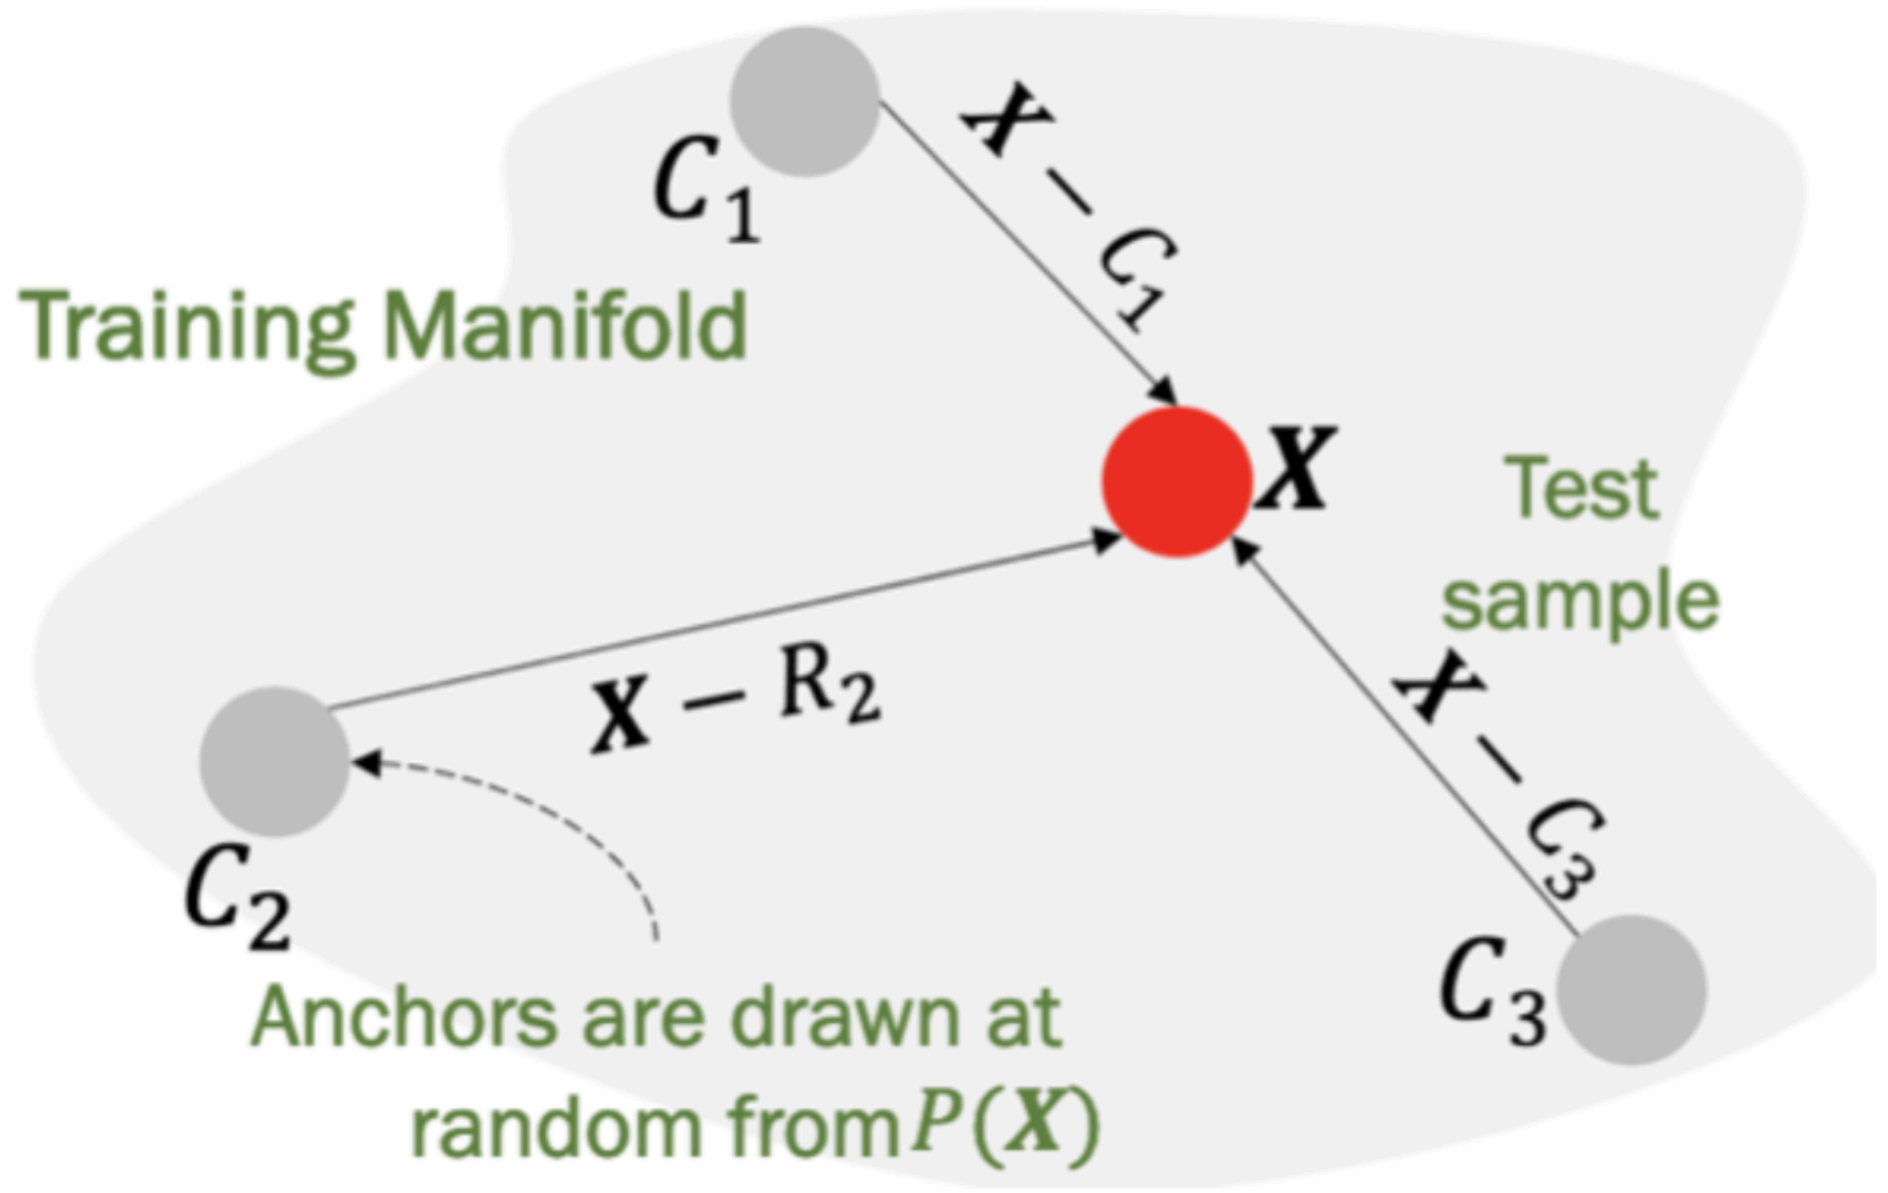

What is Stochastic Centering (ΔUQ)?

- Stochastic Centering uses ***anchoring*** to simulate the ***behavior of an ensemble*** using only a ***single trained model***.
- Anchoring creates a ***relative representation*** for an input sample ***x*** in terms of a random anchor ***c: [x – c, c].***

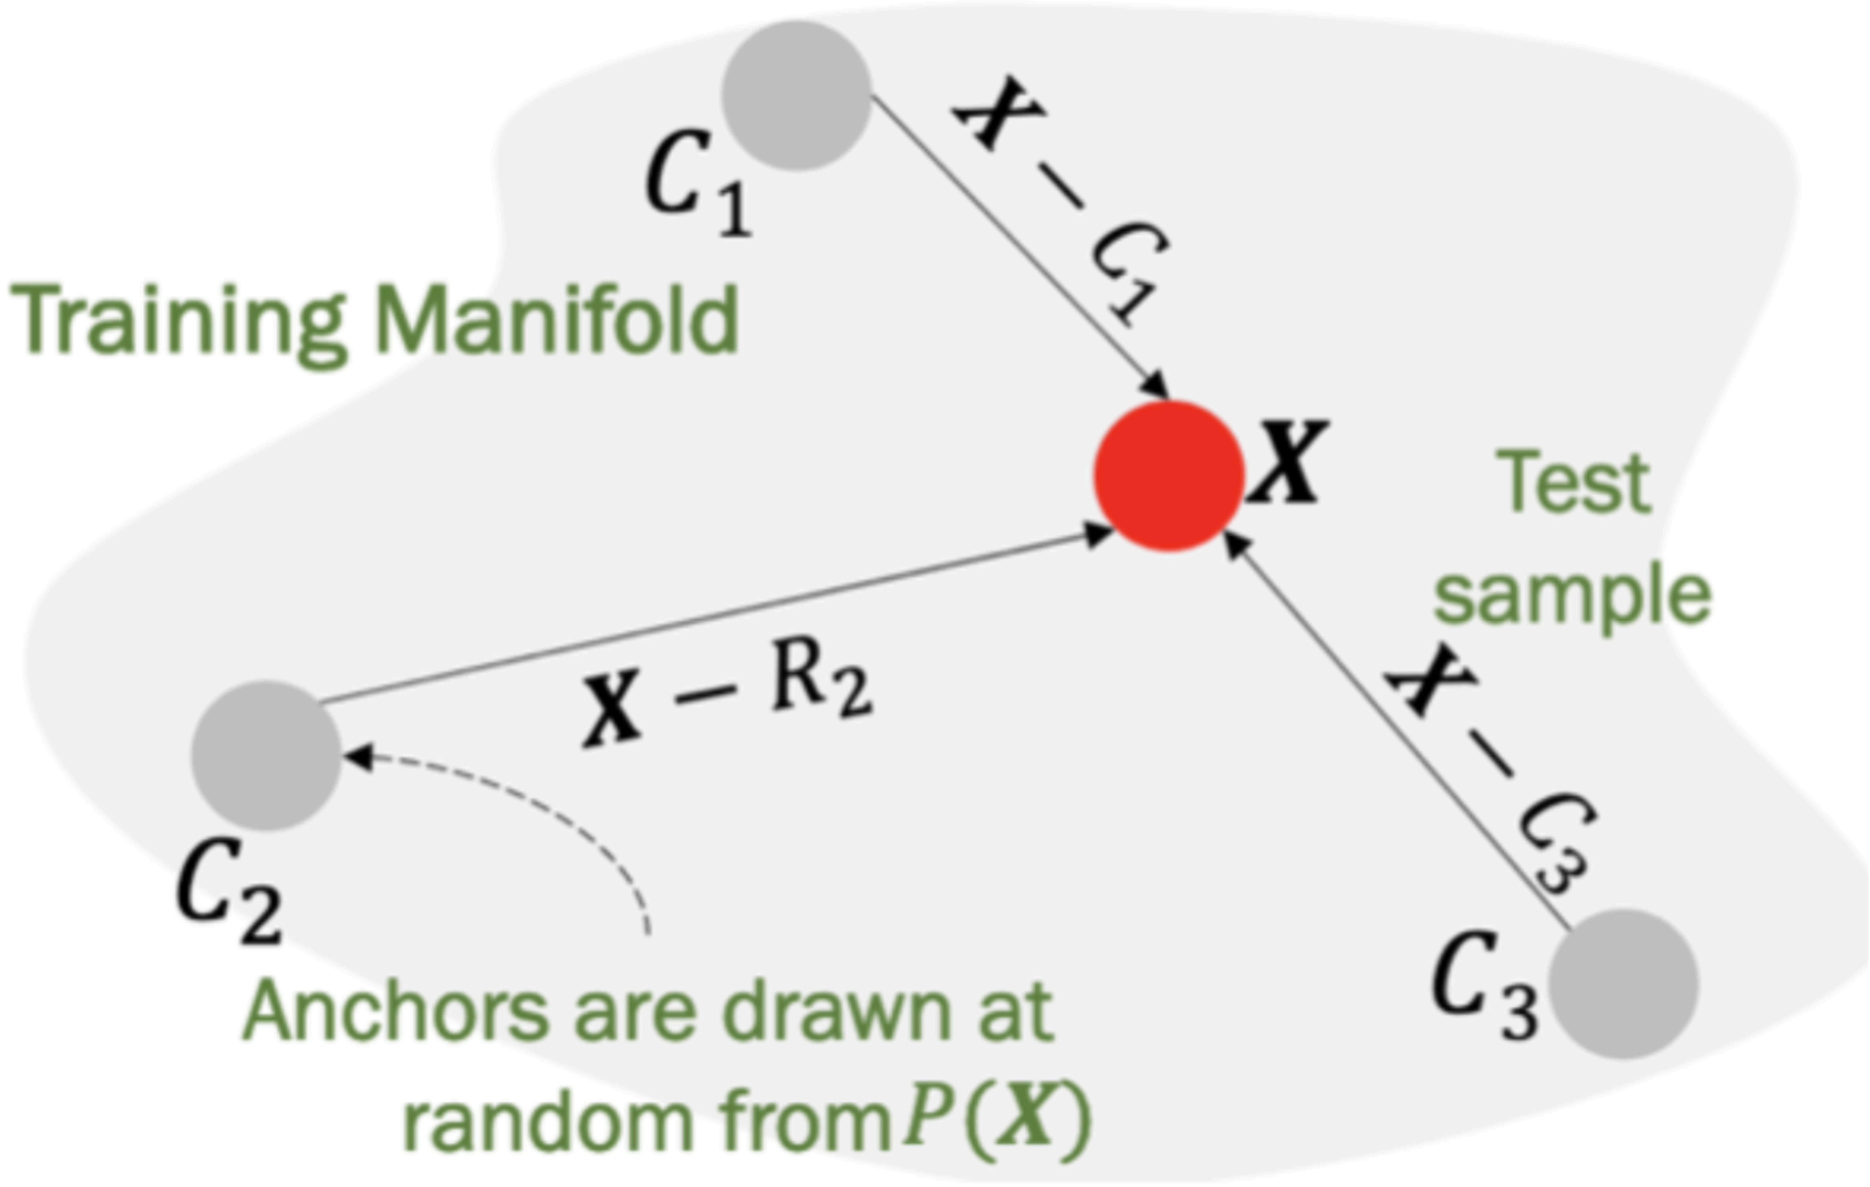

What is Stochastic Centering(ΔUQ)?

- Stochastic Centering uses ***anchoring*** to simulate the ***behavior of an ensemble*** using only a ***single trained model***.
- Anchoring creates a ***relative representation*** for an input sample ***x*** in terms of a random anchor ***c: [x – c, c].***
- During training, the anchor is ***randomized*** emulates the process of ***sampling different solutions*** from the hypothesis space.

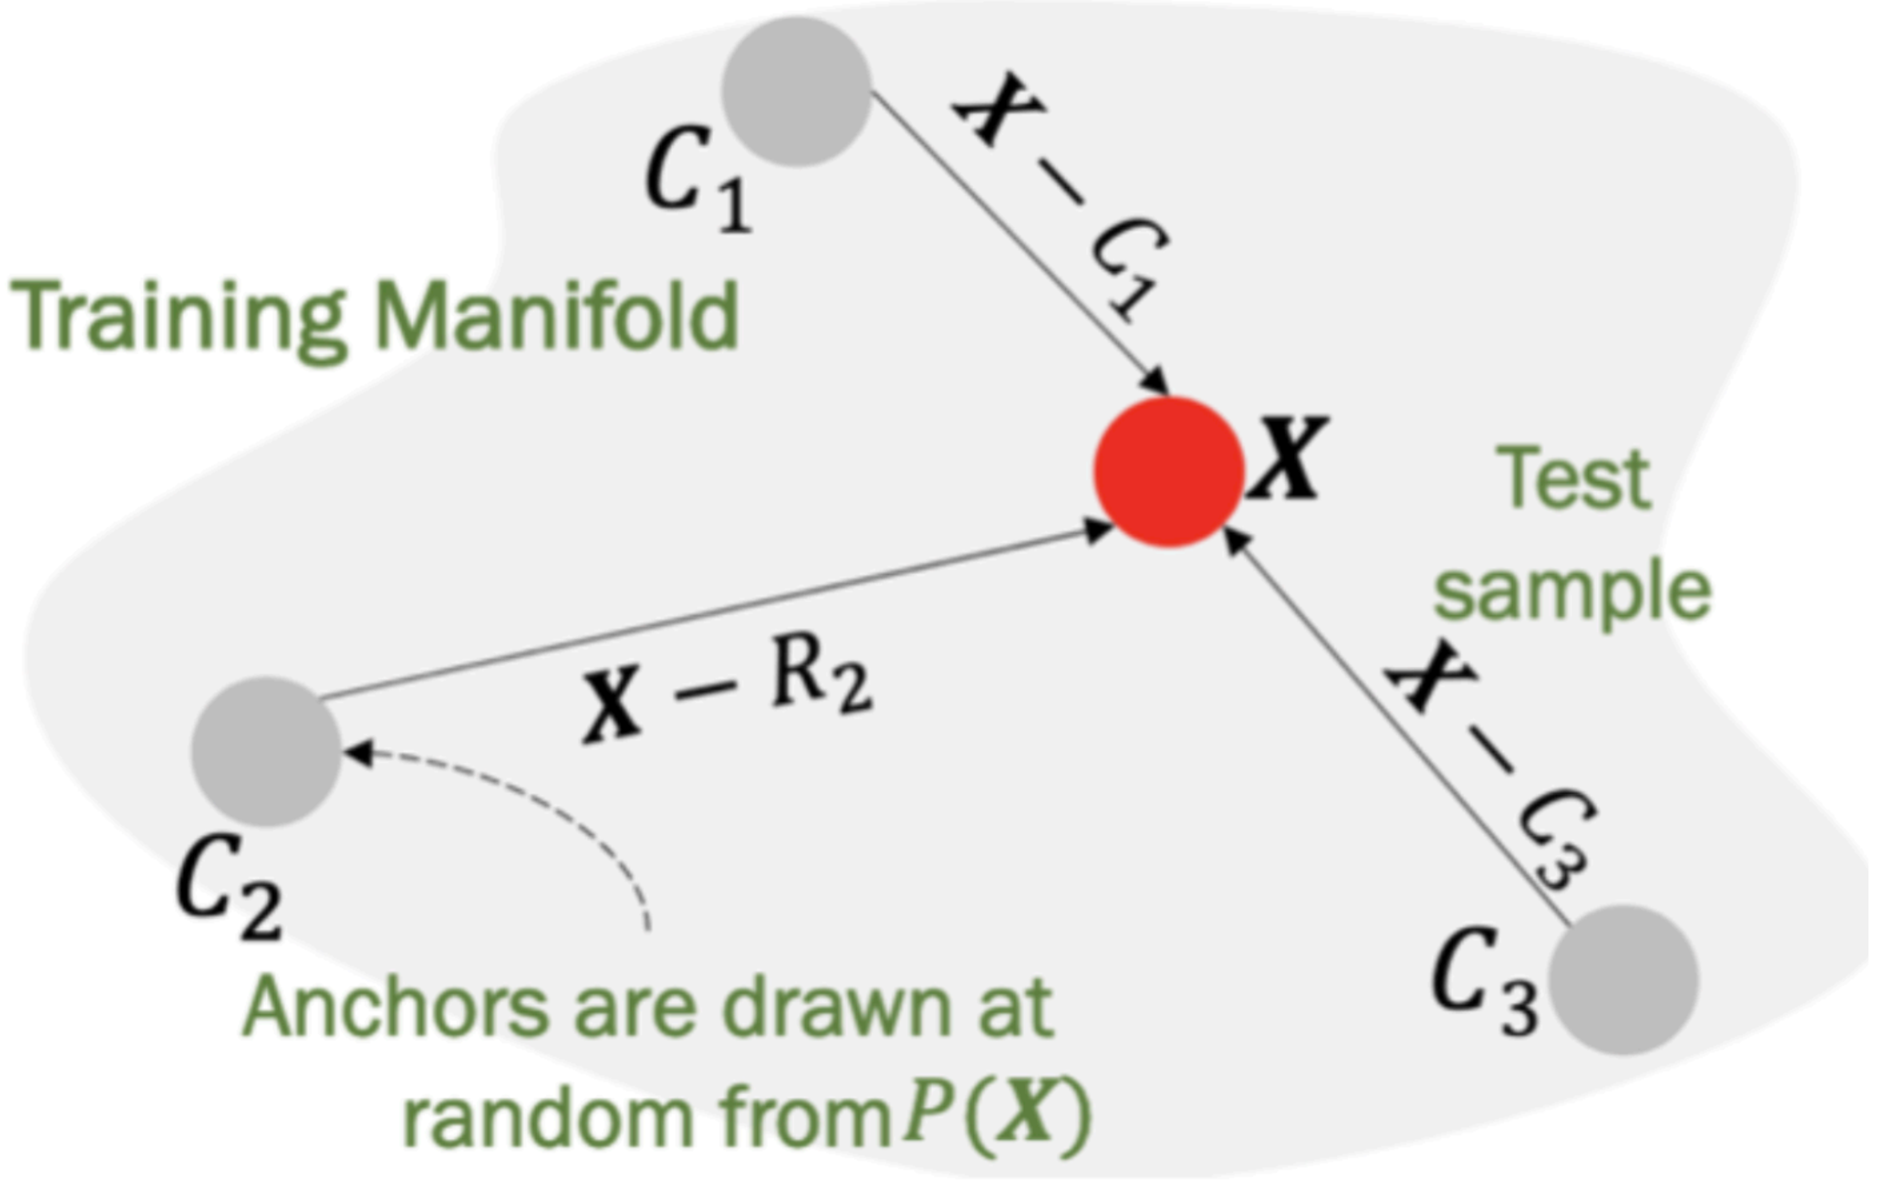

What is Stochastic Centering(ΔUQ)?

- Like a Deep Ensemble, the variance over different predictions is indicative of the uncertainty.

Predictions $$\boldsymbol{\mu}(y|\mathrm{x}) = \frac{1}{K}\sum_{k=1}^{K} f_\theta([\mathrm{x} - \mathrm{c}_k, \mathrm{c}_k])$$

Uncertainty $$\boldsymbol{\sigma}(y|\mathrm{x}) = \sqrt{\frac{1}{K-1}\sum_{k=1}^{K}(f_\theta([\mathrm{x} - \mathrm{c}_k, \mathrm{c}_k]) - \boldsymbol{\mu})^2}$$

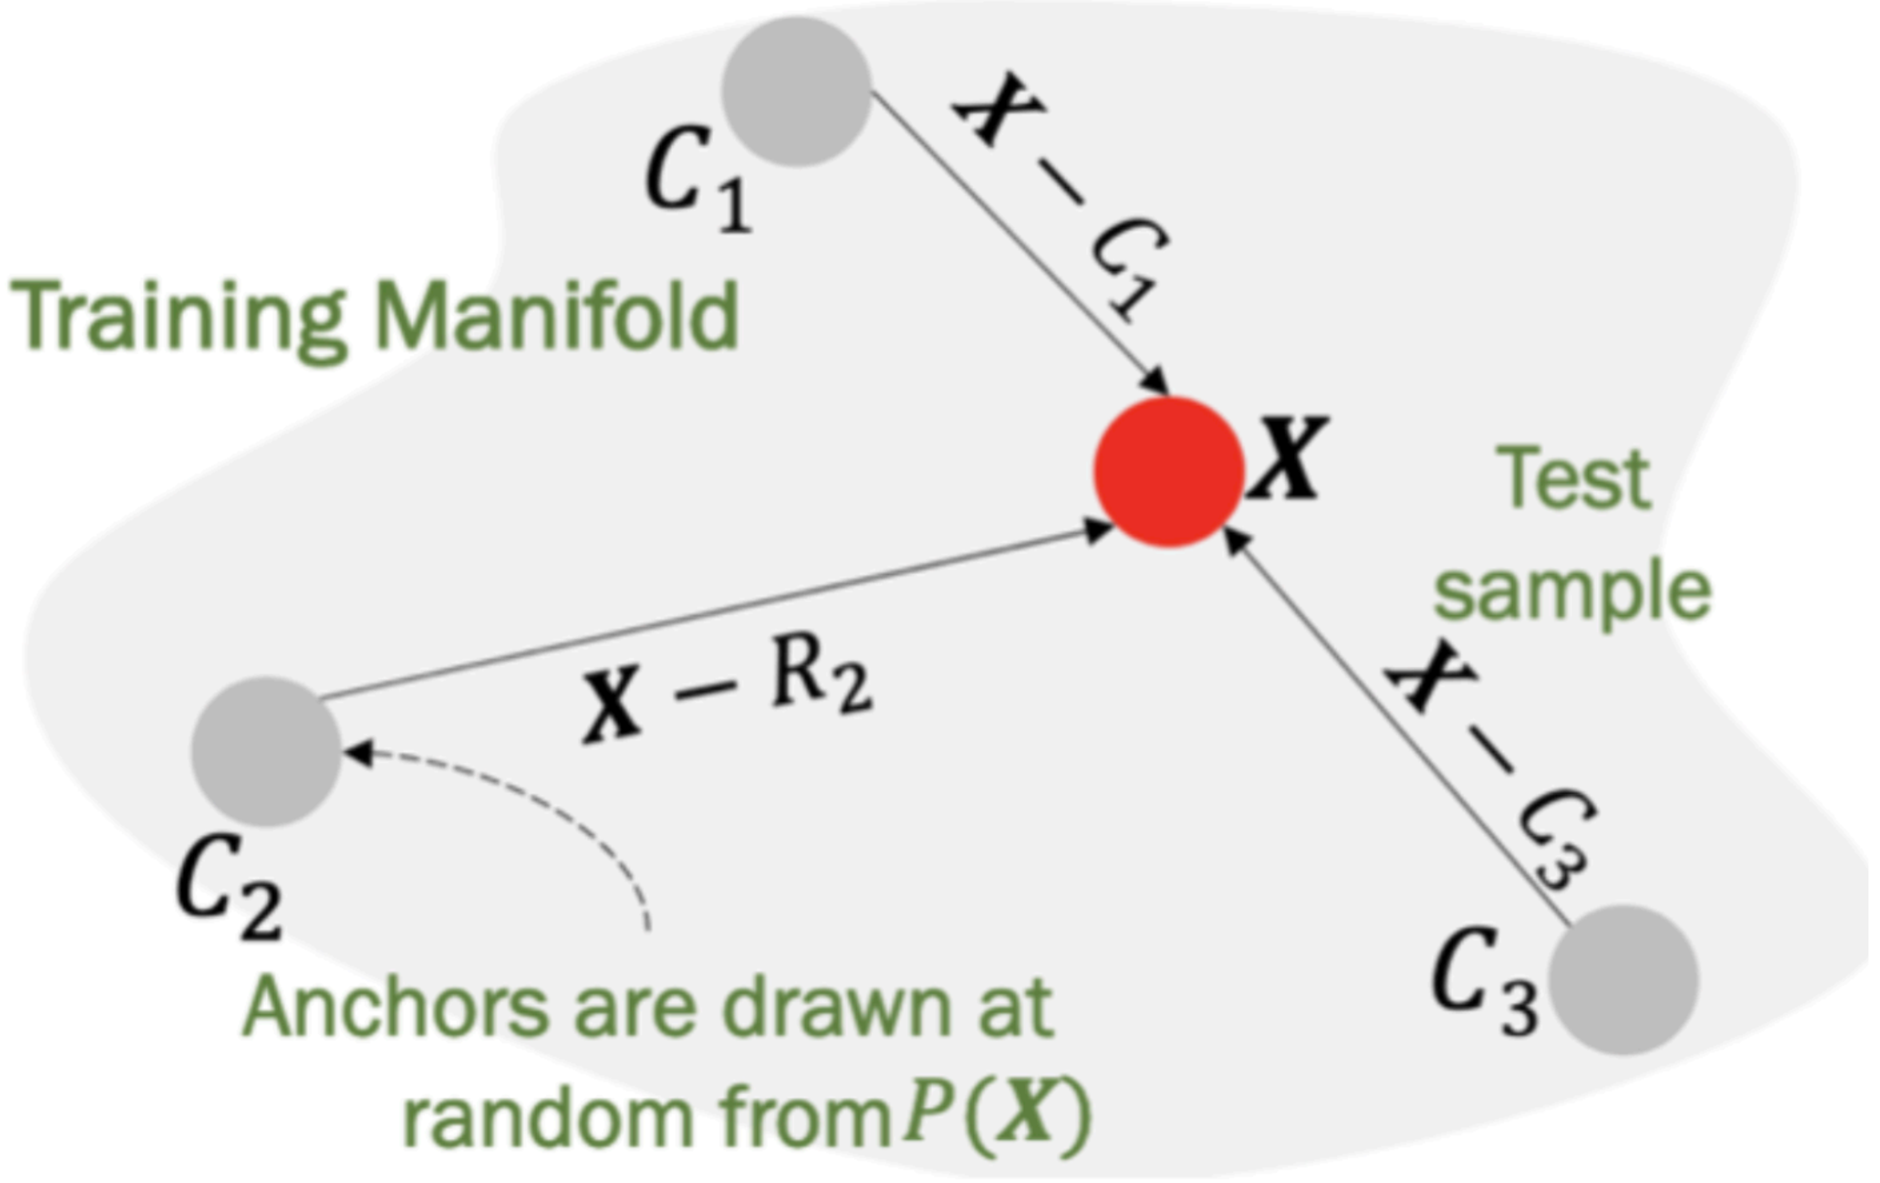

What is Stochastic Centering(ΔUQ)?

- Like a Deep Ensemble, the variance over different predictions is indicative of the uncertainty.

Predictions $$\boldsymbol{\mu}(y|\mathrm{x}) = \frac{1}{K}\sum_{k=1}^{K} f_\theta([\mathrm{x}-\mathrm{c}_k, \mathrm{c}_k])$$

Uncertainty $$\boldsymbol{\sigma}(y|\mathrm{x}) = \sqrt{\frac{1}{K-1}\sum_{k=1}^{K}(f_\theta([\mathrm{x}-\mathrm{c}_k, \mathrm{c}_k]) - \boldsymbol{\mu})^2}$$

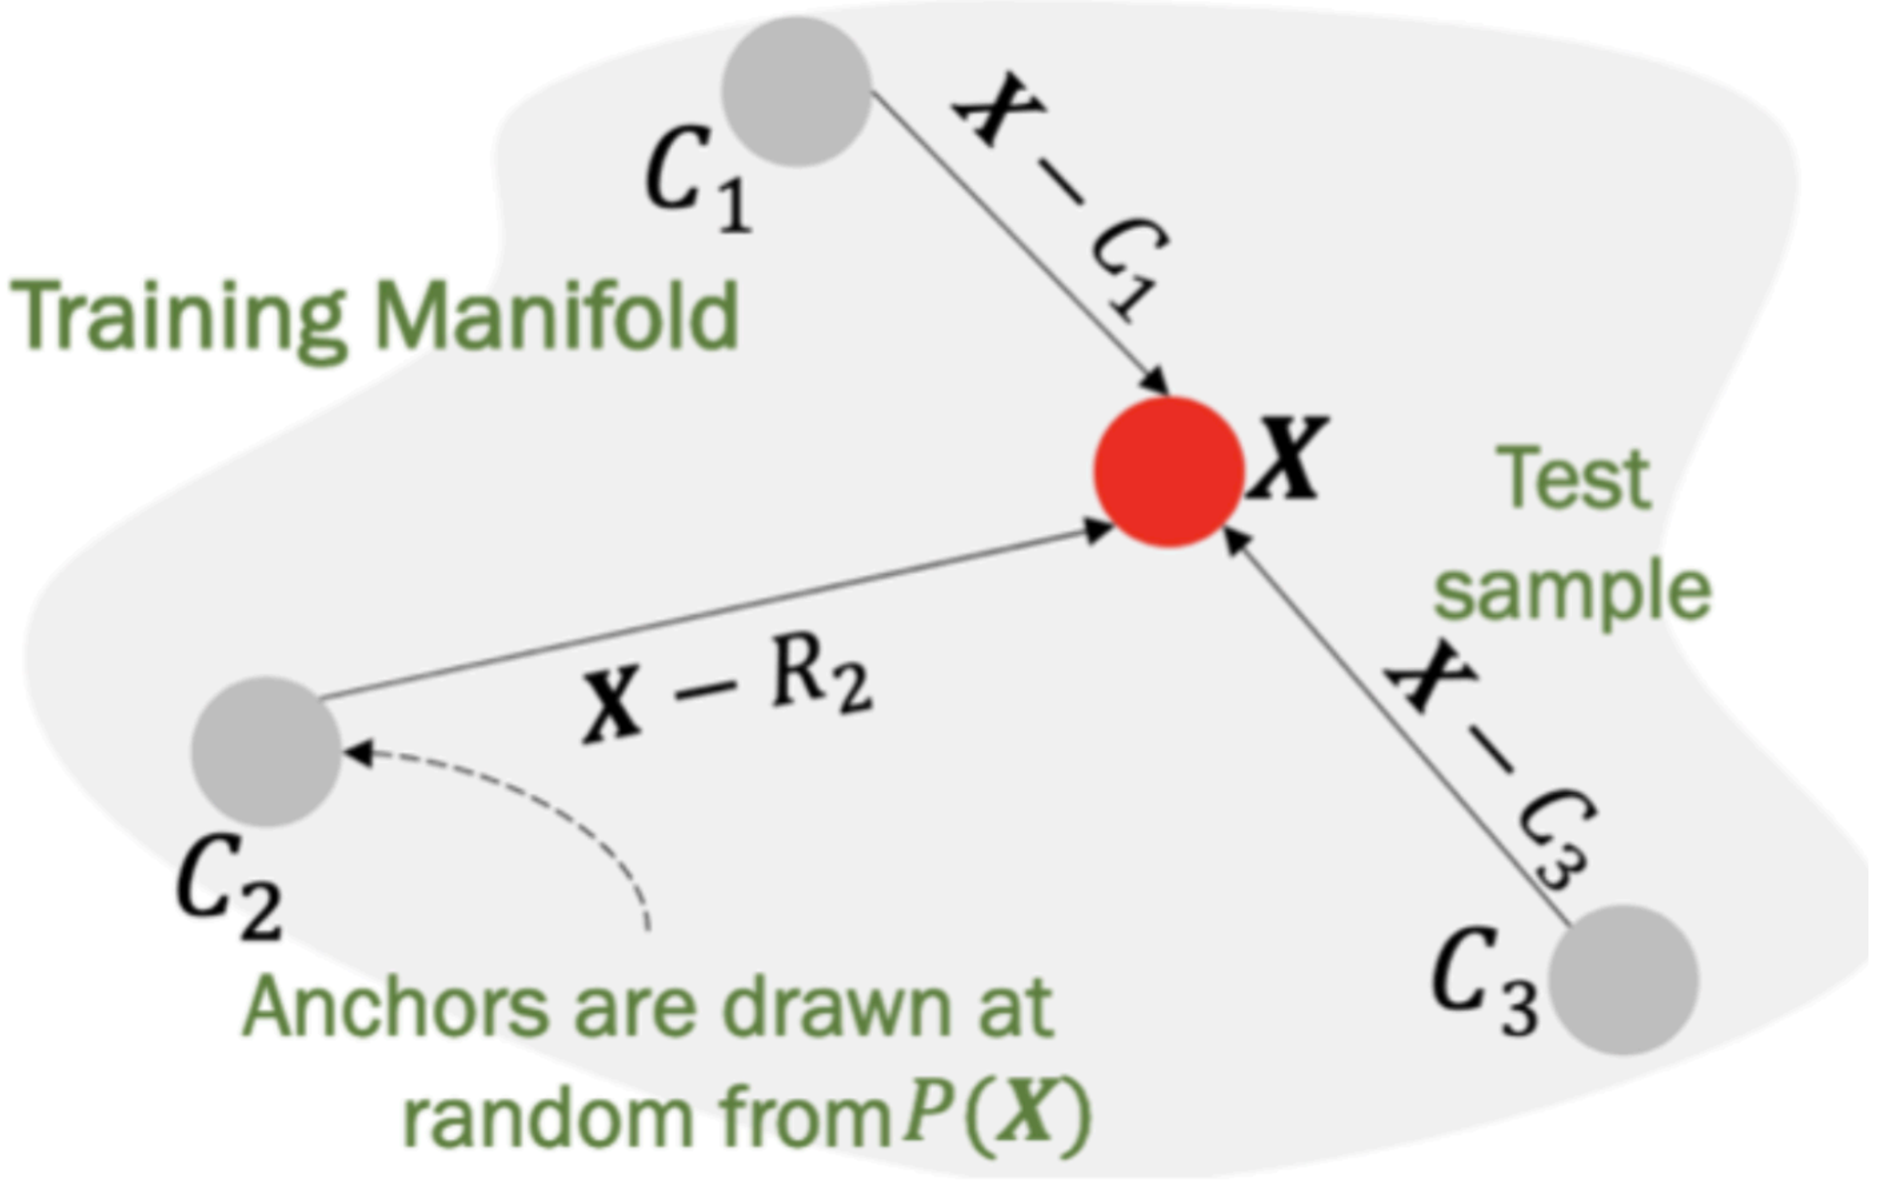

- Stochastic Centering has ***state of the art performance*** for calibration and OOD detection on vision, graph classification and node classification tasks.
- We adapt it to provide node level uncertainty that can improve link prediction performance.

E-ΔUQ (v1)

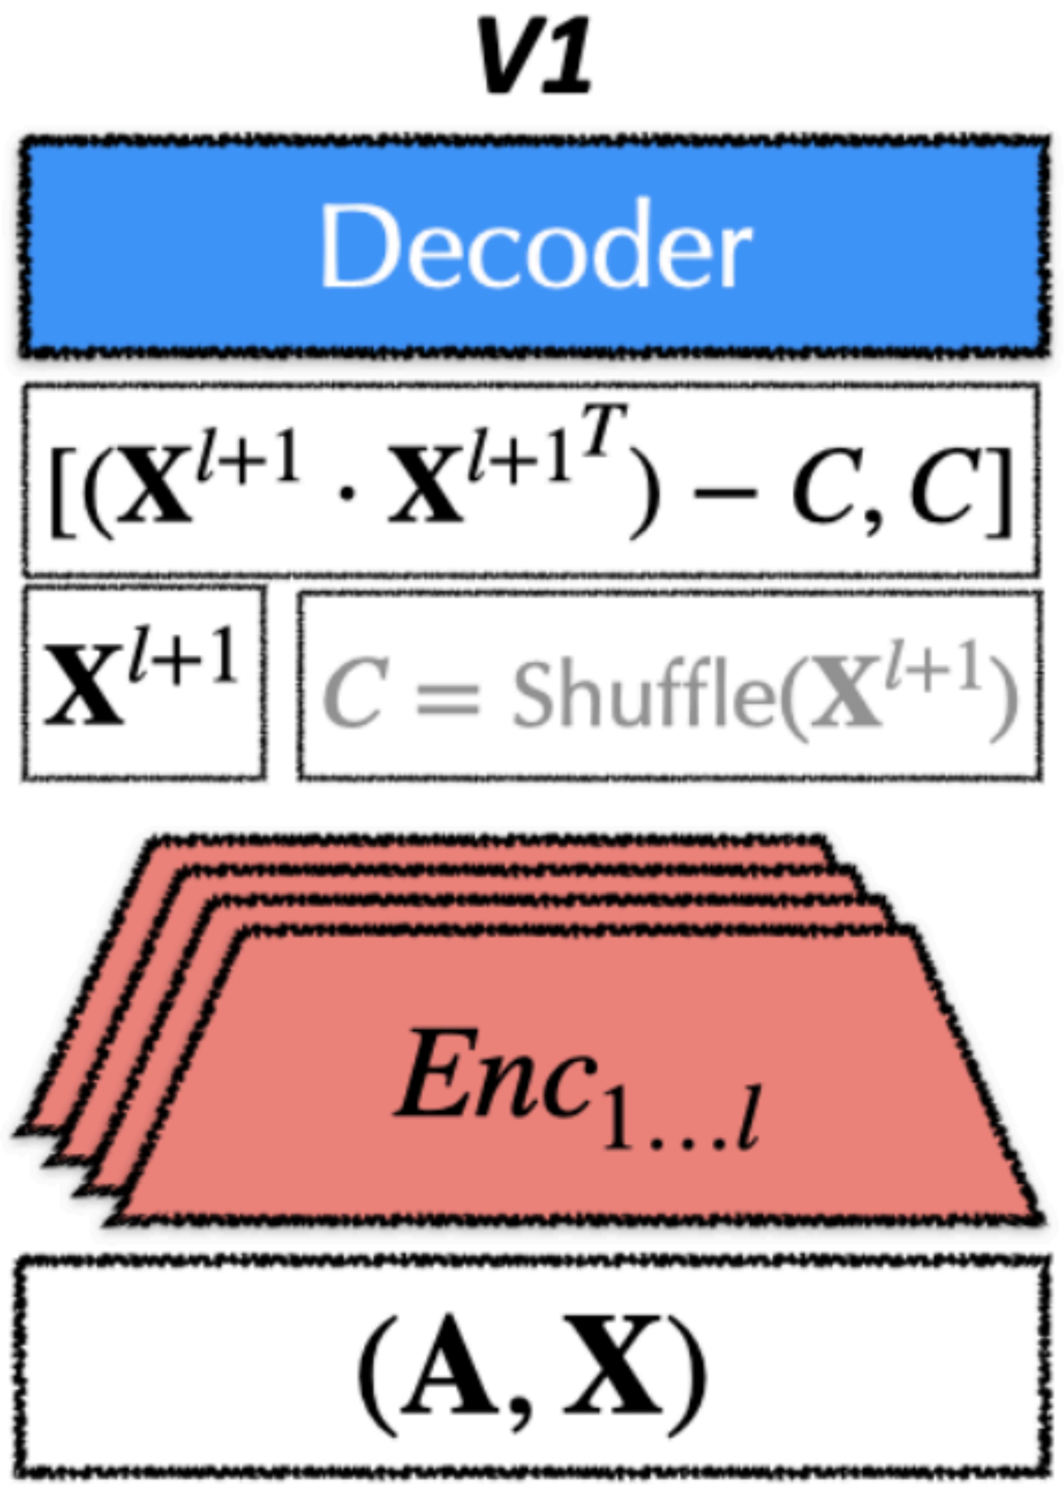

- Perform anchoring after the encoder over the ***node representations.***
- **Anchors** are sampled by shuffling the node representations.

$$\texttt{Decoder}_{dot} : [(\mathbf{x}_i - \mathbf{c}) * (\mathbf{x}_j - \mathbf{c}), \mathbf{c}]$$

$$\texttt{Decoder}_{concat} : [(\mathbf{x}_i - \mathbf{c}||\mathbf{x}_j - \mathbf{c}), \mathbf{c}]$$

- The encoder is ***deterministic***, and the decoder is ***stochastic***.

E-ΔUQ (v1)

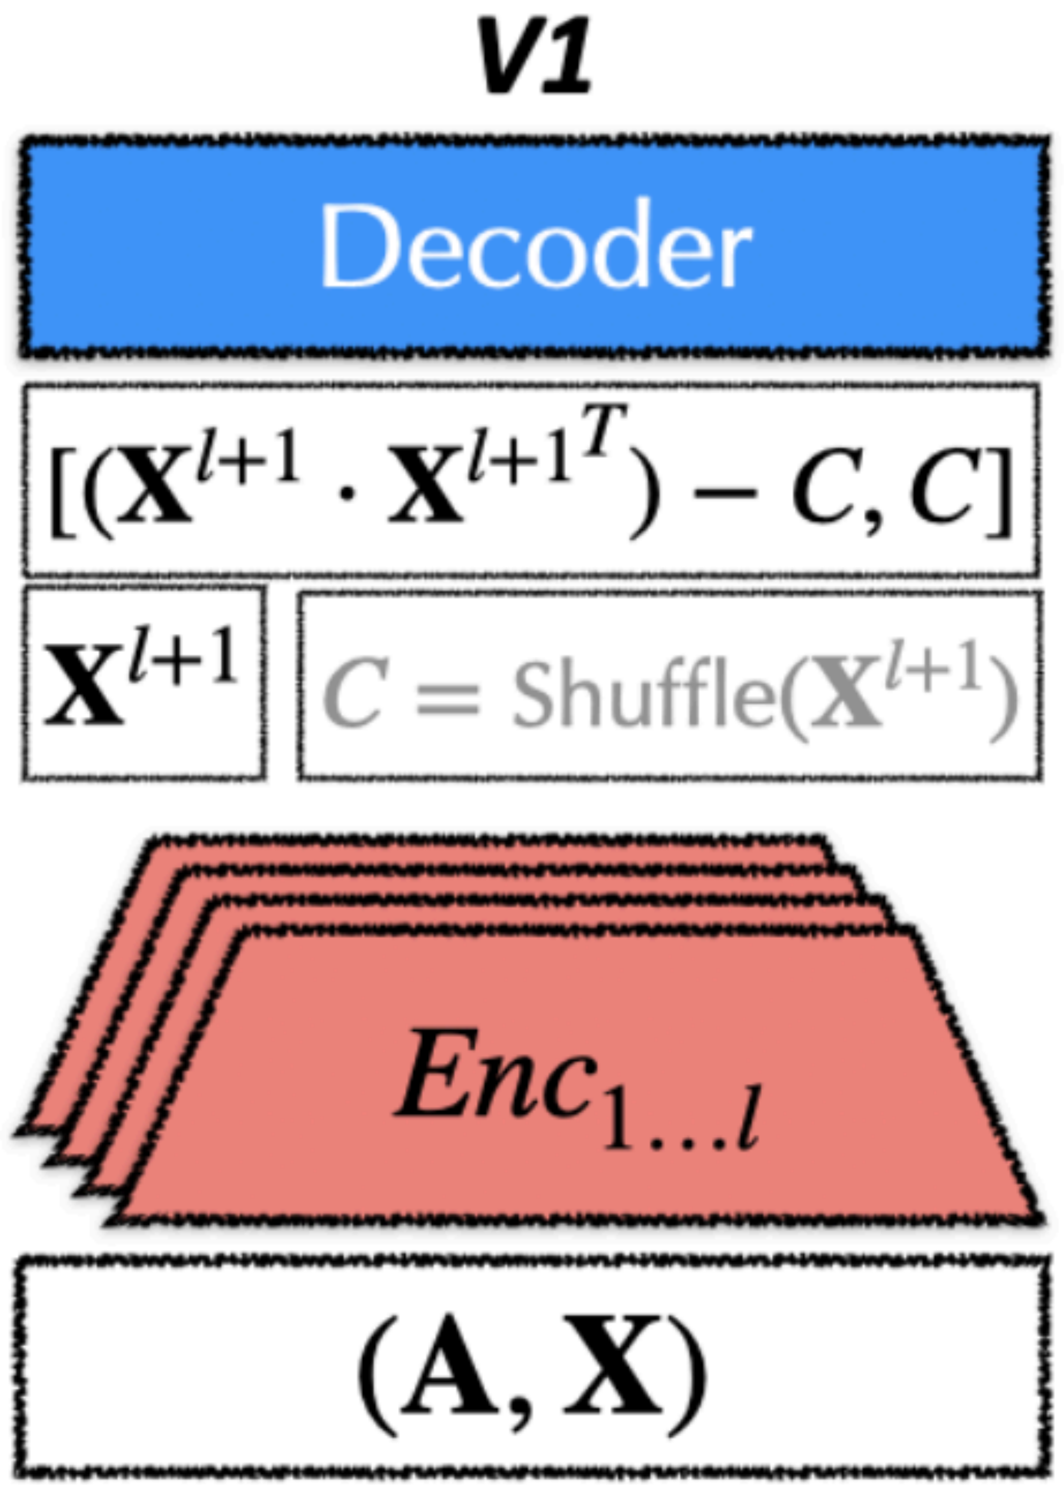

- Perform anchoring after the encoder over the ***node representations.***
- **Anchors** are sampled by shuffling the node representations.

$$\texttt{Decoder}_{dot} : [(\mathbf{x}_i - \mathbf{c}) * (\mathbf{x}_j - \mathbf{c}), \mathbf{c}]$$

$$\texttt{Decoder}_{concat} : [(\mathbf{x}_i - \mathbf{c} || \mathbf{x}_j - \mathbf{c}), \mathbf{c}]$$

- The encoder is ***deterministic***, and the decoder is ***stochastic***.

- While this is a viable extension of ΔUQ, it does not directly use the node uncertainty.

Our Contributions

1. Extending Stochastic Centering to Edge-Level Uncertainty

2. Creating Meaningful Node-level uncertainties.

3. Experimental Evaluation of Edge-ΔUQ

E-ΔUQ (v2): Partially Stochastic Encoder

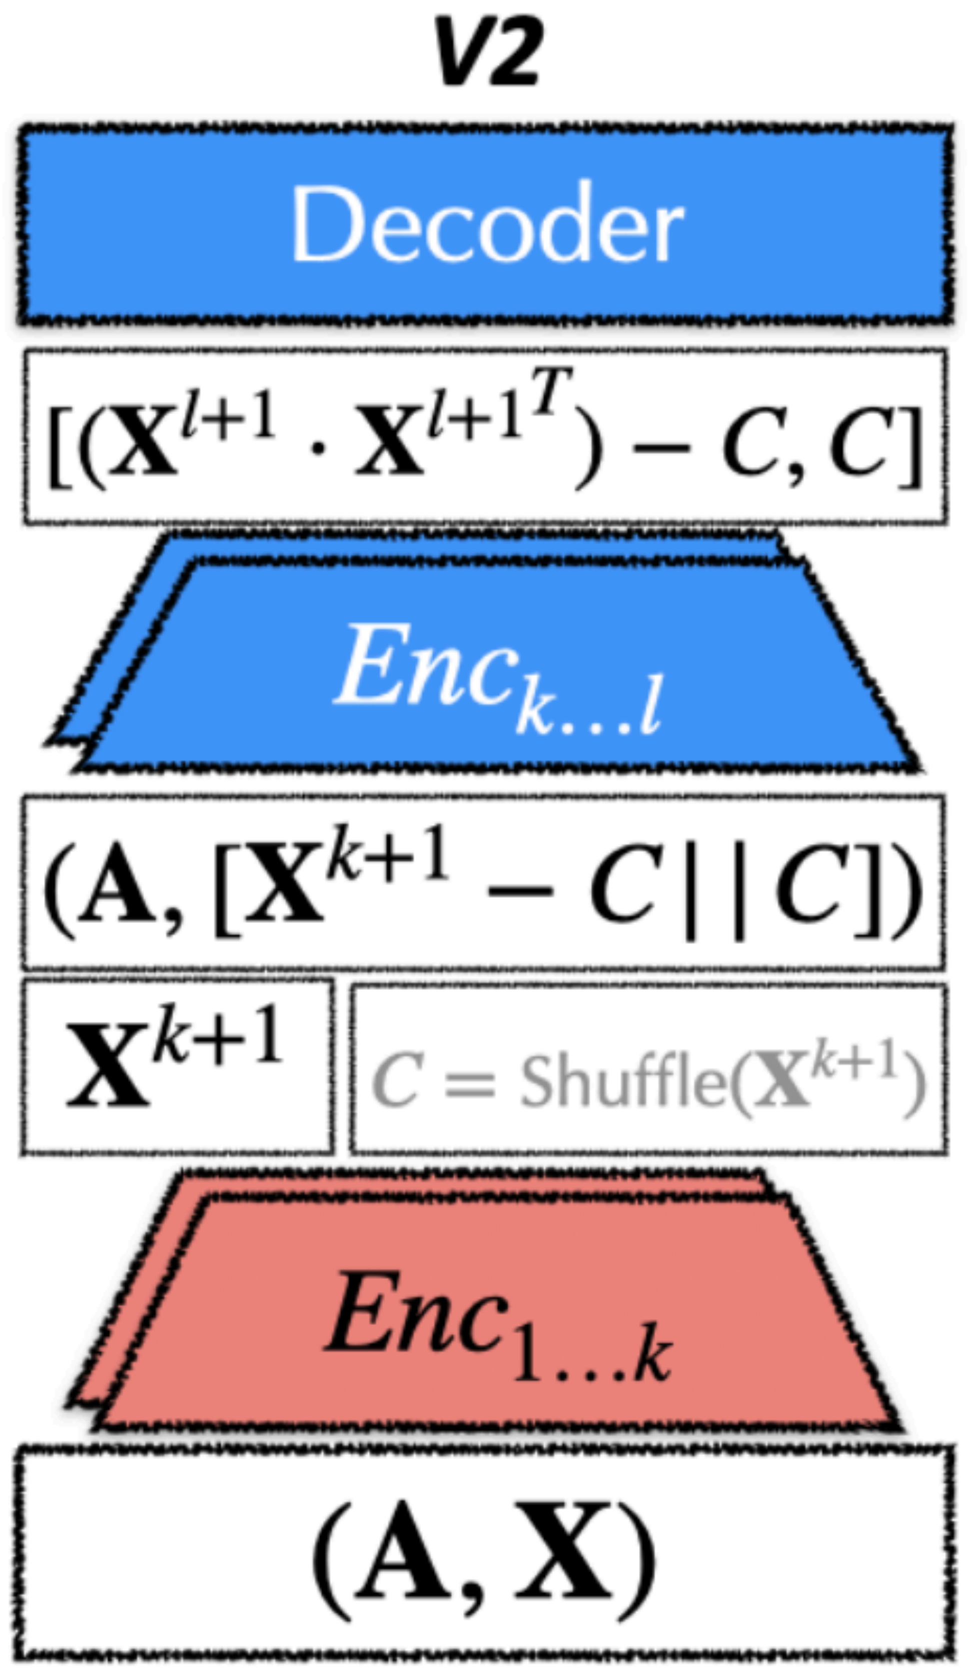

- Sampling a diverse set of hypotheses is important for ensuring useful epistemic uncertainty estimates.

- Using a stochastic encoder supports more diverse hypotheses and helps capture more node level uncertainty.

$$\mathbf{X}^{r+1} = \texttt{Encoder}^{1\dots r}(\mathbf{X}, \mathbf{A})$$
$$\mathbf{X}^{\ell+1} = \texttt{Encoder}^{r+1\dots\ell}\left([\mathbf{X}^{r+1} - \mathbf{C}, \mathbf{C}], \mathbf{A}\right)$$
$$\hat{E}_{(i,j)} = \texttt{Decoder}\left(\mathbf{X}_i^{\ell+1}, \mathbf{X}_j^{\ell+1}\right)$$

E-ΔUQ (v2): Partially Stochastic Encoder

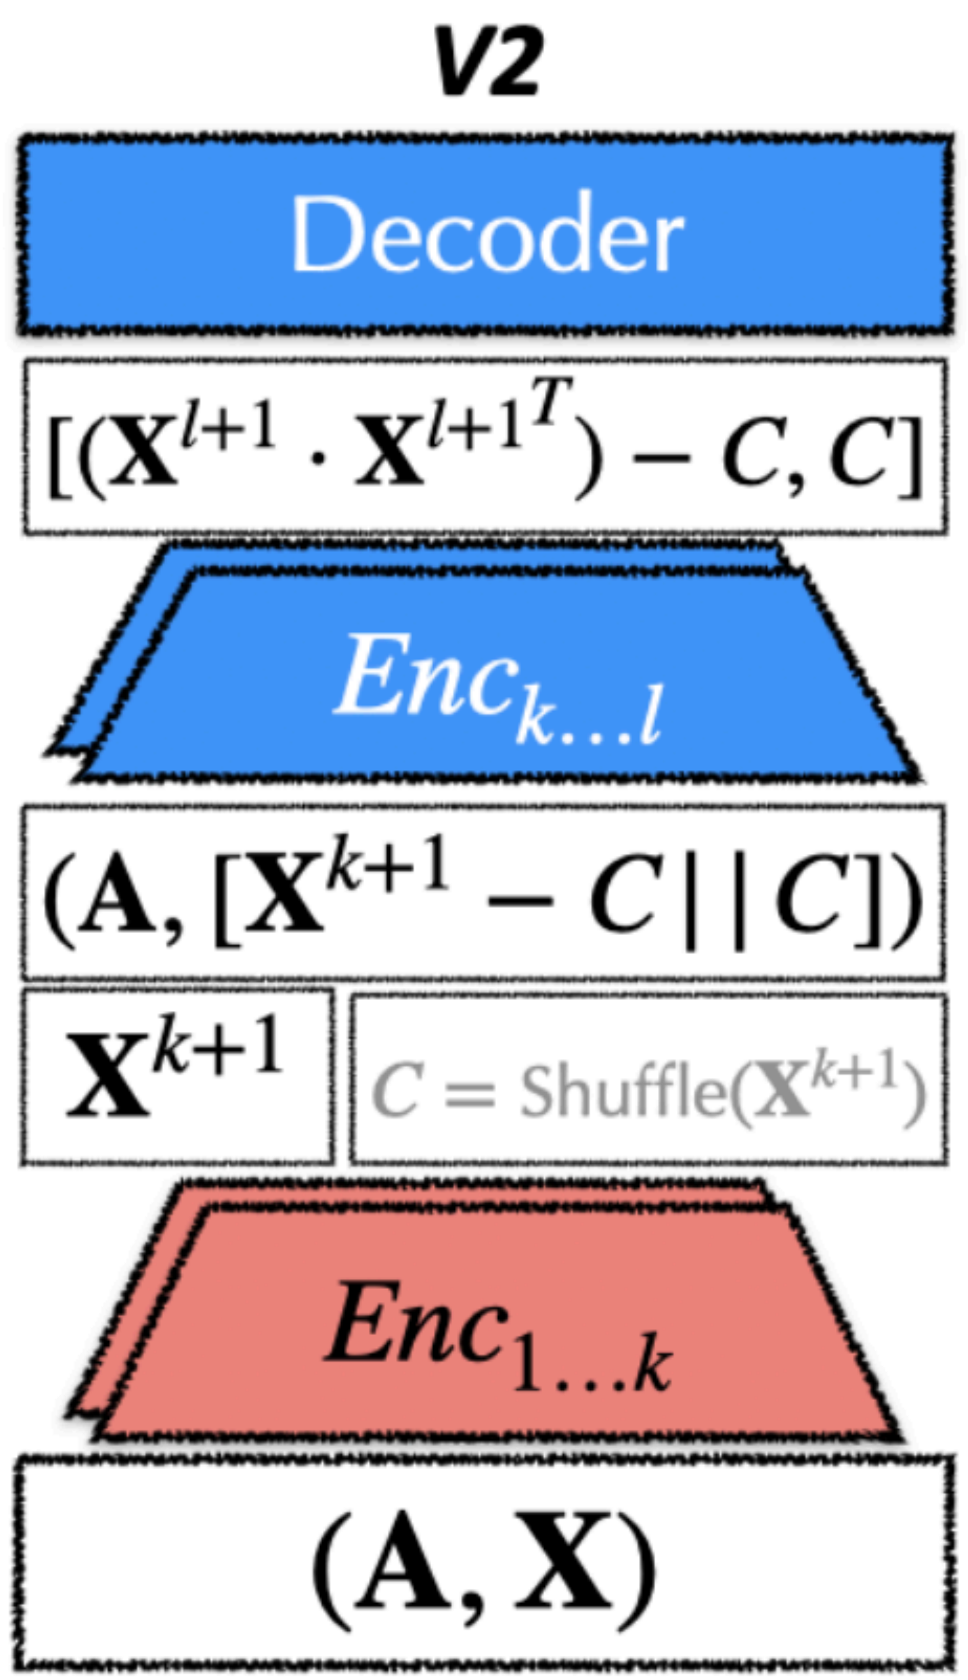

- Sampling a diverse set of hypotheses is important for ensuring useful epistemic uncertainty estimates.
- Using a stochastic encoder supports more diverse hypotheses and helps capture more node level uncertainty.

$$\mathbf{X}^{r+1} = \texttt{Encoder}^{1\ldots r}(\mathbf{X}, \mathbf{A})$$
$$\mathbf{X}^{\ell+1} = \texttt{Encoder}^{r+1\ldots\ell}\left([\mathbf{X}^{r+1} - \mathbf{C}, \mathbf{C}], \mathbf{A}\right)$$
$$\hat{E}_{(i,j)} = \texttt{Decoder}\left(\mathbf{X}_i^{\ell+1}, \mathbf{X}_j^{\ell+1}\right)$$

While we know have more diverse node level uncertainty, there is no apriori guarantee that these are calibrated!

E-ΔUQ (v3):
Partially Stochastic Encoder + Node Level Pretraining

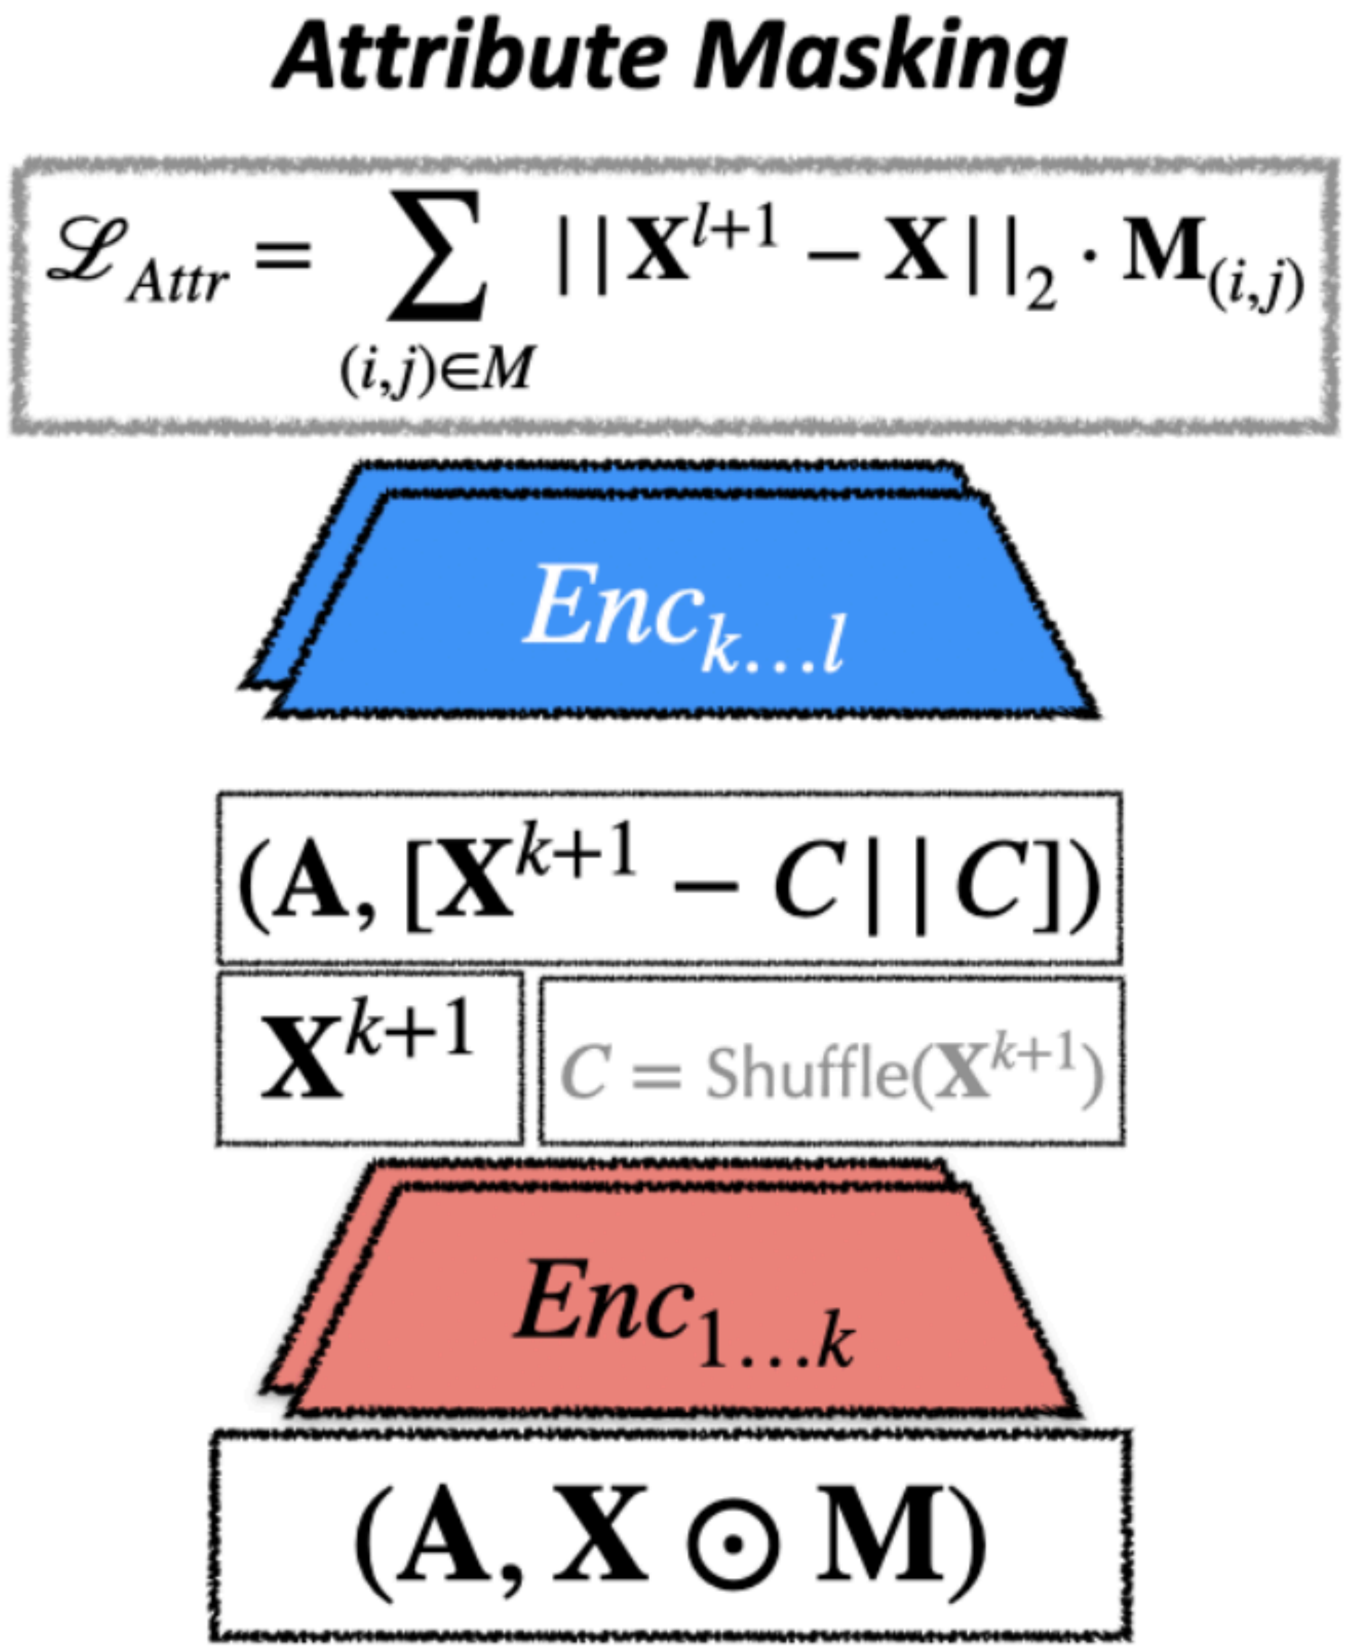

- To improve node level uncertainty, we use an ***auxiliary feature reconstruction*** task.

$$\mathbf{X}^{r+1} = \texttt{Encoder}^{1\ldots r}(\mathbf{X} \odot \mathbf{M}, , \mathbf{A})$$
$$\mathbf{X}^{\ell+1} = \texttt{Encoder}^{r+1\ldots\ell}\left([\mathbf{X}^{r+1} - \mathbf{C}, \mathbf{C}], \mathbf{A}\right)$$
$$\mathcal{L}_{Attr} = \sum_{(i,j)\in\mathbf{M}} ||\mathbf{X}^{\ell+1} - \mathbf{X}||_2 \cdot \mathbf{M}_{(i,j)}$$

- The model is ***trained end to end*** with the additional task at a negligible loss.

E-ΔUQ: Stochastic Centering for Link Prediction

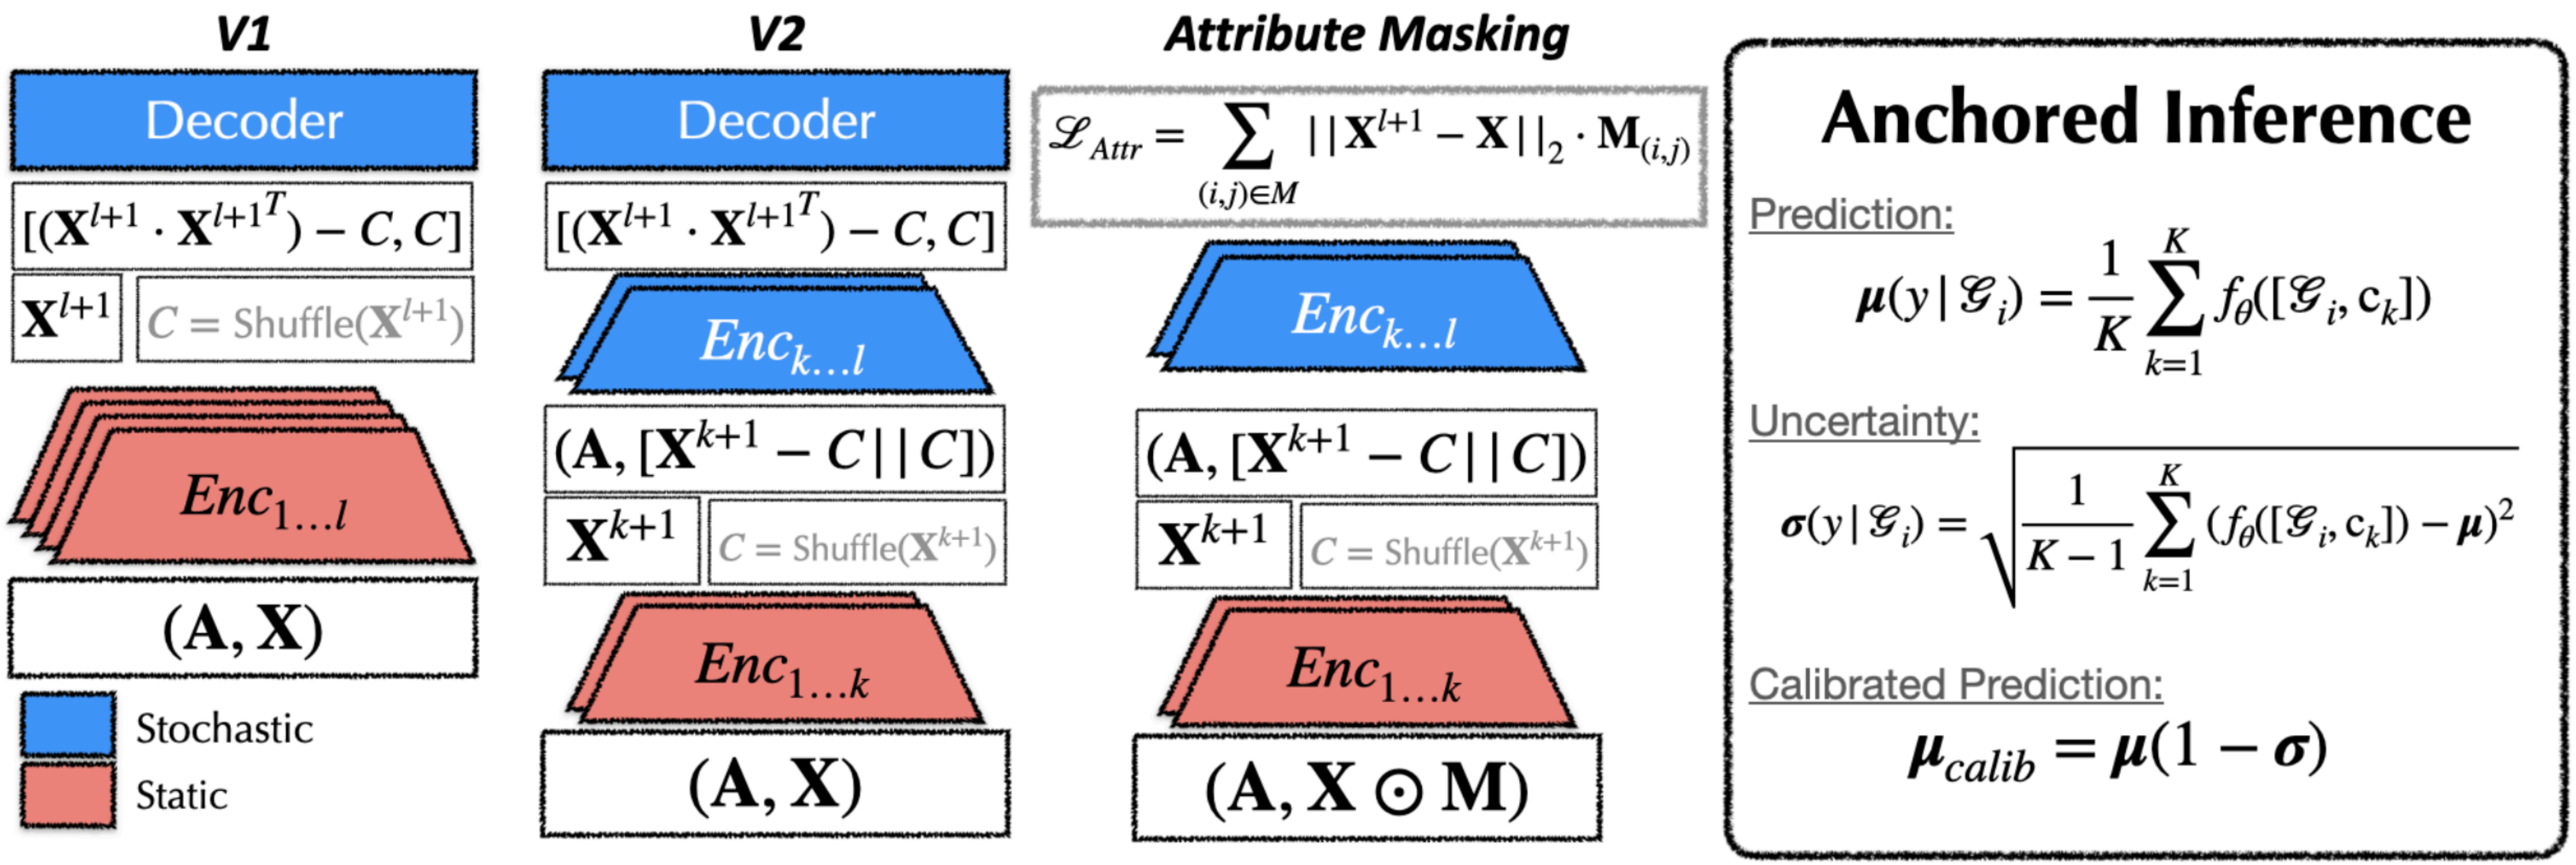

Our Contributions

1. Extending Stochastic Centering to Edge-Level Uncertainty

2. Creating Meaningful Node-level uncertainties.

3. Experimental Evaluation of Edge-ΔUQ

Experimental Evaluation

Dataset	Method	AUPR (↑)	ECE (↓)
Citeseer	E-ΔUQ (v3)		
	E-ΔUQ (v2)		
	E-ΔUQ (v1)		
	Vanilla		
Cora	E-ΔUQ (v3)		
	E-ΔUQ (v2)		
	E-ΔUQ (v1)		
	Vanilla		
Pubmed	E-ΔUQ (v3)		
	E-ΔUQ (v2)		
	E-ΔUQ (v1)		
	Vanilla		

Experimental Evaluation

Dataset	Method	AUPR (↑)	ECE (↓)
Citeseer	E-ΔUQ (v3)	0.8409 ±0.0115	**0.2591** ±0.0178
	E-ΔUQ (v2)	**0.8548** ±0.0076	0.2833 ±0.0075
	E-ΔUQ (v1)	0.8070 ±0.0218	0.3056 ±0.0109
	Vanilla	0.8236 ±0.0115	0.3002 ±0.0062
Cora	E-ΔUQ (v3)	0.8886 ±0.0042	**0.1554** ±0.0060
	E-ΔUQ (v2)	0.8888 ±0.0062	0.1731 ±0.0181
	E-ΔUQ (v1)	0.8598 ±0.0207	0.2640 ±0.0125
	Vanilla	**0.8936** ±0.0066	0.3503 ±0.0146
Pubmed	E-ΔUQ (v3)	0.8775 ±0.0098	0.1818 ±0.0048
	E-ΔUQ (v2)	0.8701 ±0.0016	**0.1538** ±0.0059
	E-ΔUQ (v1)	**0.9069** ±0.0063	0.1801 ±0.0117
	Vanilla	0.8897 ±0.0091	0.1980 ±0.0035

- **Obs 1:** E-ΔUQ improves the calibration on all datasets over the vanilla model.

Experimental Evaluation

Dataset	Method	AUPR (↑)	ECE (↓)
Citeseer	E-ΔUQ (v3)	0.8409 ±0.0115	**0.2591** ±0.0178
	E-ΔUQ (v2)	**0.8548** ±0.0076	0.2833 ±0.0075
	E-ΔUQ (v1)	0.8070 ±0.0218	0.3056 ±0.0109
	Vanilla	0.8236 ±0.0115	0.3002 ±0.0062
Cora	E-ΔUQ (v3)	0.8886 ±0.0042	**0.1554** ±0.0060
	E-ΔUQ (v2)	0.8888 ±0.0062	0.1731 ±0.0181
	E-ΔUQ (v1)	0.8598 ±0.0207	0.2640 ±0.0125
	Vanilla	**0.8936** ±0.0066	0.3503 ±0.0146
Pubmed	E-ΔUQ (v3)	0.8775 ±0.0098	0.1818 ±0.0048
	E-ΔUQ (v2)	0.8701 ±0.0016	**0.1538** ±0.0059
	E-ΔUQ (v1)	**0.9069** ±0.0063	0.1801 ±0.0117
	Vanilla	0.8897 ±0.0091	0.1980 ±0.0035

- **Obs 1:** E-ΔUQ improves the calibration on all datasets over the vanilla model.
- **Obs 2:** E-ΔUQ perform comparably on ***AUPR (best 2/3).***

Experimental Evaluation

Dataset	Method	AUPR (↑)	ECE (↓)
Citeseer	E-ΔUQ (v3)	0.8409 ±0.0115	**0.2591** ±0.0178
	E-ΔUQ (v2)	**0.8548** ±0.0076	0.2833 ±0.0075
	E-ΔUQ (v1)	0.8070 ±0.0218	0.3056 ±0.0109
	Vanilla	0.8236 ±0.0115	0.3002 ±0.0062
Cora	E-ΔUQ (v3)	0.8886 ±0.0042	**0.1554** ±0.0060
	E-ΔUQ (v2)	0.8888 ±0.0062	0.1731 ±0.0181
	E-ΔUQ (v1)	0.8598 ±0.0207	0.2640 ±0.0125
	Vanilla	**0.8936** ±0.0066	0.3503 ±0.0146
Pubmed	E-ΔUQ (v3)	0.8775 ±0.0098	0.1818 ±0.0048
	E-ΔUQ (v2)	0.8701 ±0.0016	**0.1538** ±0.0059
	E-ΔUQ (v1)	**0.9069** ±0.0063	0.1801 ±0.0117
	Vanilla	0.8897 ±0.0091	0.1980 ±0.0035

- **Obs 1:** E-ΔUQ improves the calibration on all datasets over the vanilla model.
- **Obs 2:** E-ΔUQ perform comparably on ***AUPR (best 2/3).***
- **Obs 3:** ***E-ΔUQ (v3)*** obtains the best calibration on 2/3 datasets.

Contributions

- Extending Stochastic Centering to Edge-Level Uncertainty
- Creating Meaningful Node-level uncertainties
- Experimental Evaluation of Edge-ΔUQ

Thank you!

Questions?: pujat@umich.edu
pujacomputes.github.io