Does Invariant Graph Learning via Environment Augmentation Learn Invariance?

Yongqiang Chen
CUHK, Tencent AI Lab

with Yatao Bian, Kaiwen Zhou, Binghui Xie, Bo Han, and James Cheng

Out-of-Distribution Generalization on Graphs

OOD generalization on graphs is fundamentally more **challenging** than that on Euclidean data:

$f_{\text{GNN}}(\{$ 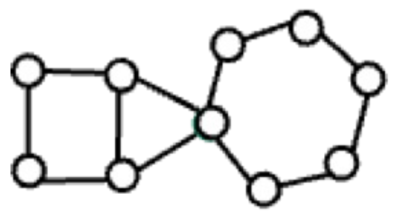 $\}$, $\{$  $\})$ = **“House”**

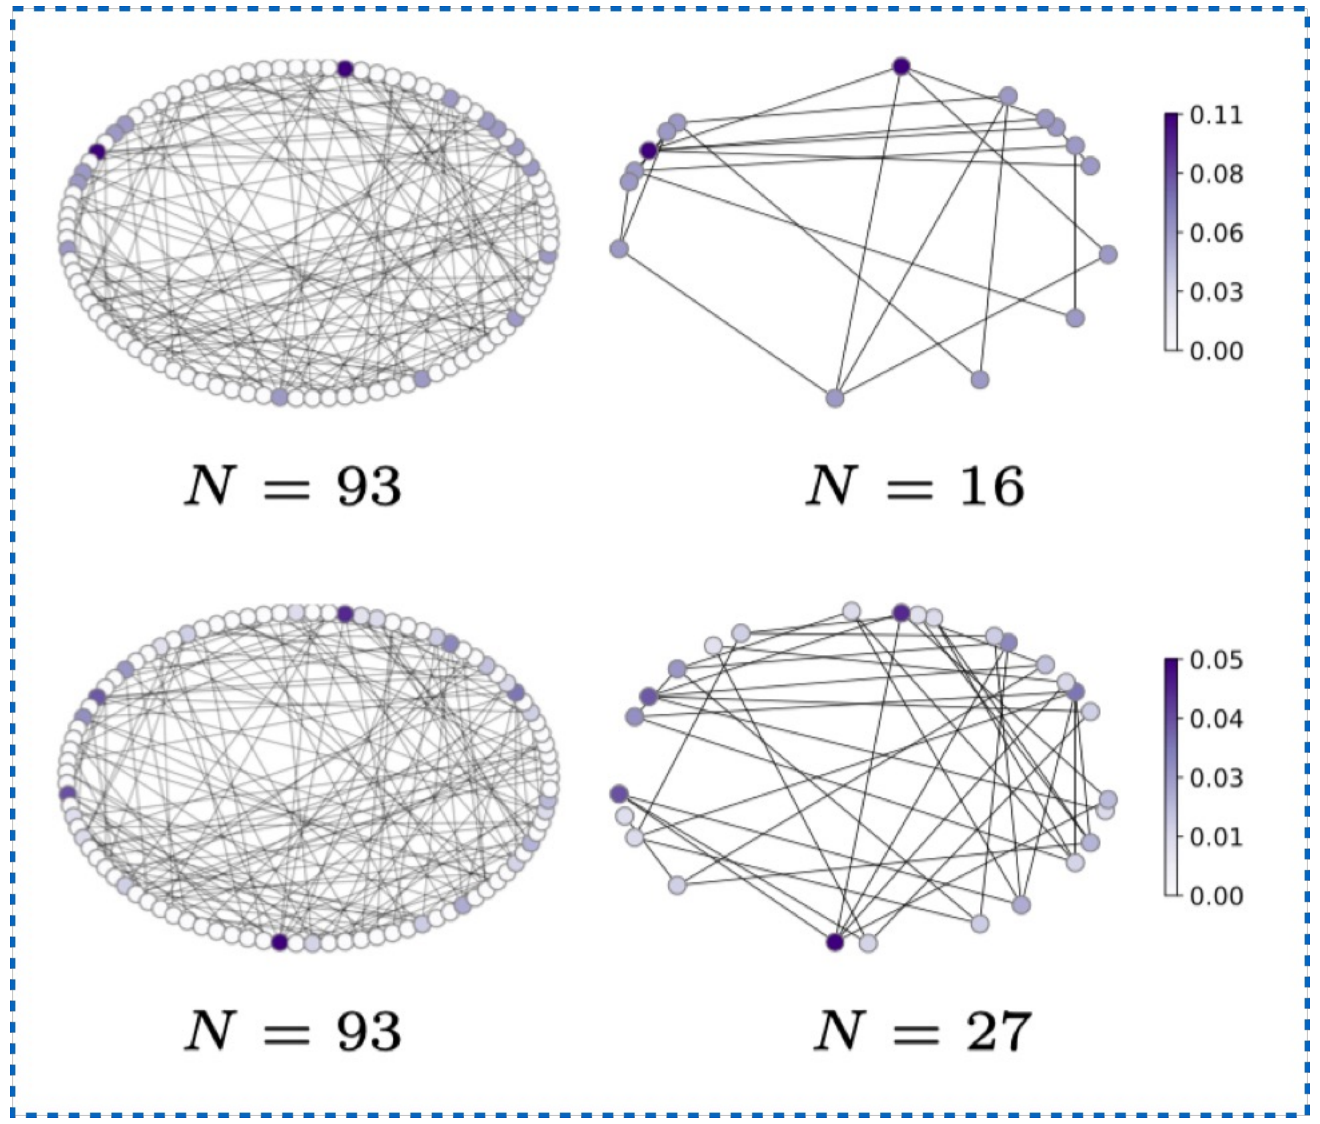

Structure-level shifts

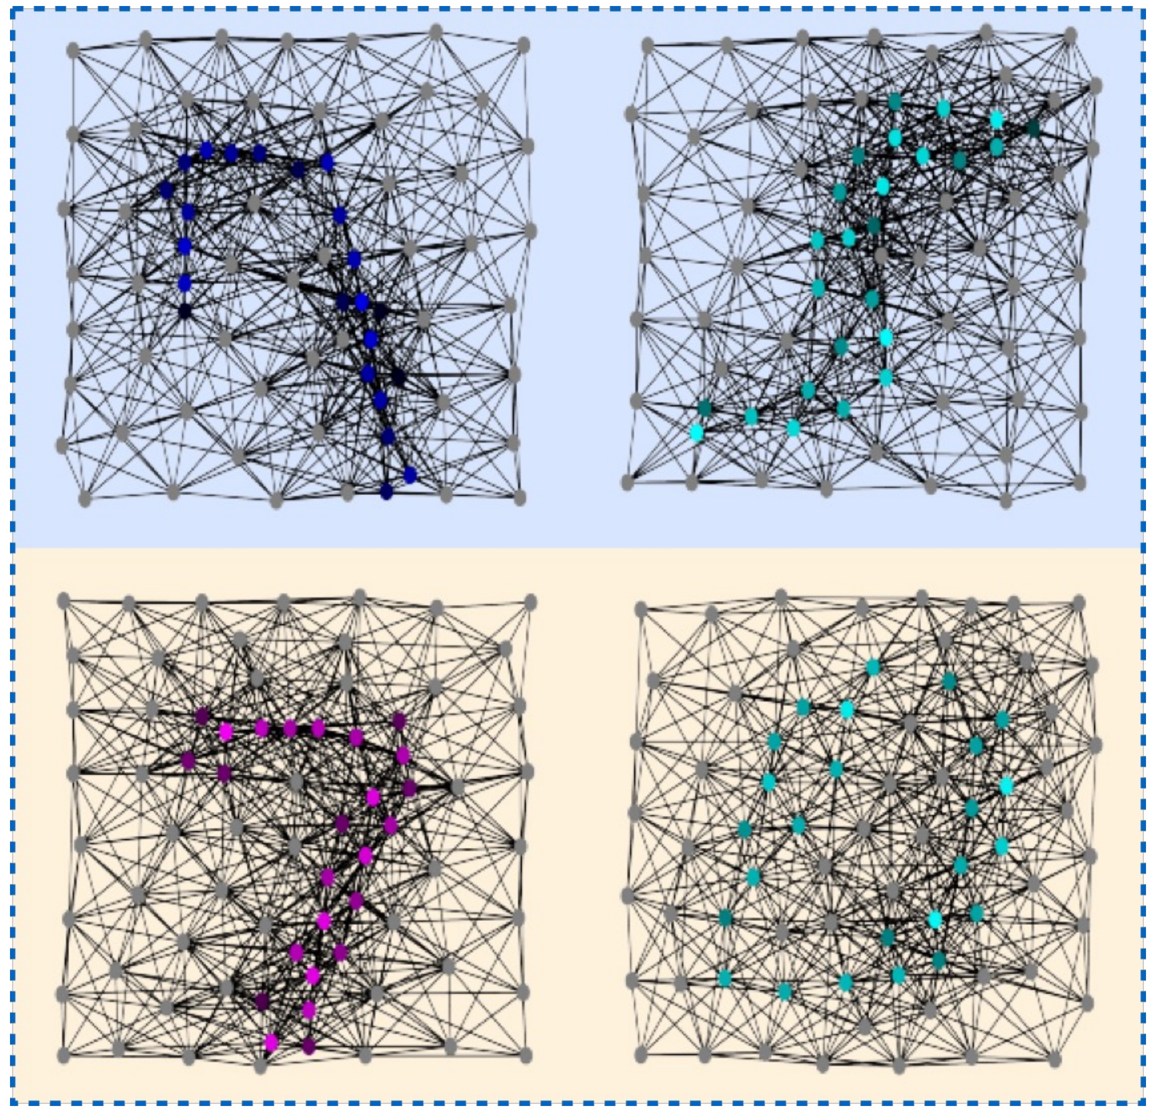

Attribute-level shifts

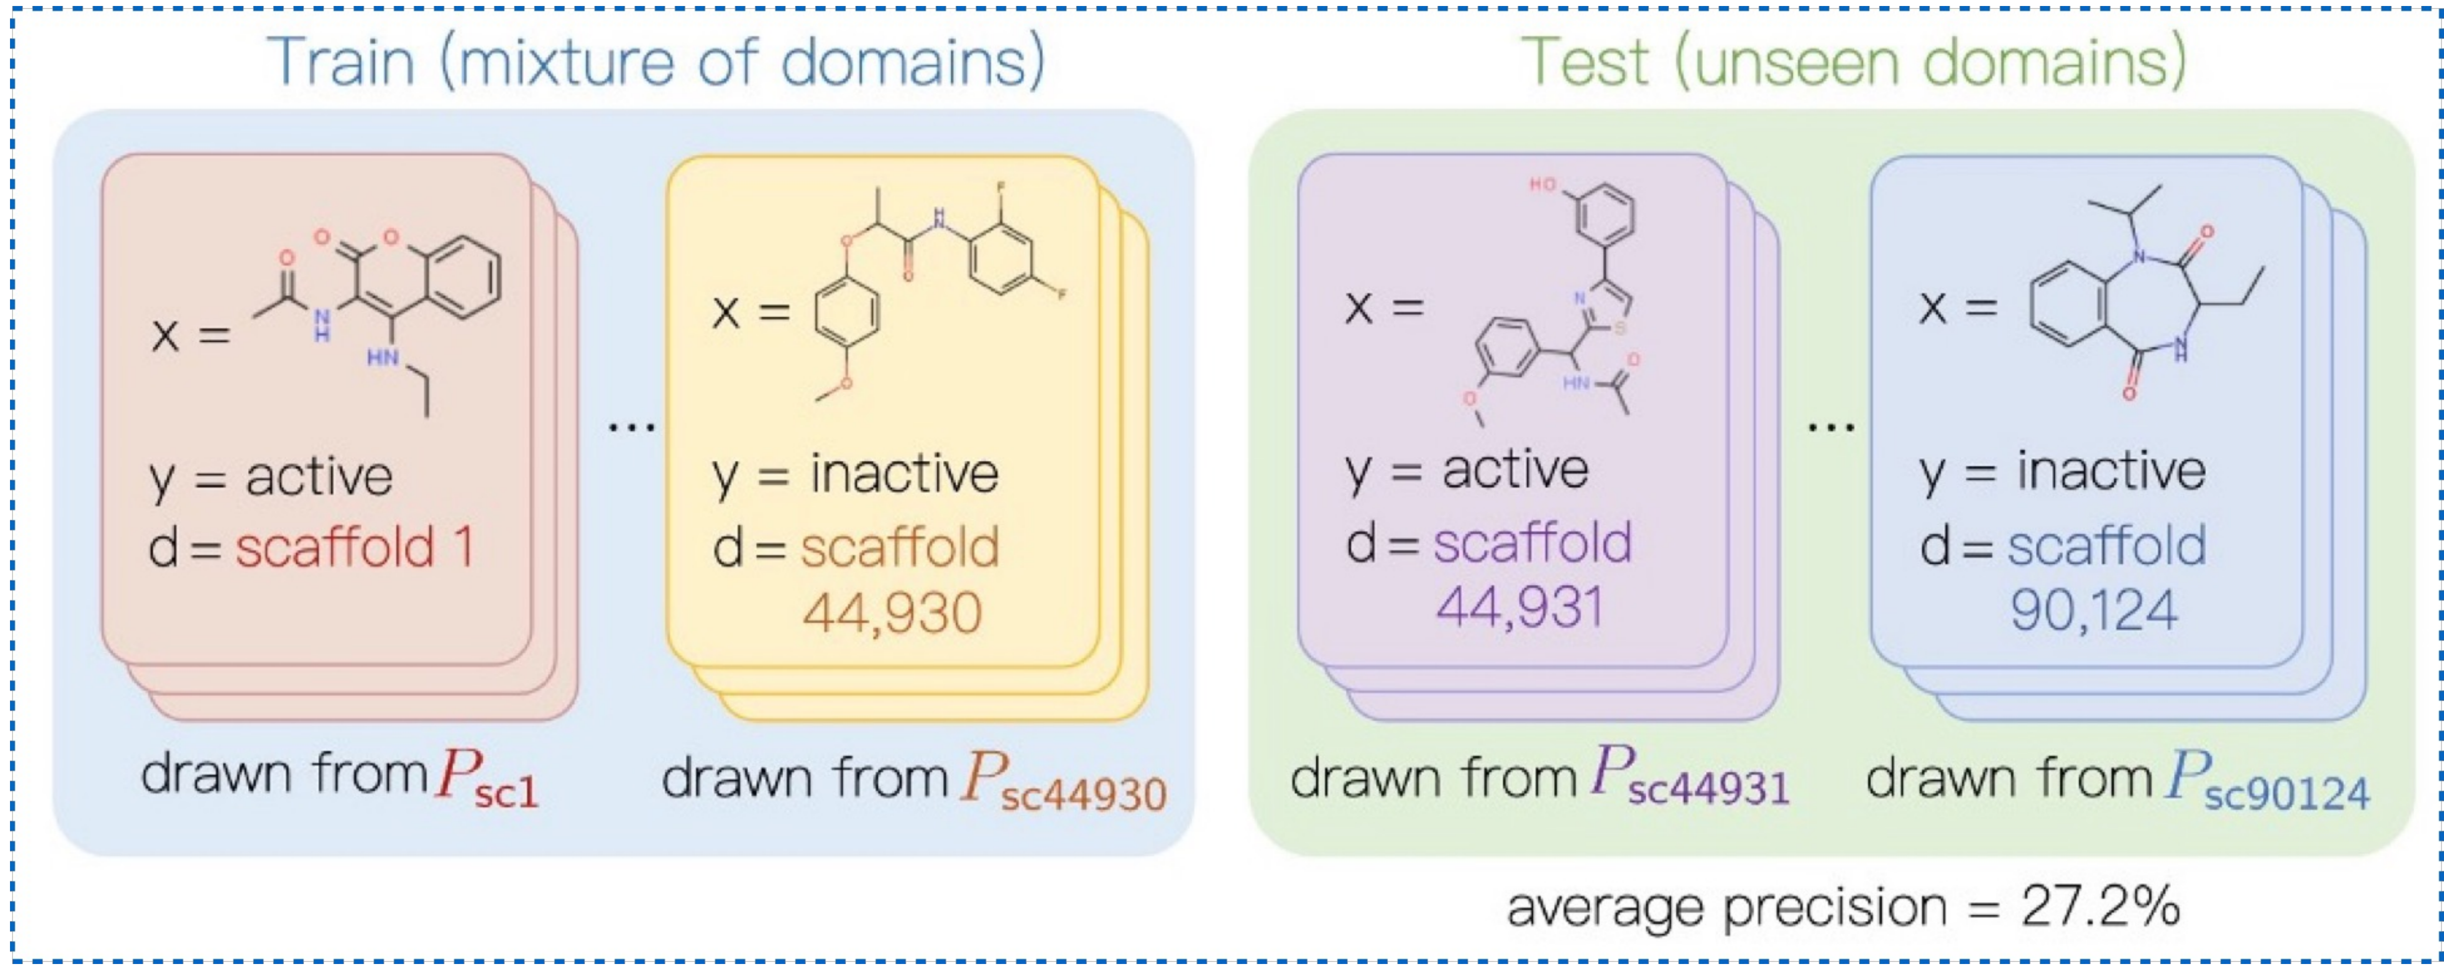

Mixture of structure-level and attribute-level shifts

(Knyazev et al. 2019; Hu et al., 2020; Koh et al., 2021; Gui et al., 2022; Chen et al., 2022)

Out-of-Distribution Generalization on Graphs

Invariant graph representation learning aims to identify an **invariant subgraph** among graphs from different **environments** or **domains**:

Environment #1: Class “House”

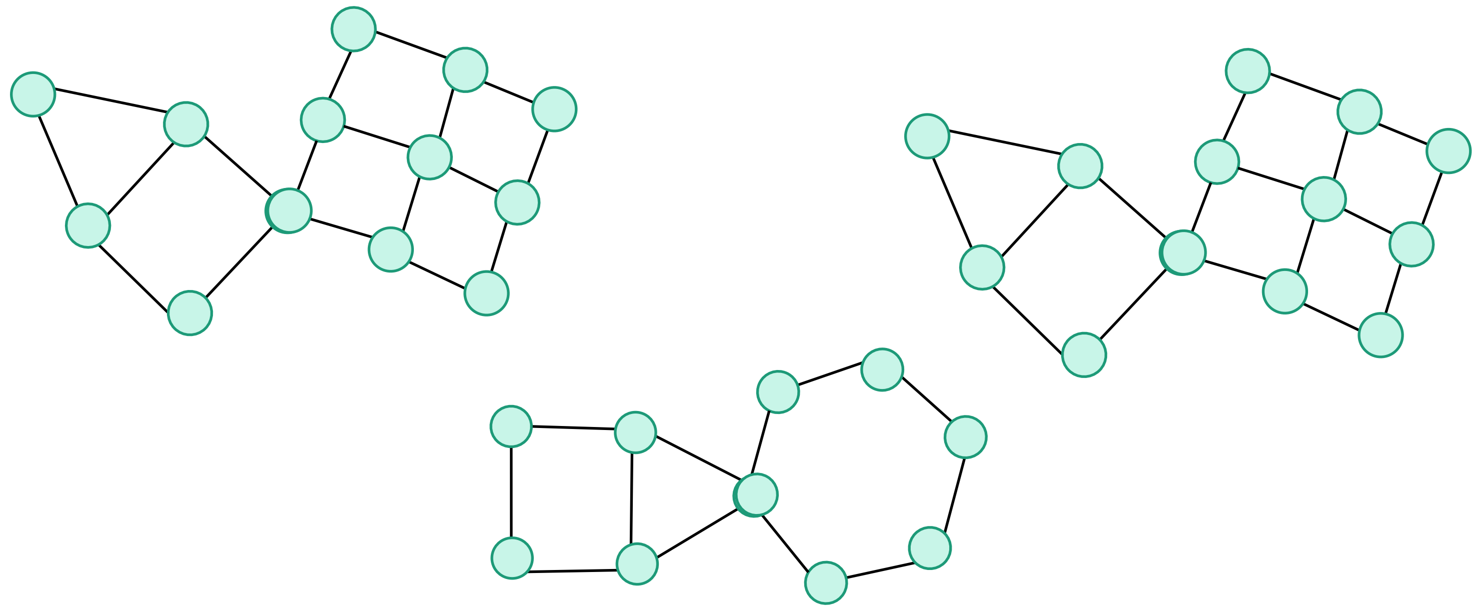

Environment #2: Class “House”

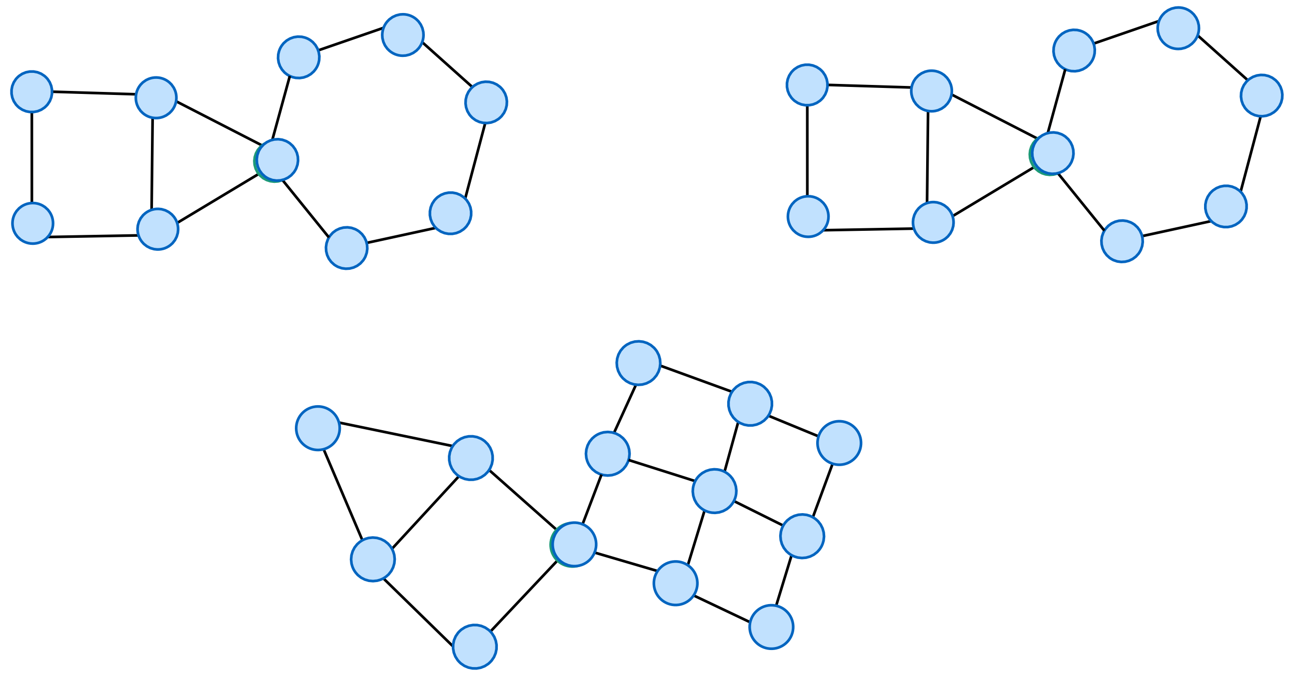

(Wu et al., 2022ab; Miao et al., 2022; Chen et al., 2022)

Out-of-Distribution Generalization on Graphs

Invariant graph representation learning aims to identify an **invariant subgraph** among graphs from different **environments** or **domains**:

Environment #1: Class “House”

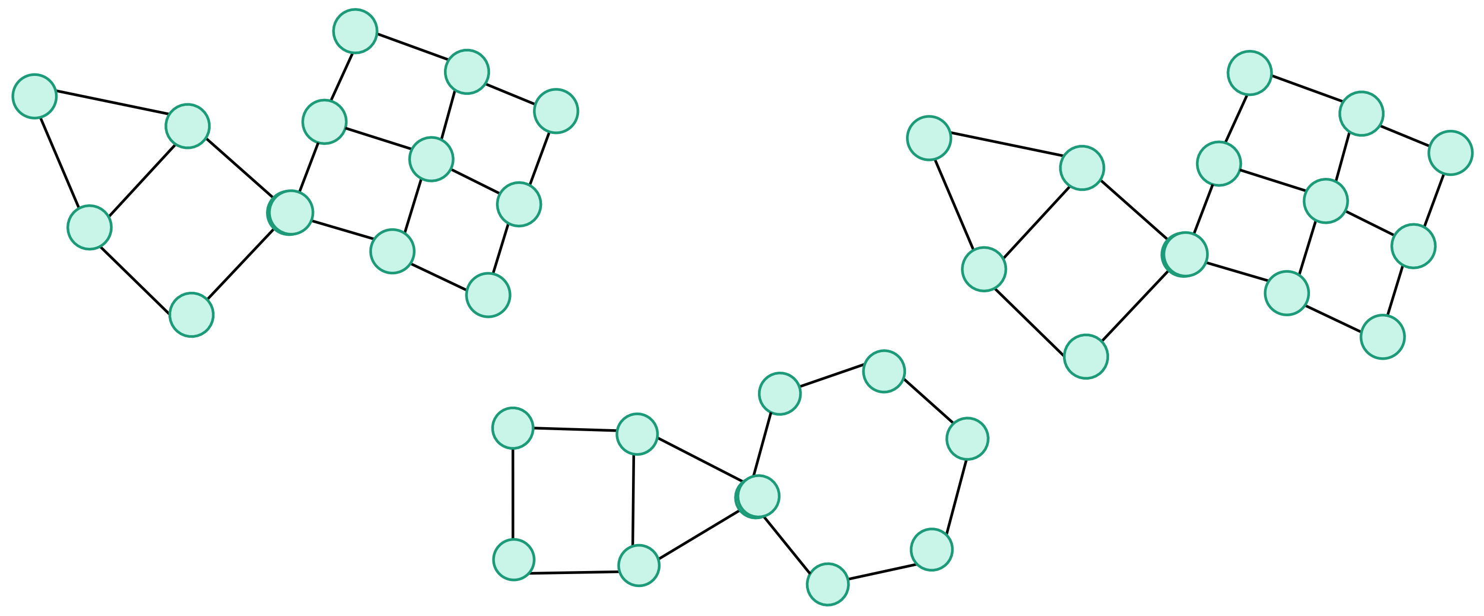

Environment #2: Class “House”

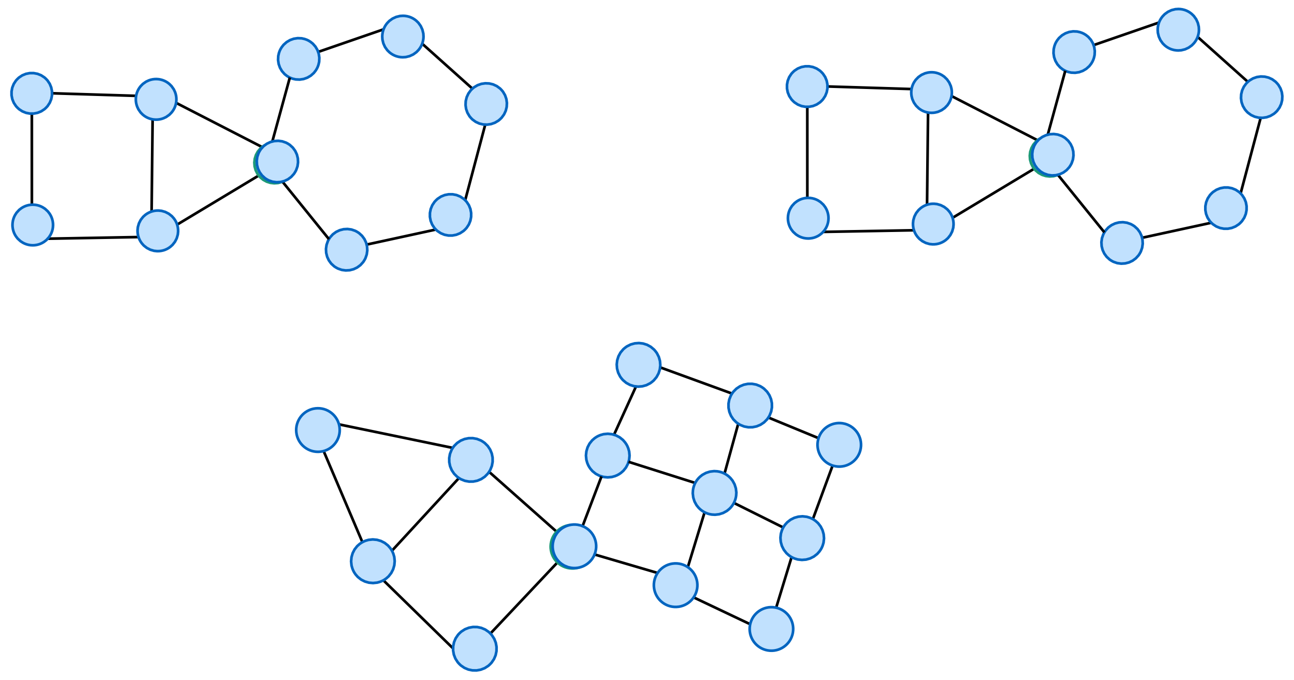

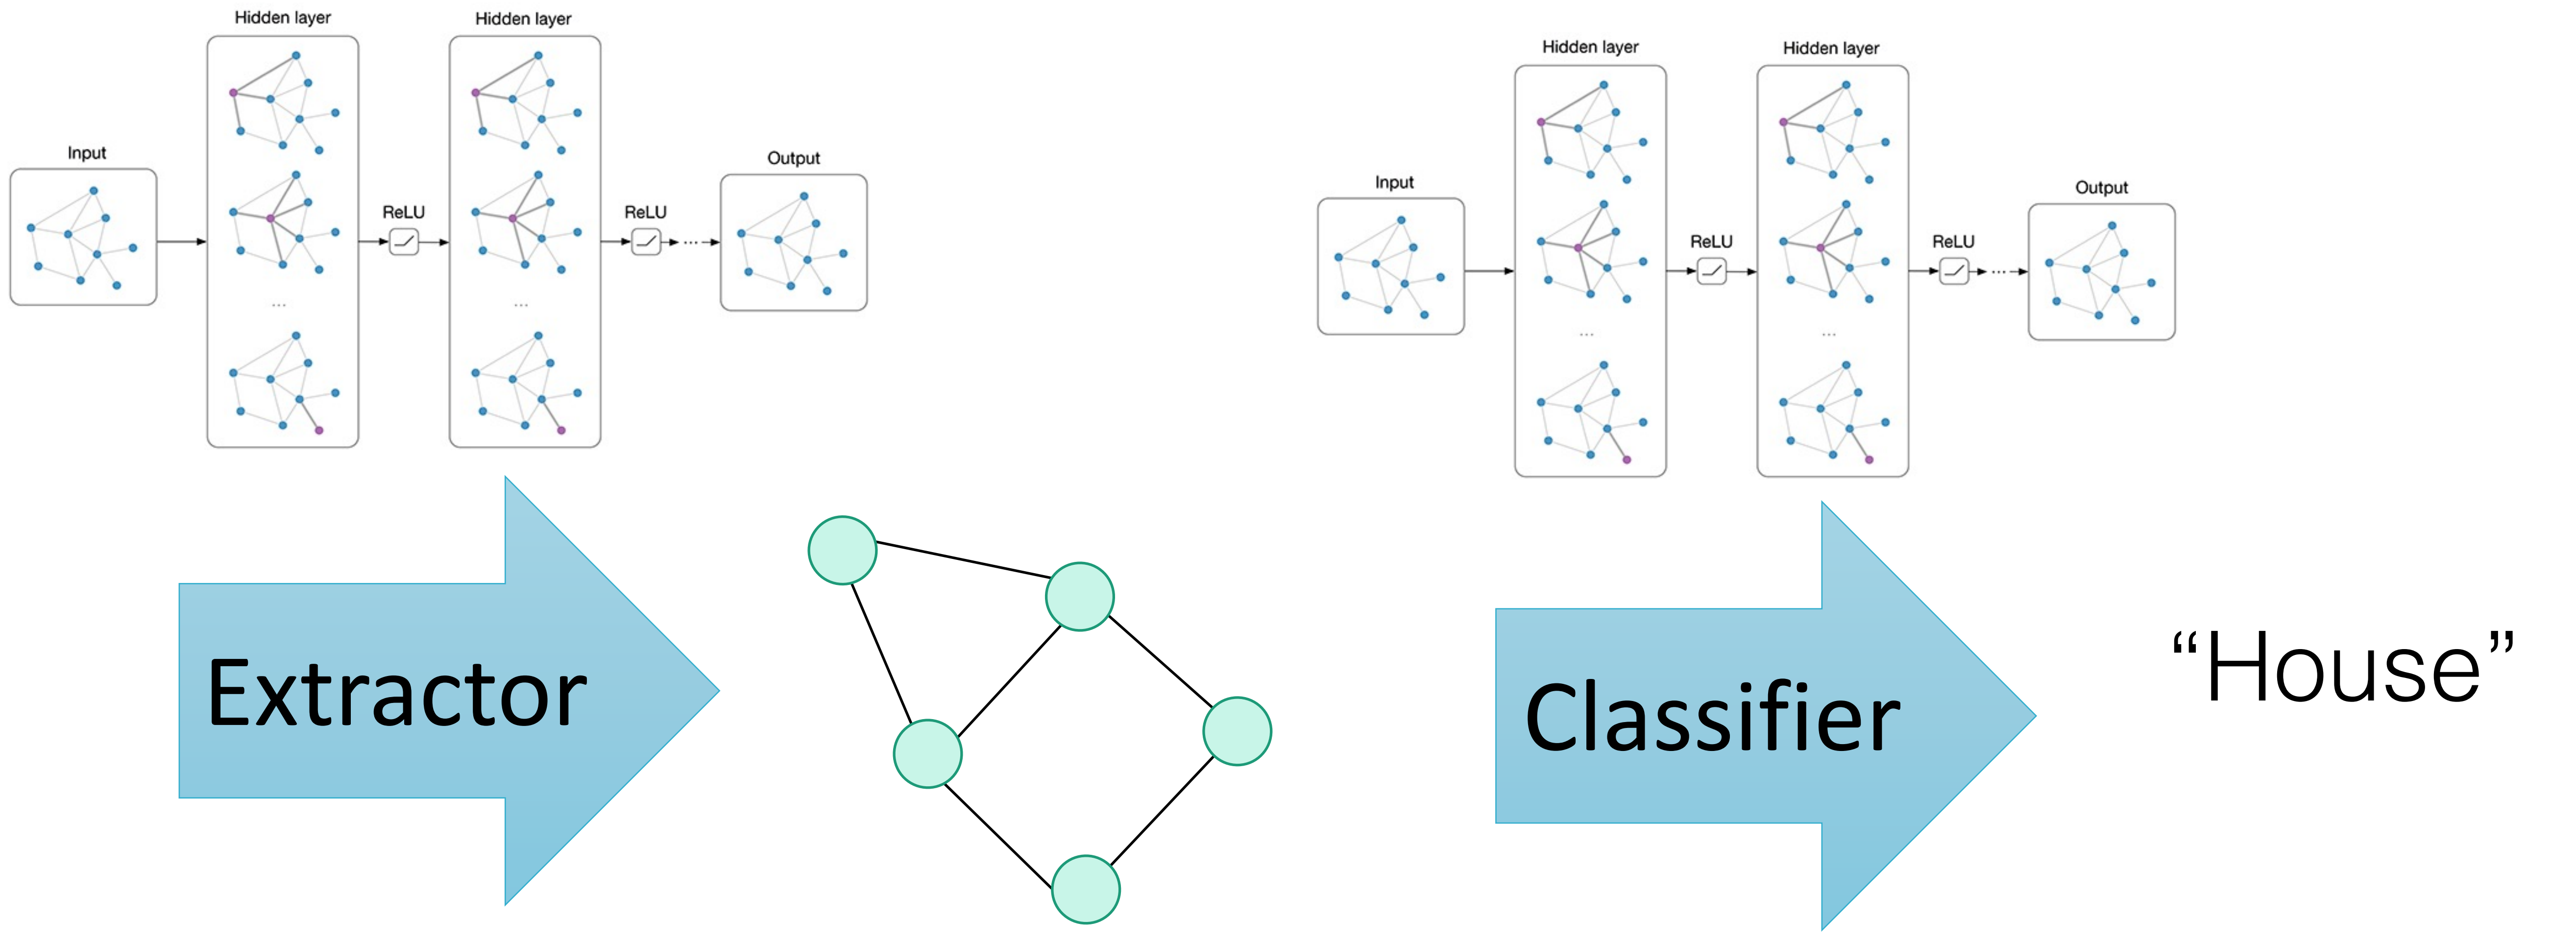

Extracted Invariant Subgraph

(Wu et al., 2022ab; Miao et al., 2022; Chen et al., 2022)

The “Free Lunch Dilemma” in OOD Generalization on Graphs

However, the **environments** or **domains**: information are usually expensive to obtain for graph structured data:

Environment #?: Class “House”

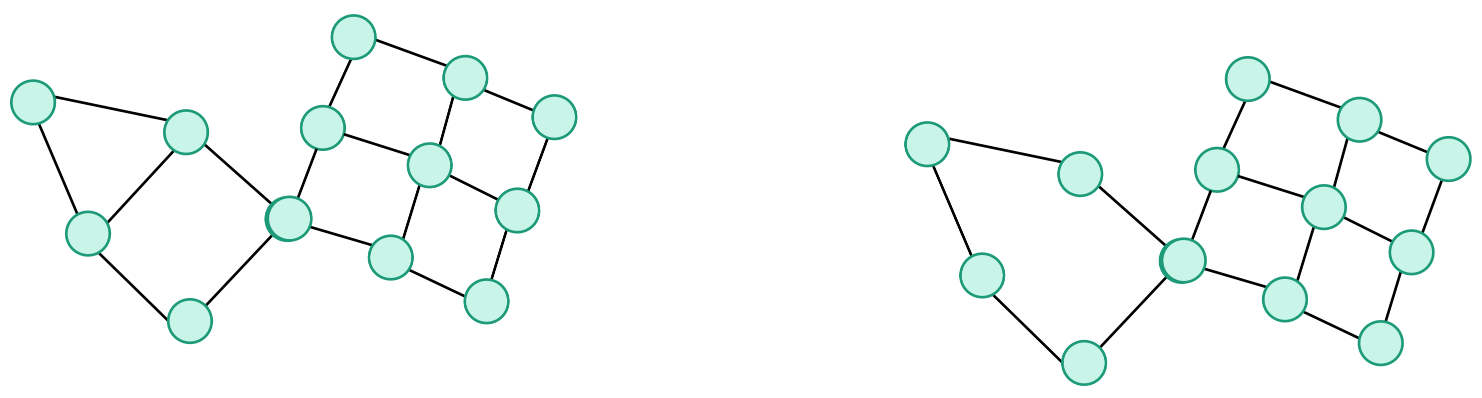

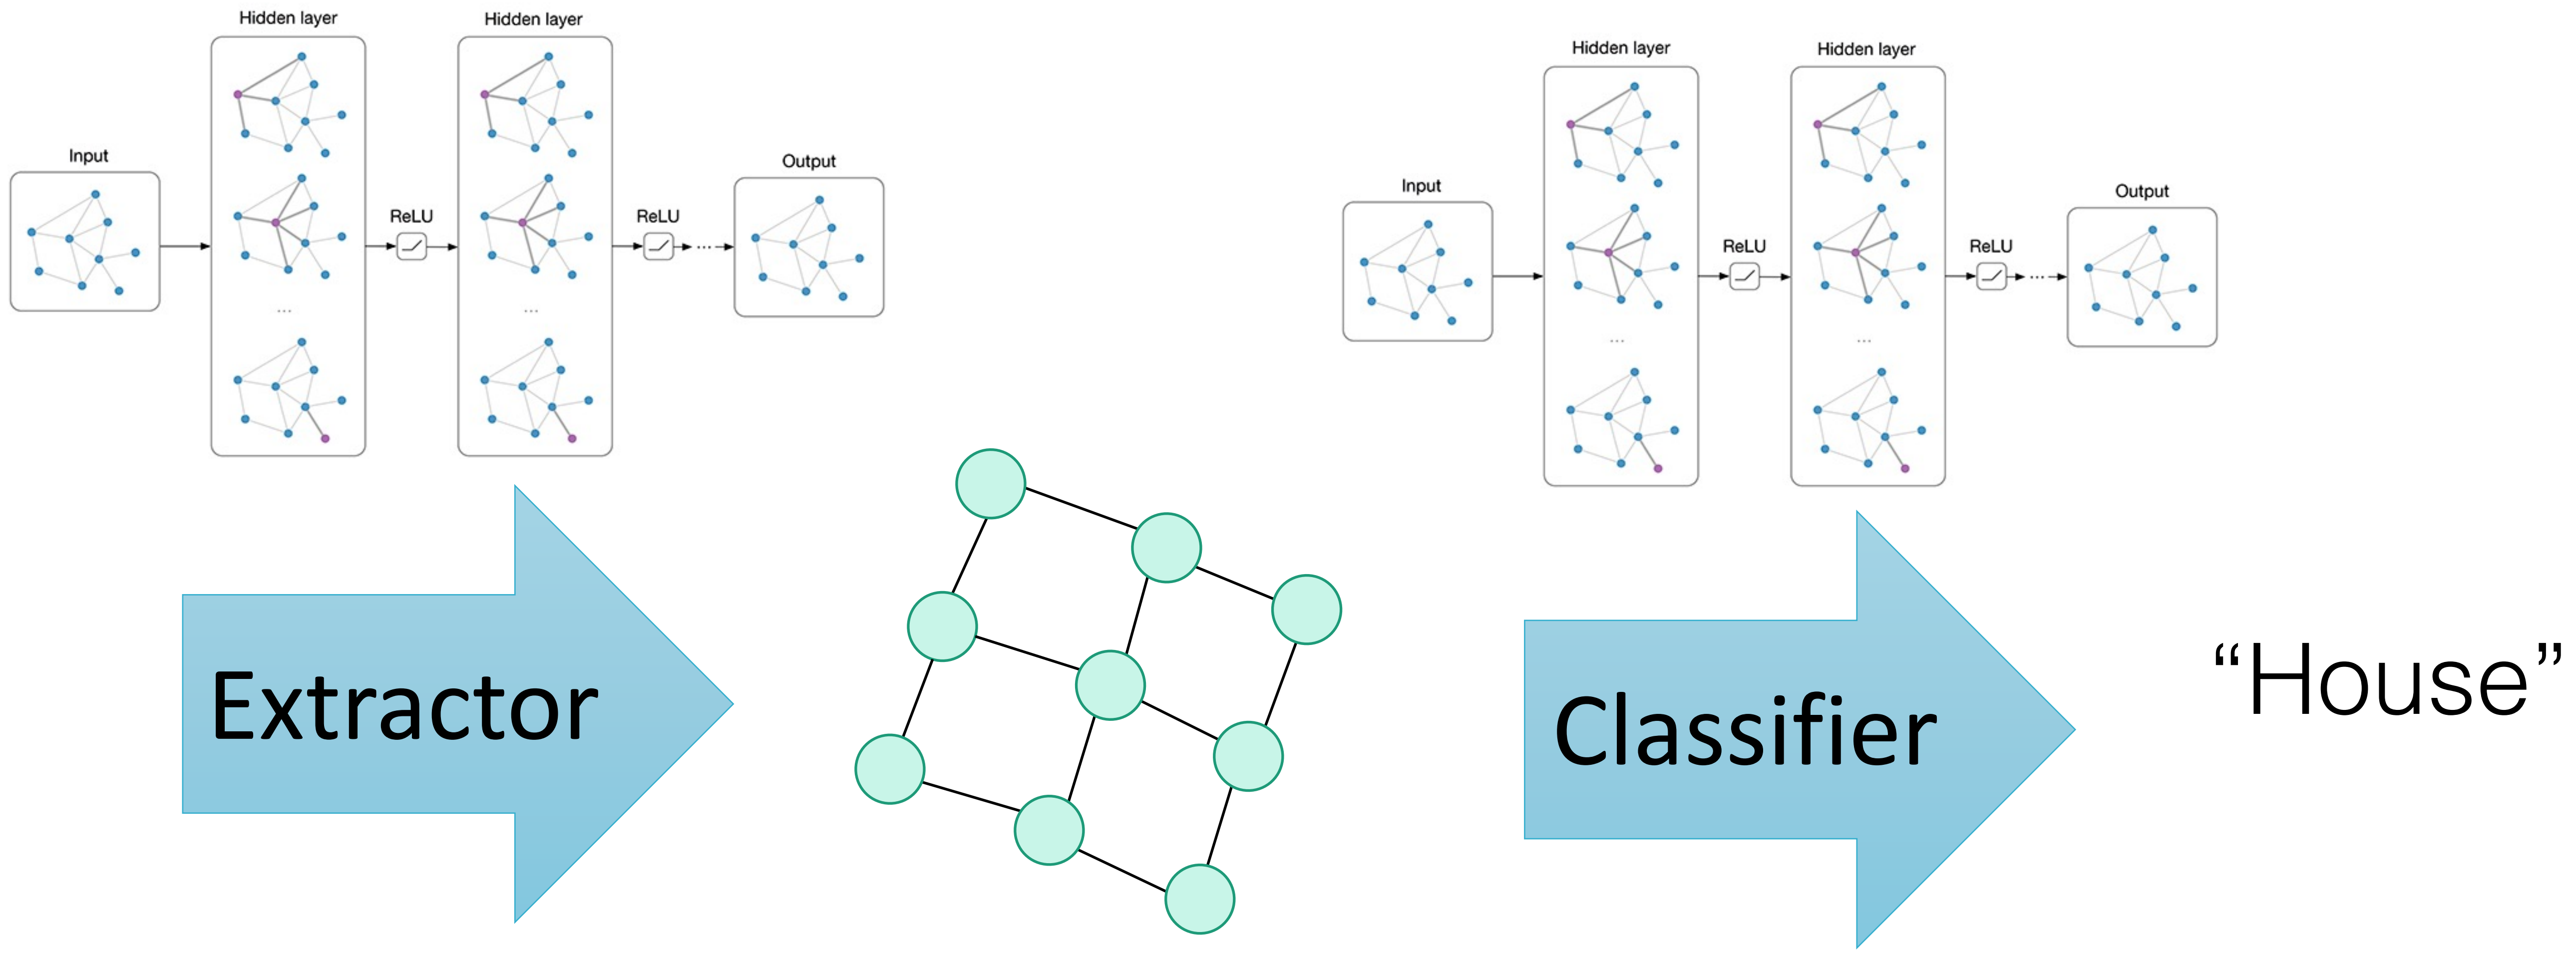

Environment #?: Class “House”

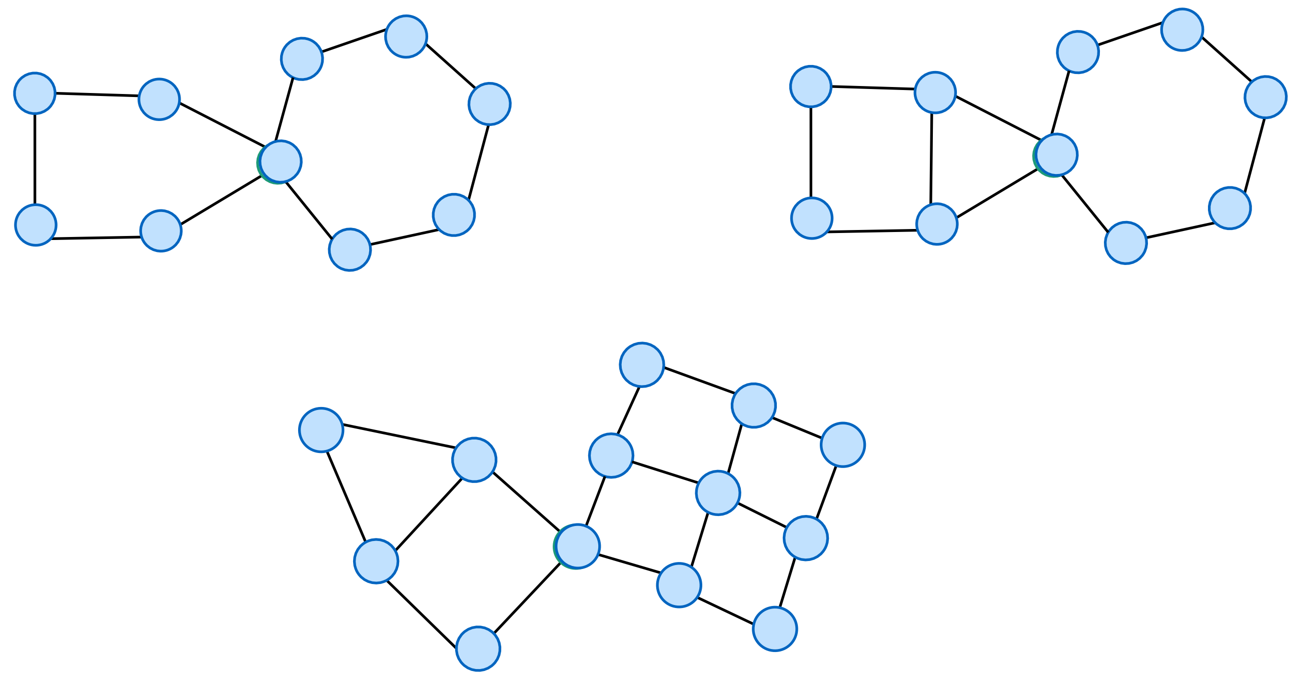

Extracted “Invariant” Subgraph

(Wu et al., 2022ab; Miao et al., 2022; Chen et al., 2022)

Is it possible to augment the environment information to enable OOD generalization on graphs?

The "Free Lunch Dilemma" in OOD Generalization on Graphs

Let us considering the data generative model with **Partial Informative Invariant Features**:

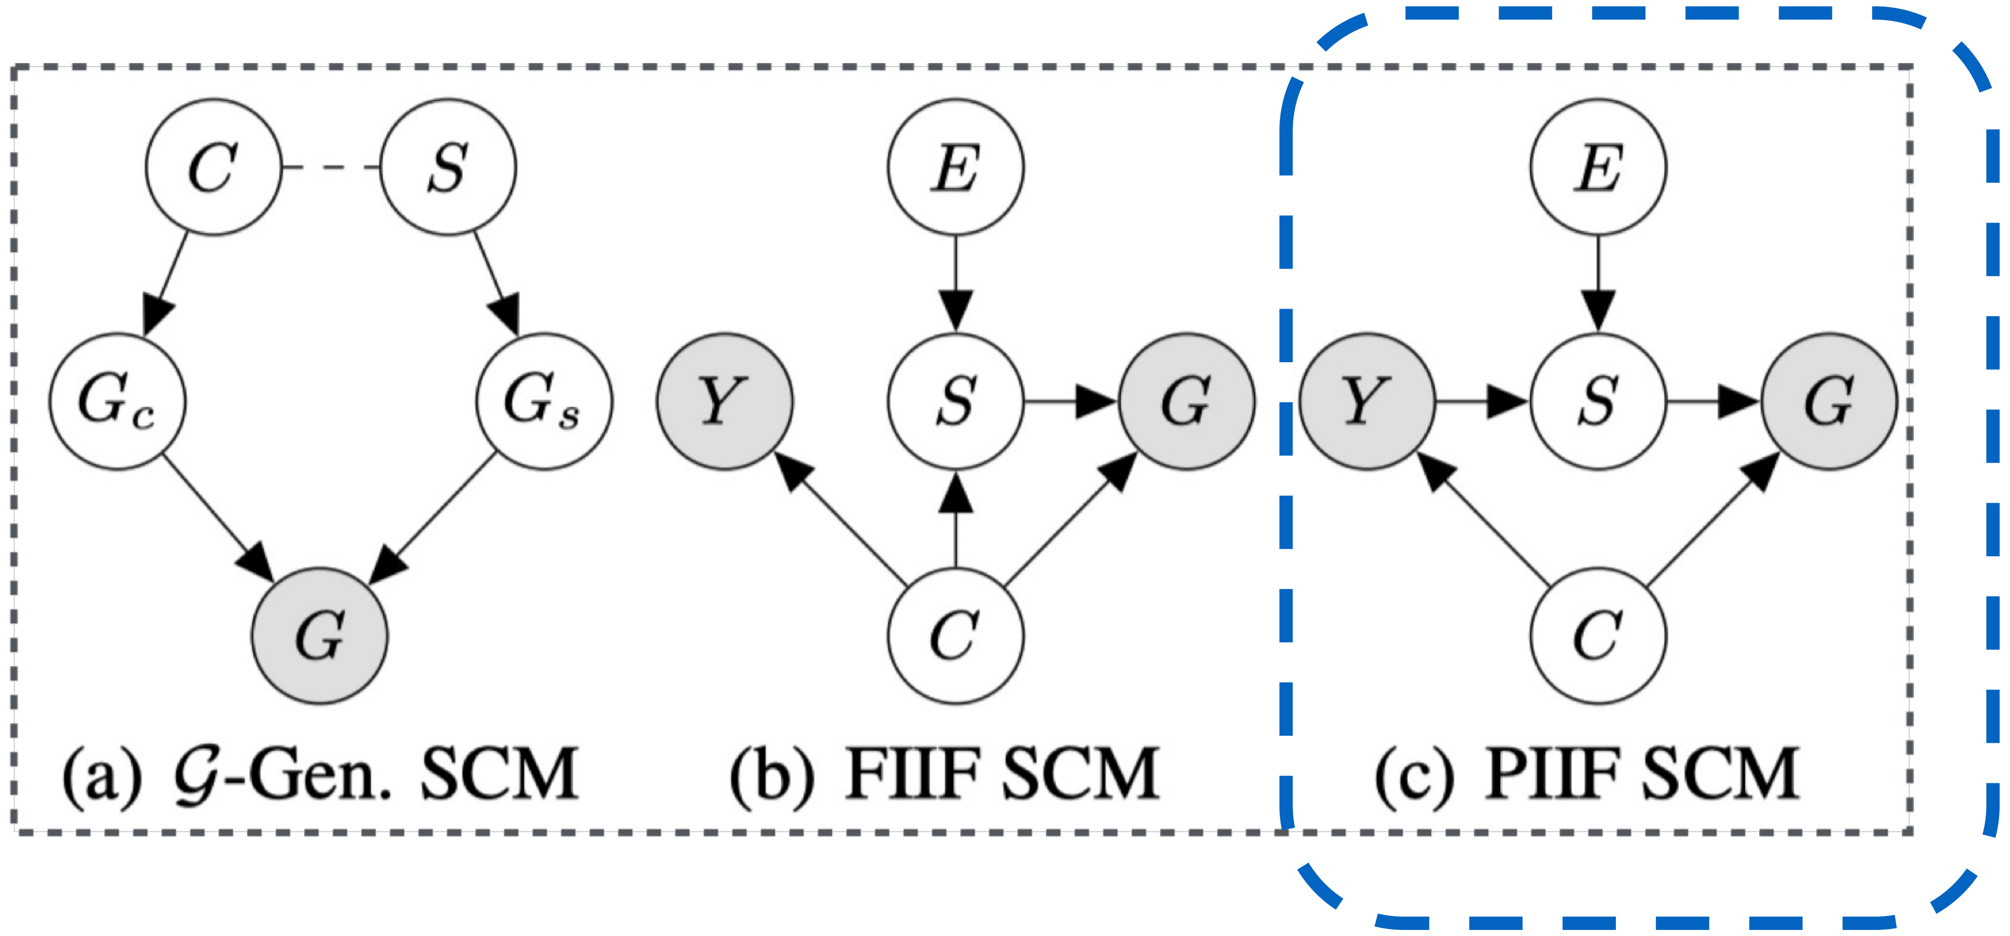

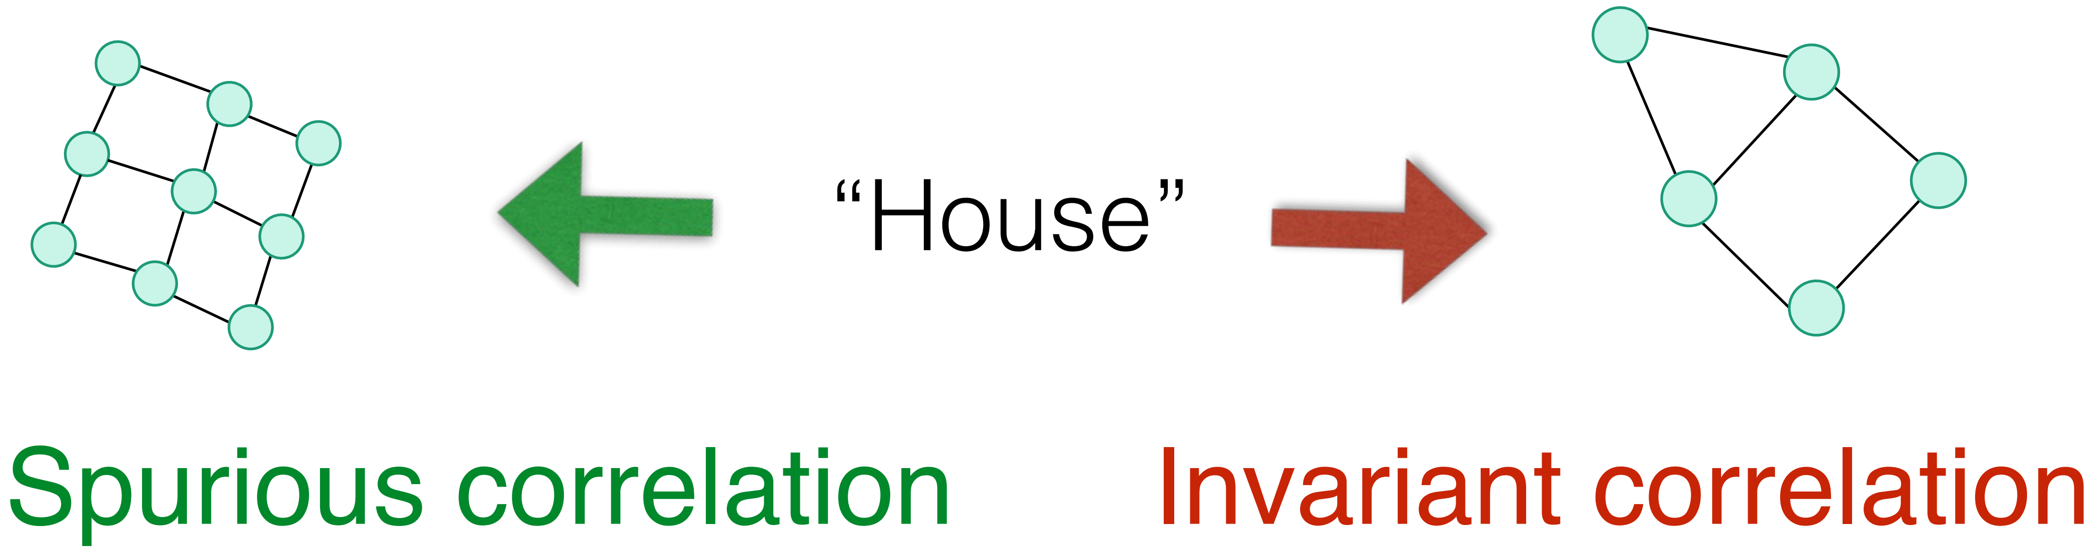

(Yu et al., 2021; Miao et al., 2022;)

The “Free Lunch Dilemma” in OOD Generalization on Graphs

One line of works aim to generate **new environments** based on the existing extracted subgraphs:

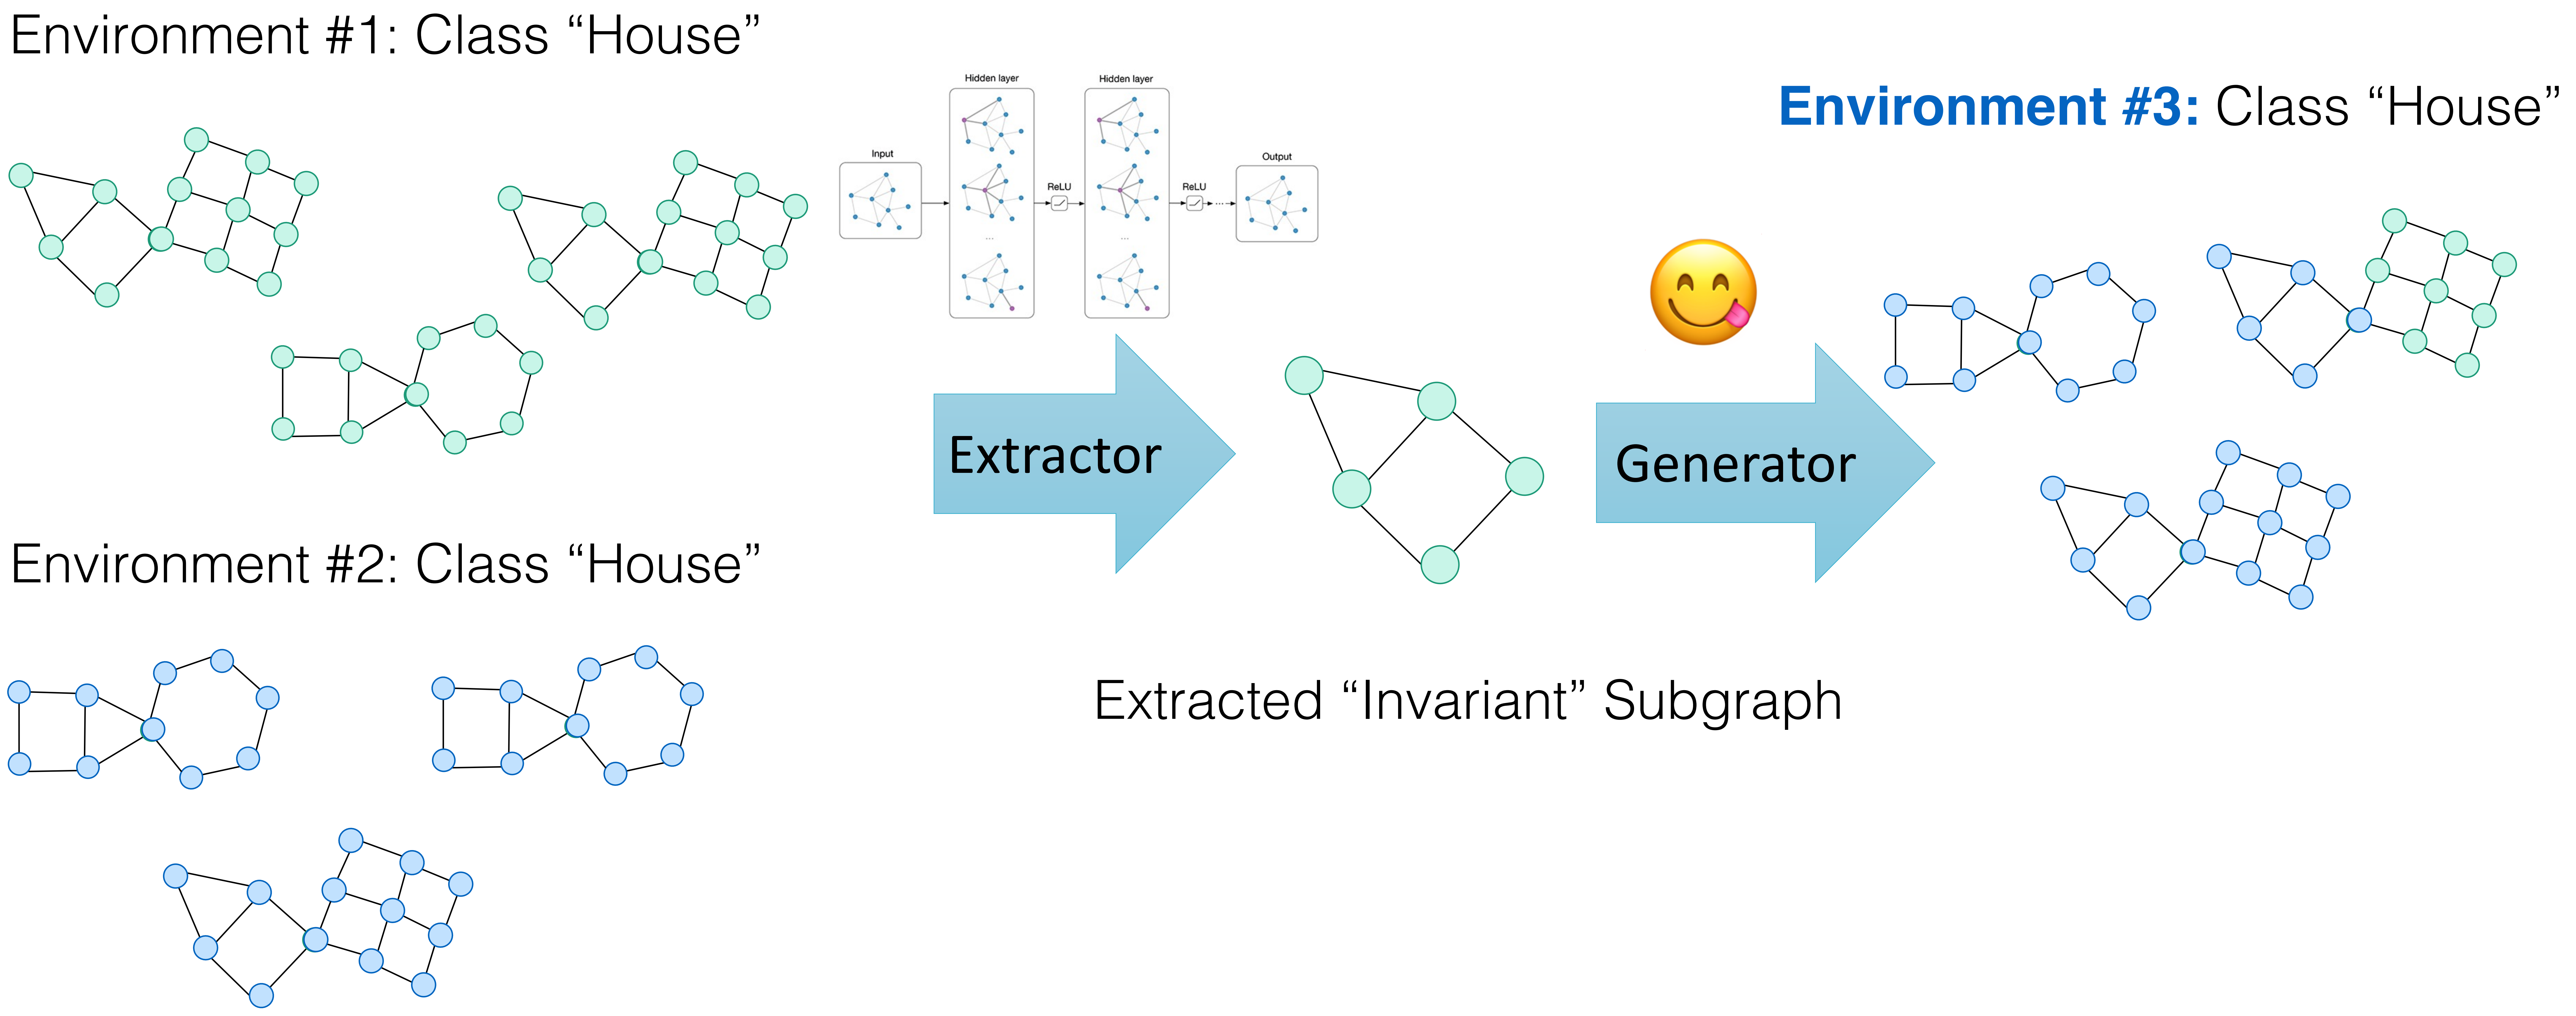

(Wu et al., 2022ab; Liu et al., 2022)

The “Free Lunch Dilemma” in OOD Generalization on Graphs

One line of works aim to generate **new environments** based on the existing extracted subgraphs:

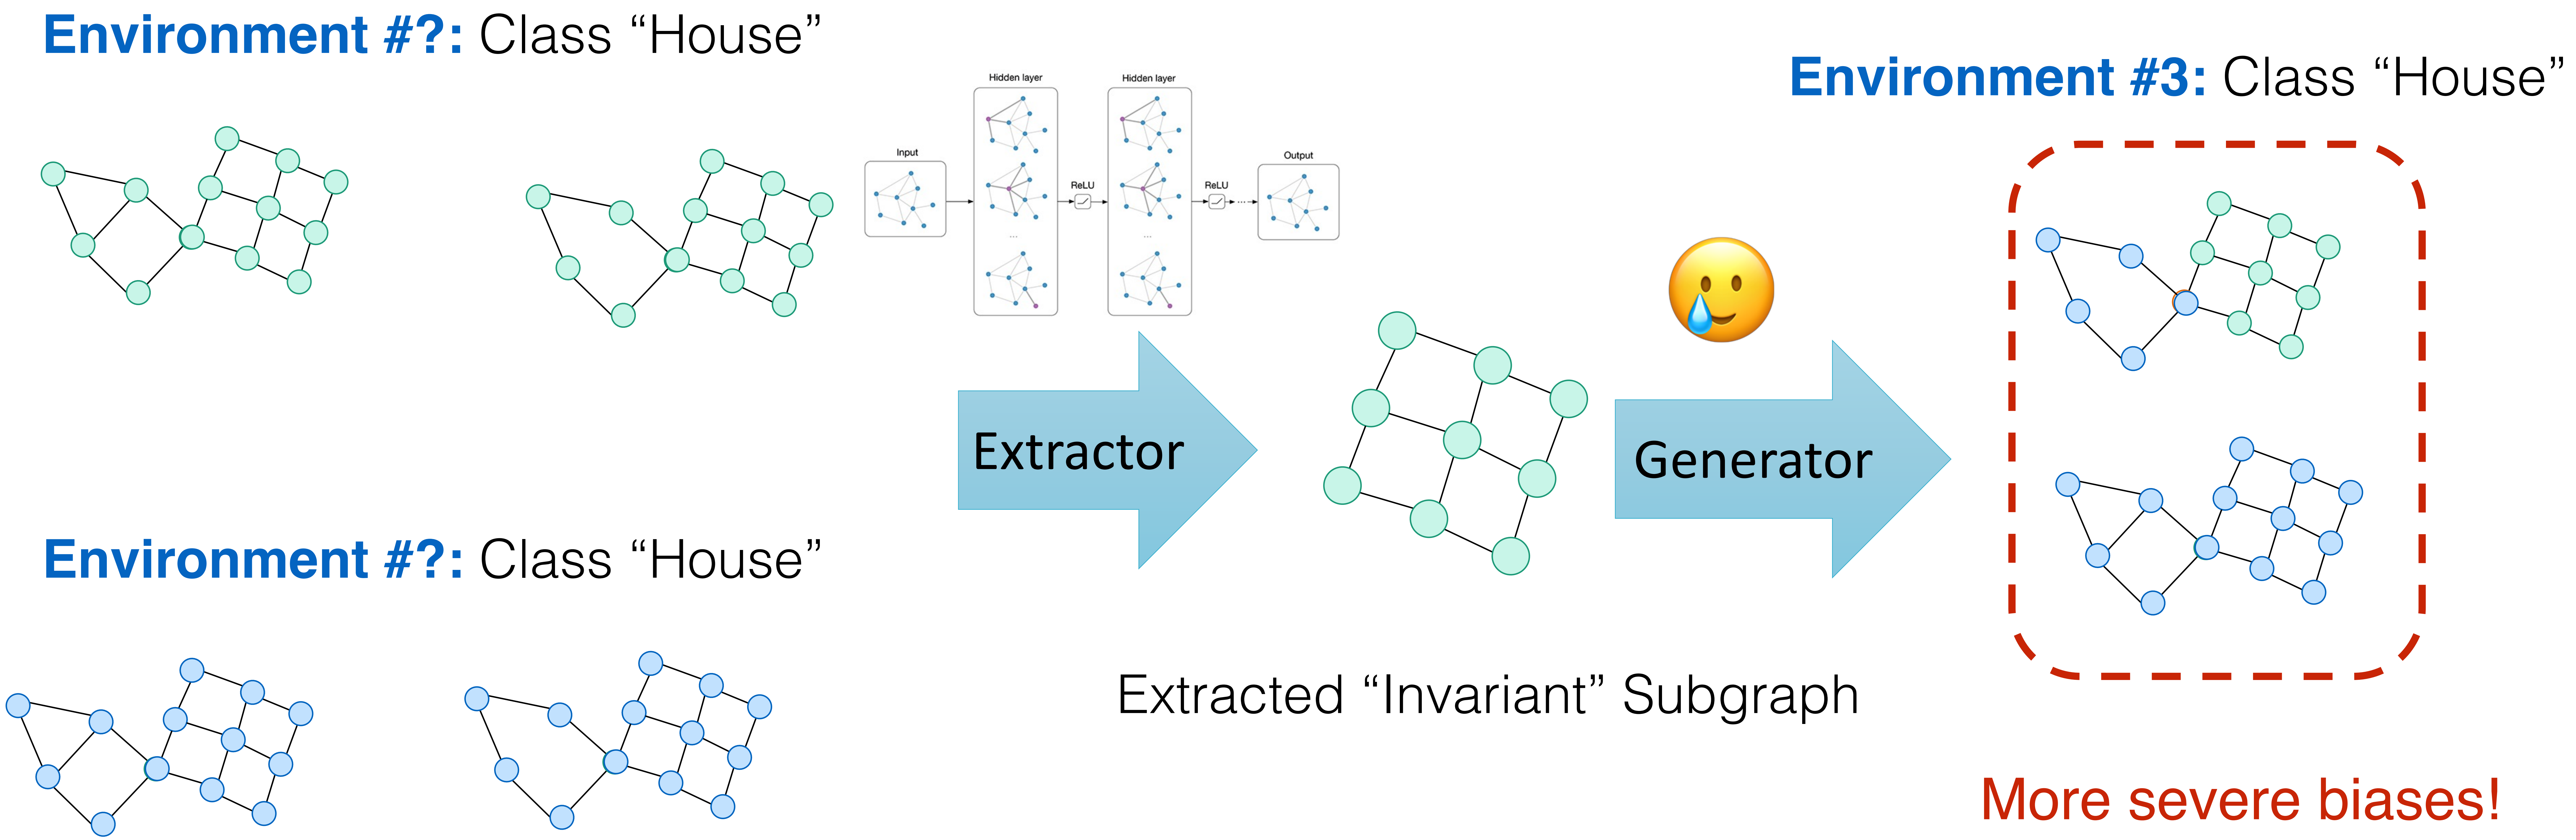

(Wu et al., 2022ab; Liu et al., 2022)

The “Free Lunch Dilemma” in OOD Generalization on Graphs

Another line of works aim to **infer environment labels** for learning the underlying invariance:

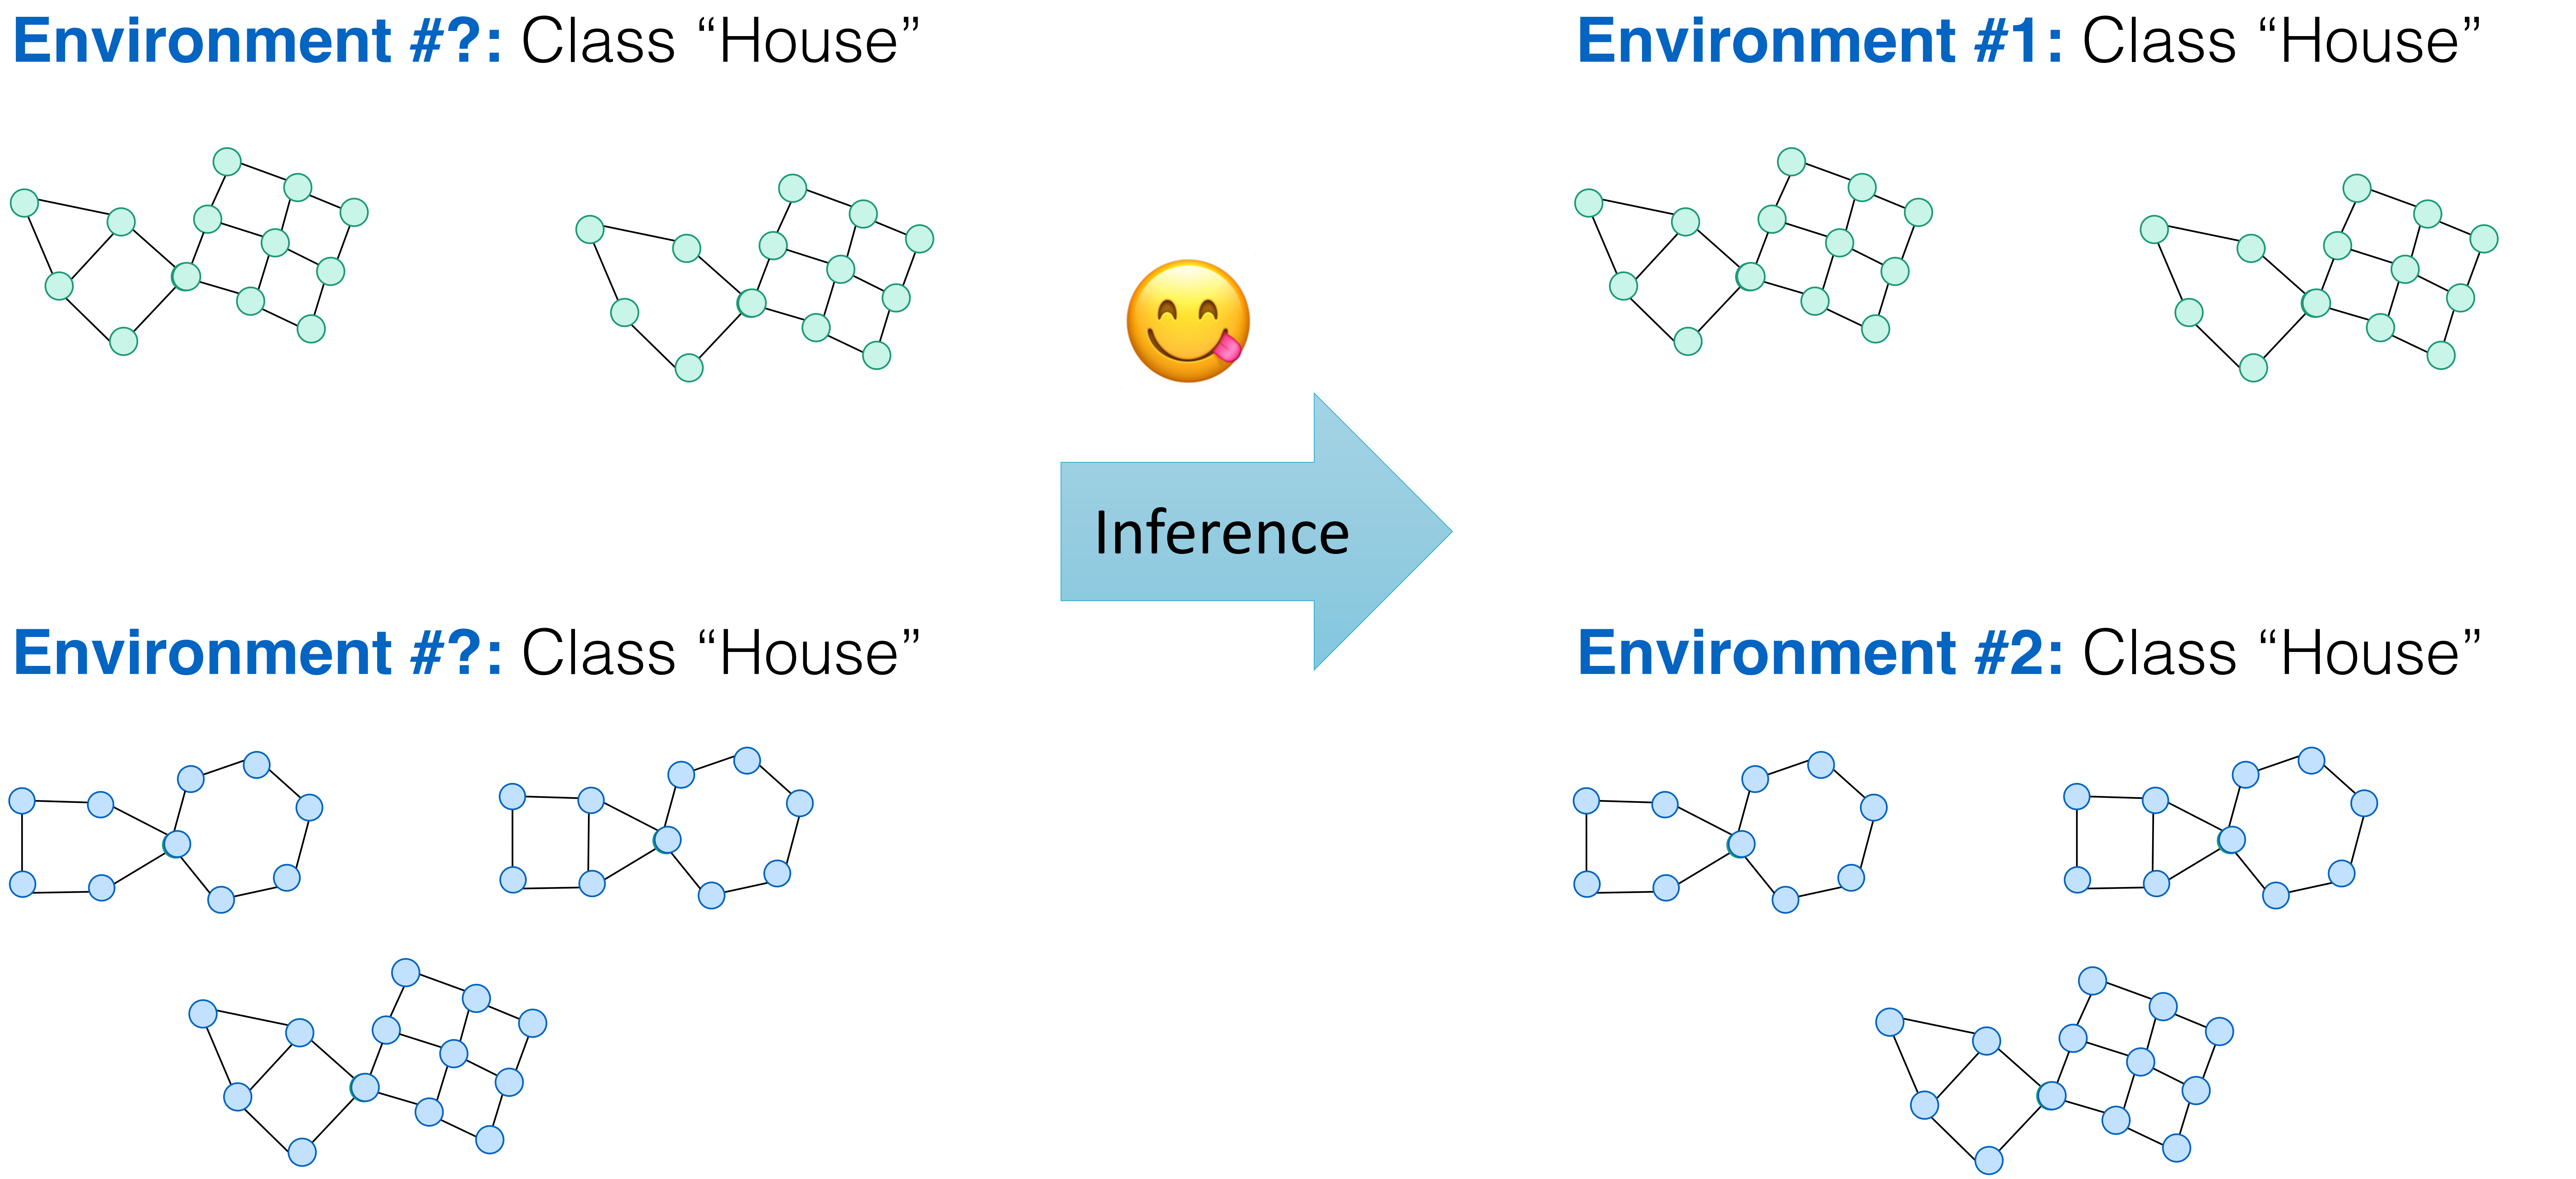

(Li et al., 2022; Yang et al., 2022)

The “Free Lunch Dilemma” in OOD Generalization on Graphs

Another line of works aim to **infer environment labels** for learning the underlying invariance:

Environment #?: Class “Grid”

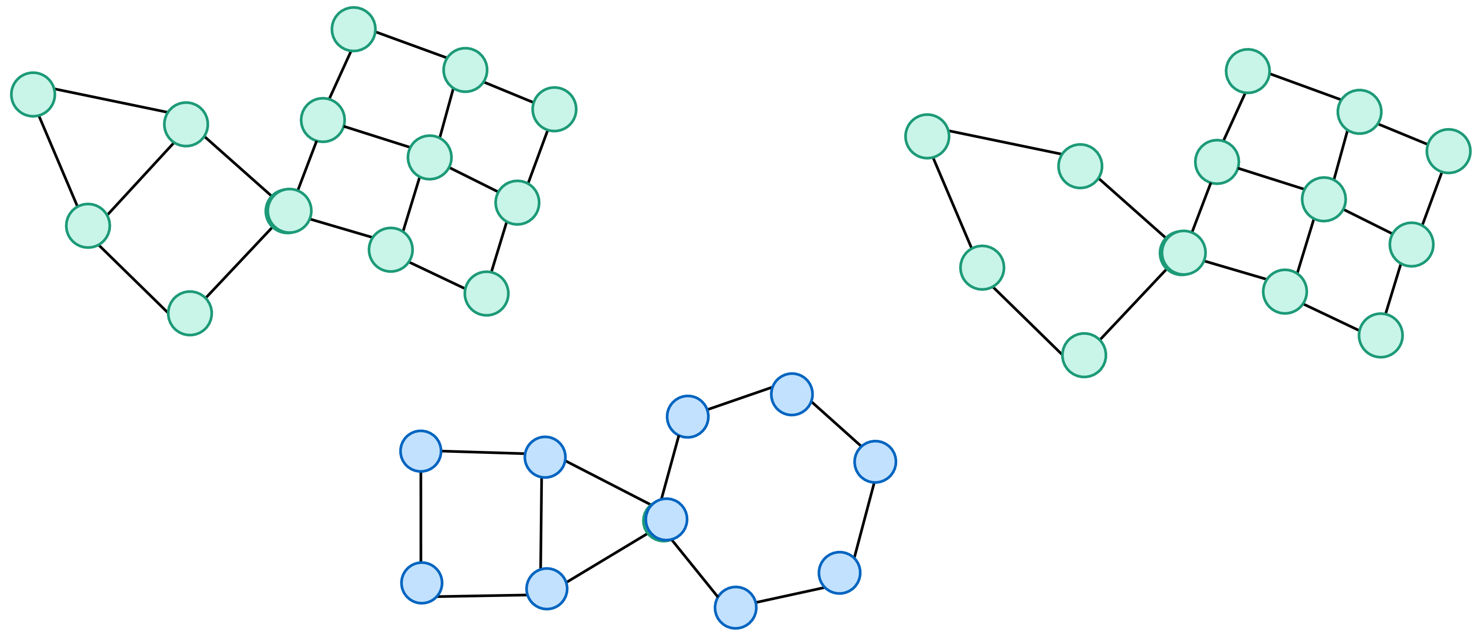

Environment #?: Class “Grid”

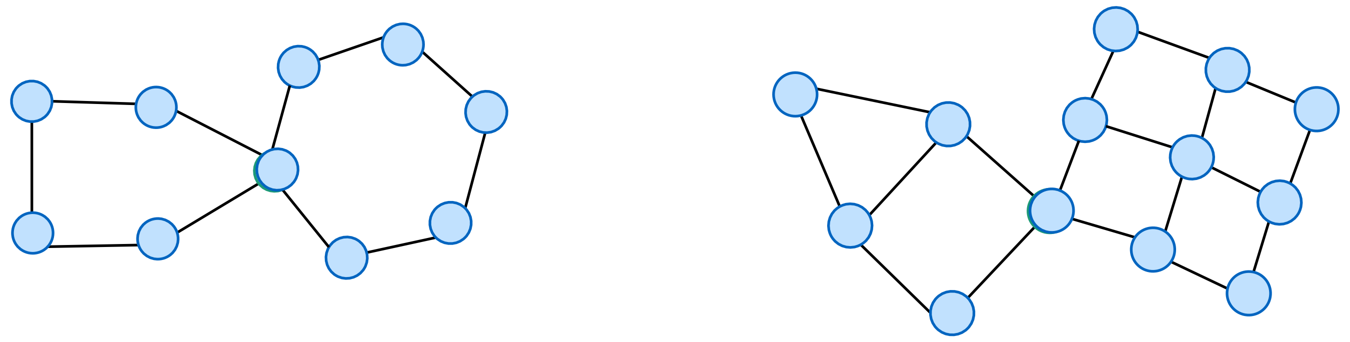

Environment #1: Class “Grid”

Environment #2: Class “Grid”

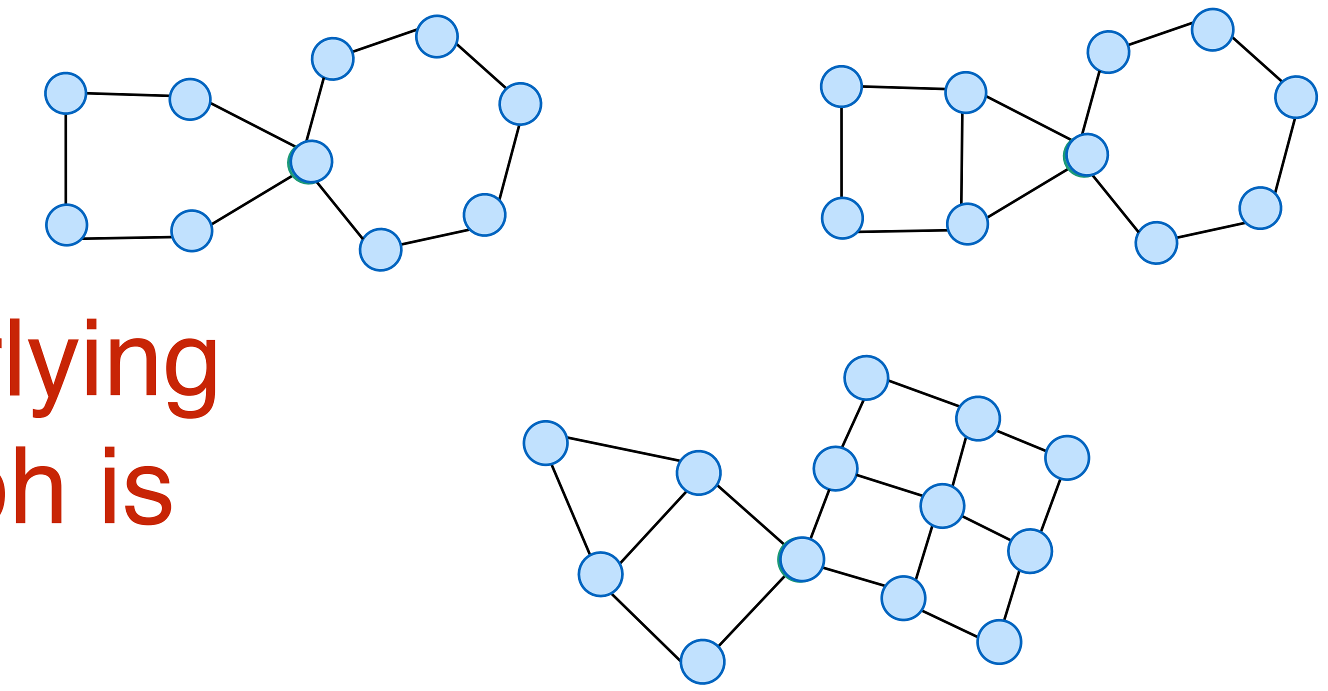

What if the underlying invariant subgraph is REVERSED?

(*Li et al., 2022; Yang et al., 2022)*

Impossibility Results for OOD Generalization on Graphs

OOD generalization on graphs is fundamentally more **challenging** than that on Euclidean data:

Environment #?: Class “House”

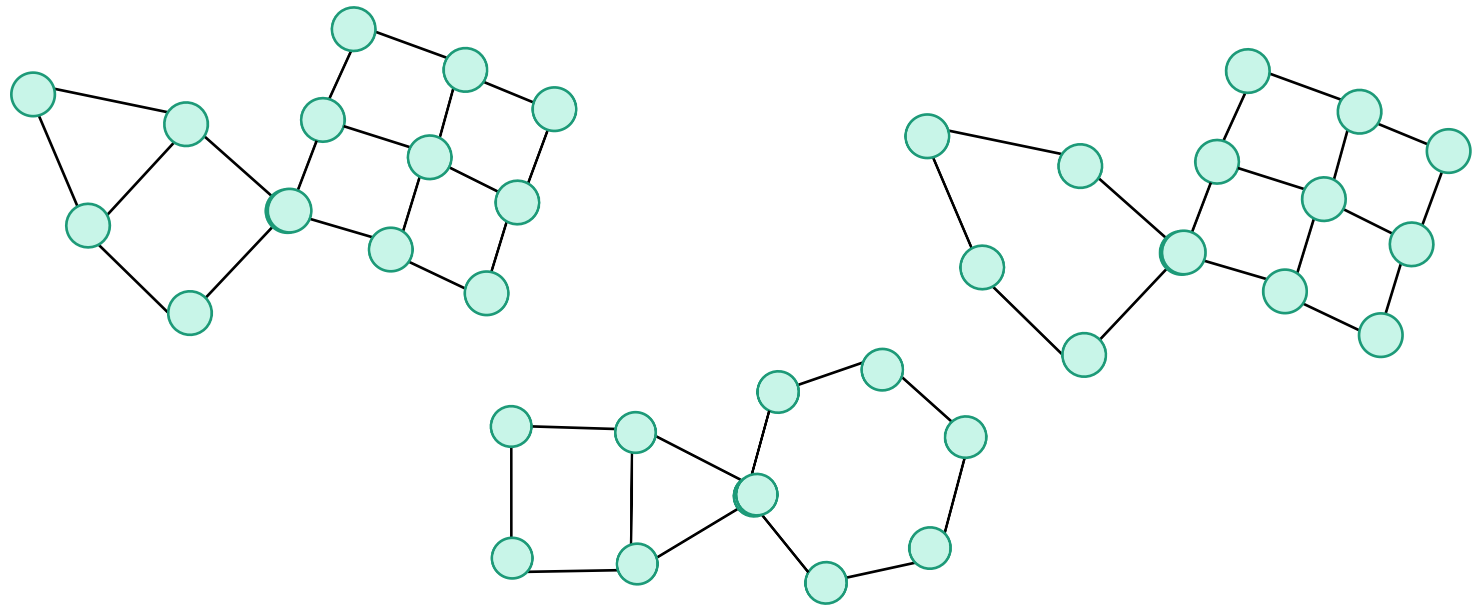

Environment #?: Class “House”

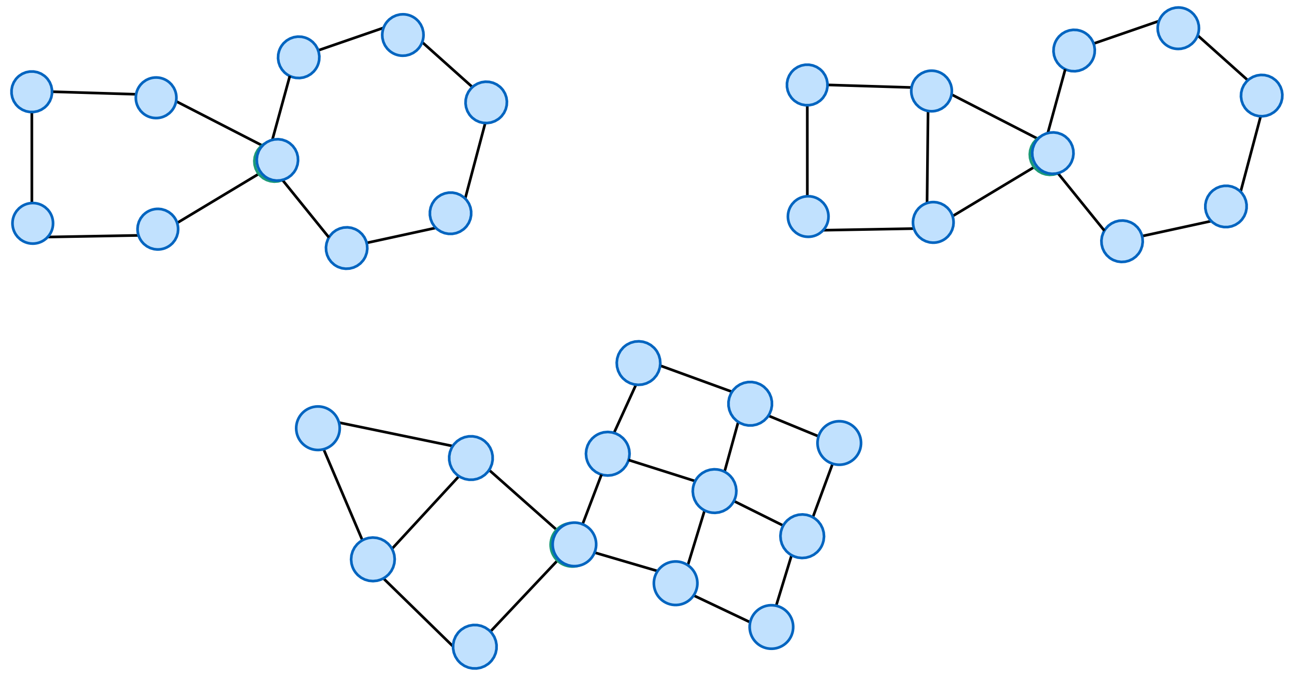

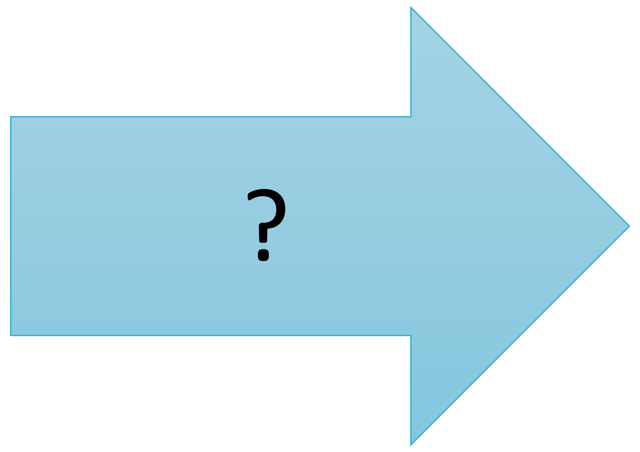

No Free Lunch in Graph OOD (Informal)
It is fundamentally ***impossible*** to identify the underlying invariant subgraph without further inductive biases.

What are the minimal sufficient inductive biases for invariant graph representation learning?

Failures of Environment Generation

How can we address **environments generation** failures?

Environment #?: Class “House”

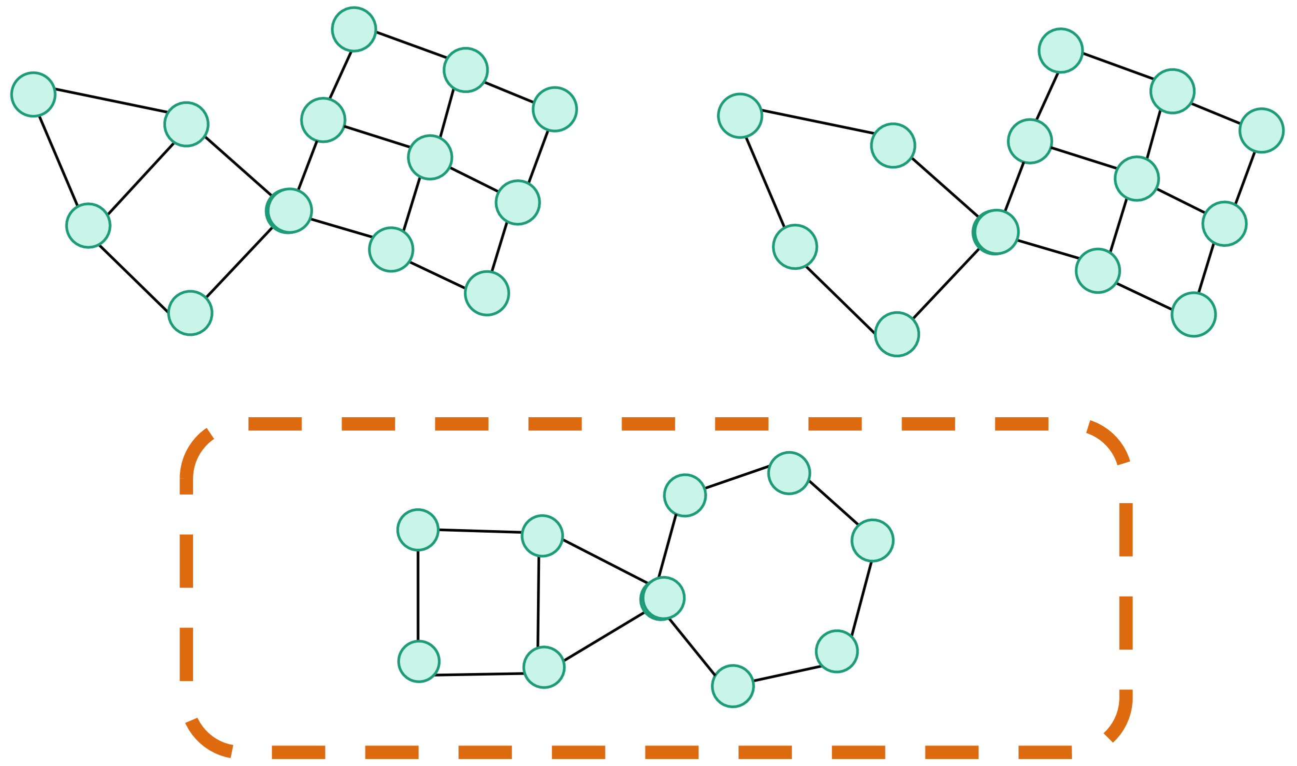

Environment #?: Class “House”

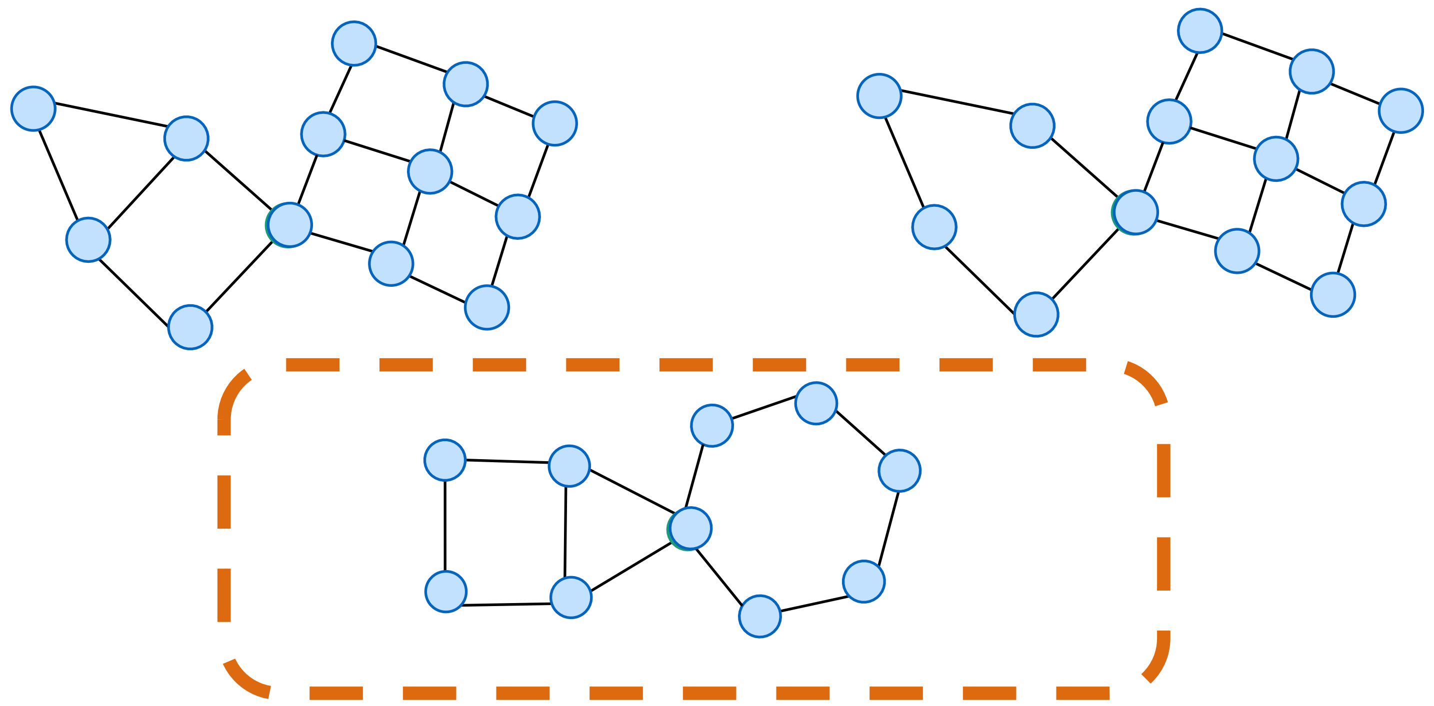

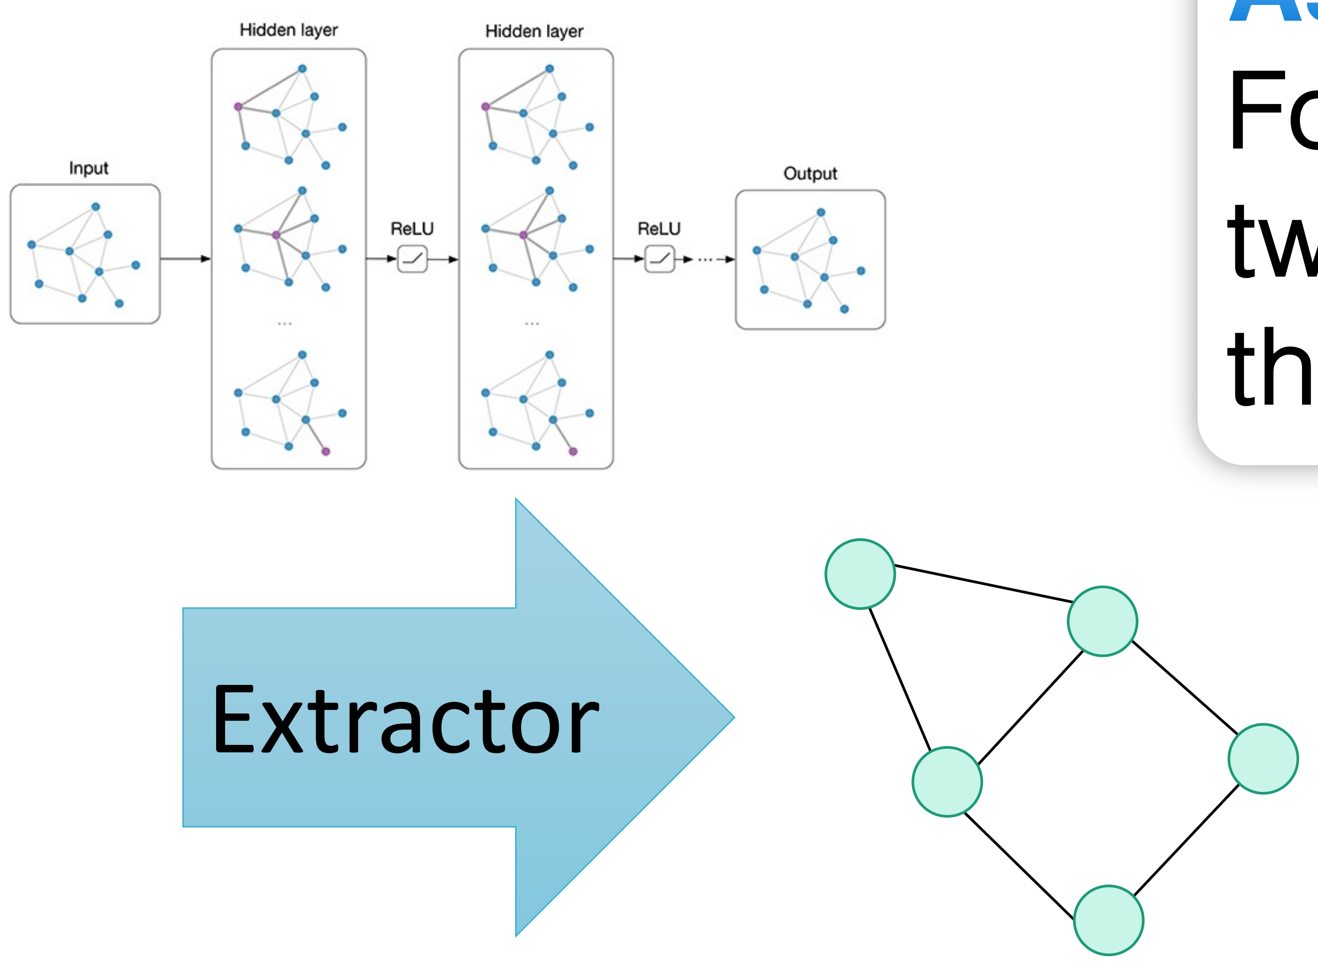

Extracted Invariant Subgraph

Assumption 1 (Variation Sufficiency)
For any spurious subgraph, there exists two underlying environments, such that the spurious correlation varies.

(Wu et al., 2022ab; Liu et al., 2022)

Failures of Environment Inference

How can we address **environment inference** failures?

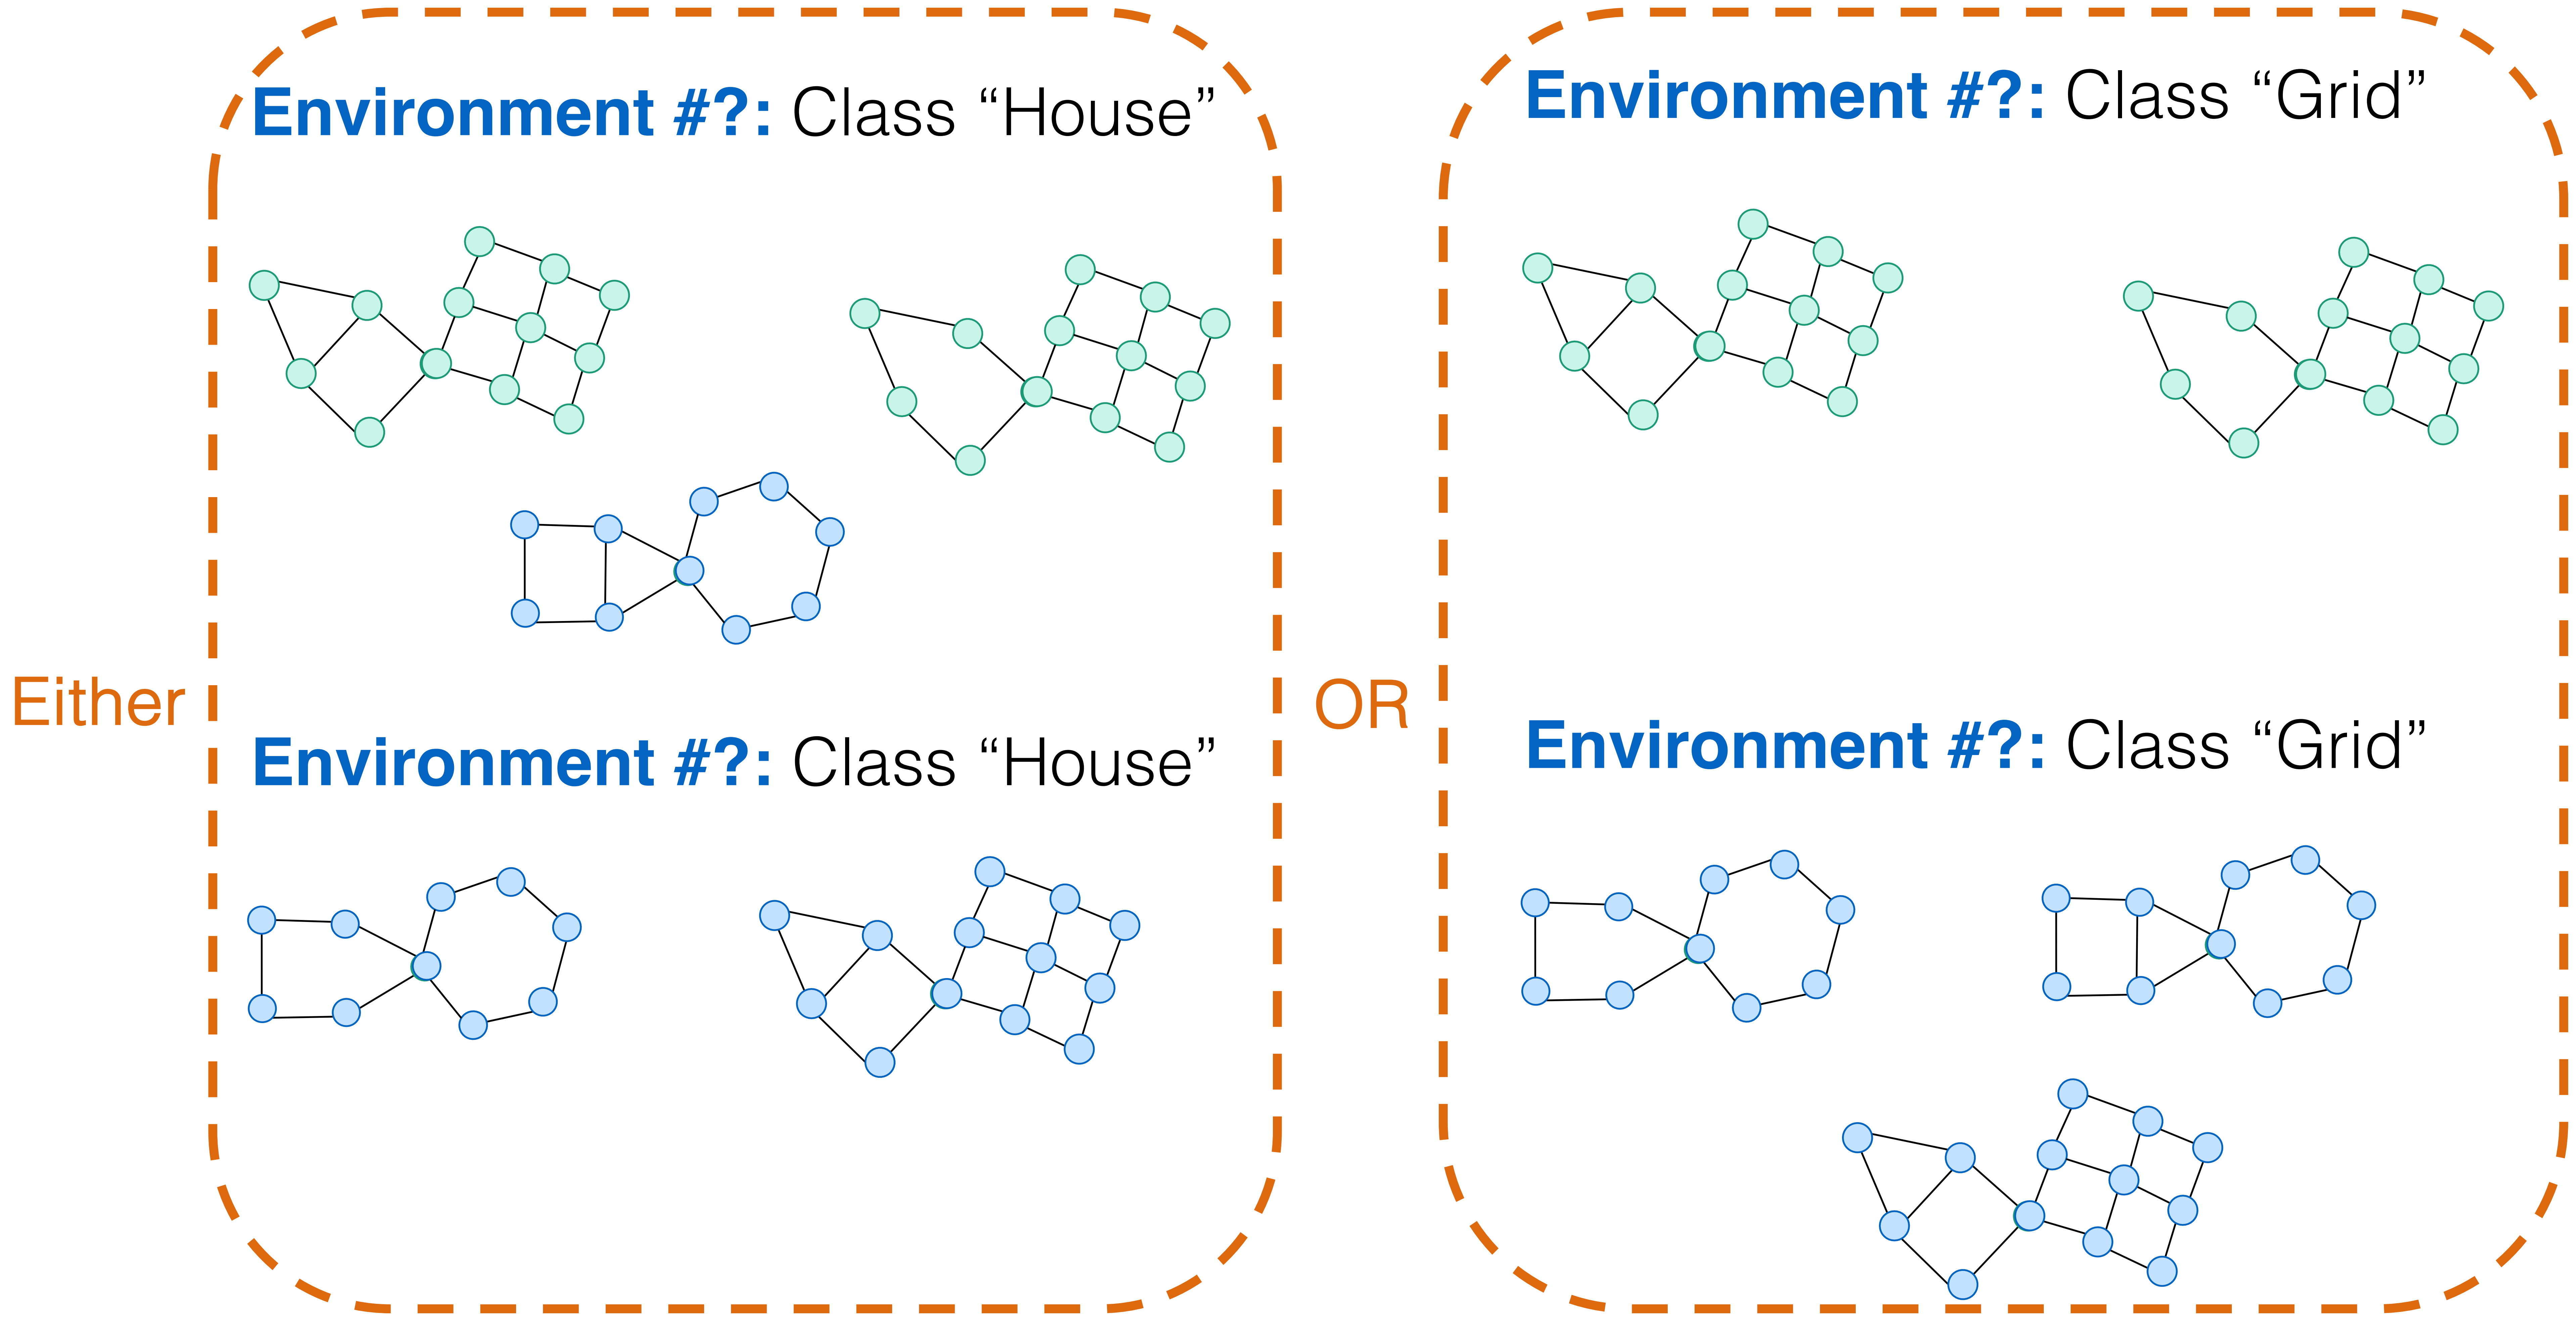

Assumption 2 (Variation Consistency)
For all environments, either spurious correlation is stronger or weaker.

(Li et al., 2022; Yang et al., 2022)

Invariant Graph Learning with Minimal Assumptions

How can we address **environment inference** failures?

Assumption 1 (Variation Sufficiency)
For any spurious subgraph, there exists two underlying environments, such that the spurious correlation varies.

Environment Generation?

Environment Inference?

*More assumptions needed!

Assumption 2 (Variation Consistency)
For all environments, either spurious correlation is stronger or weaker.

ZIN: When and How to Learn Invariance by Environment Inference?

Invariant Graph Learning with Minimal Assumptions

How can we address **environment inference** failures?

Assumption 1 (Variation Sufficiency)
For any spurious subgraph, there exists two underlying environments, such that the spurious correlation varies.

Environment Generation?

More assumptions needed!

Environment Inference?

Assumption 2 (Variation Consistency)
For all environments, either spurious correlation is stronger or weaker.

Spurious correlation stronger:
DisC (Fan et al., 2022)

Invariant correlation stronger:
CIGA (Chen et al., 2022)

(Lin et al., 2022; Fan et al., 2022; Chen et al., 2022)

GALA: invariant GrAph Learning Assistant

To begin with, we need to first understand the reasons for the failures of CIGA:

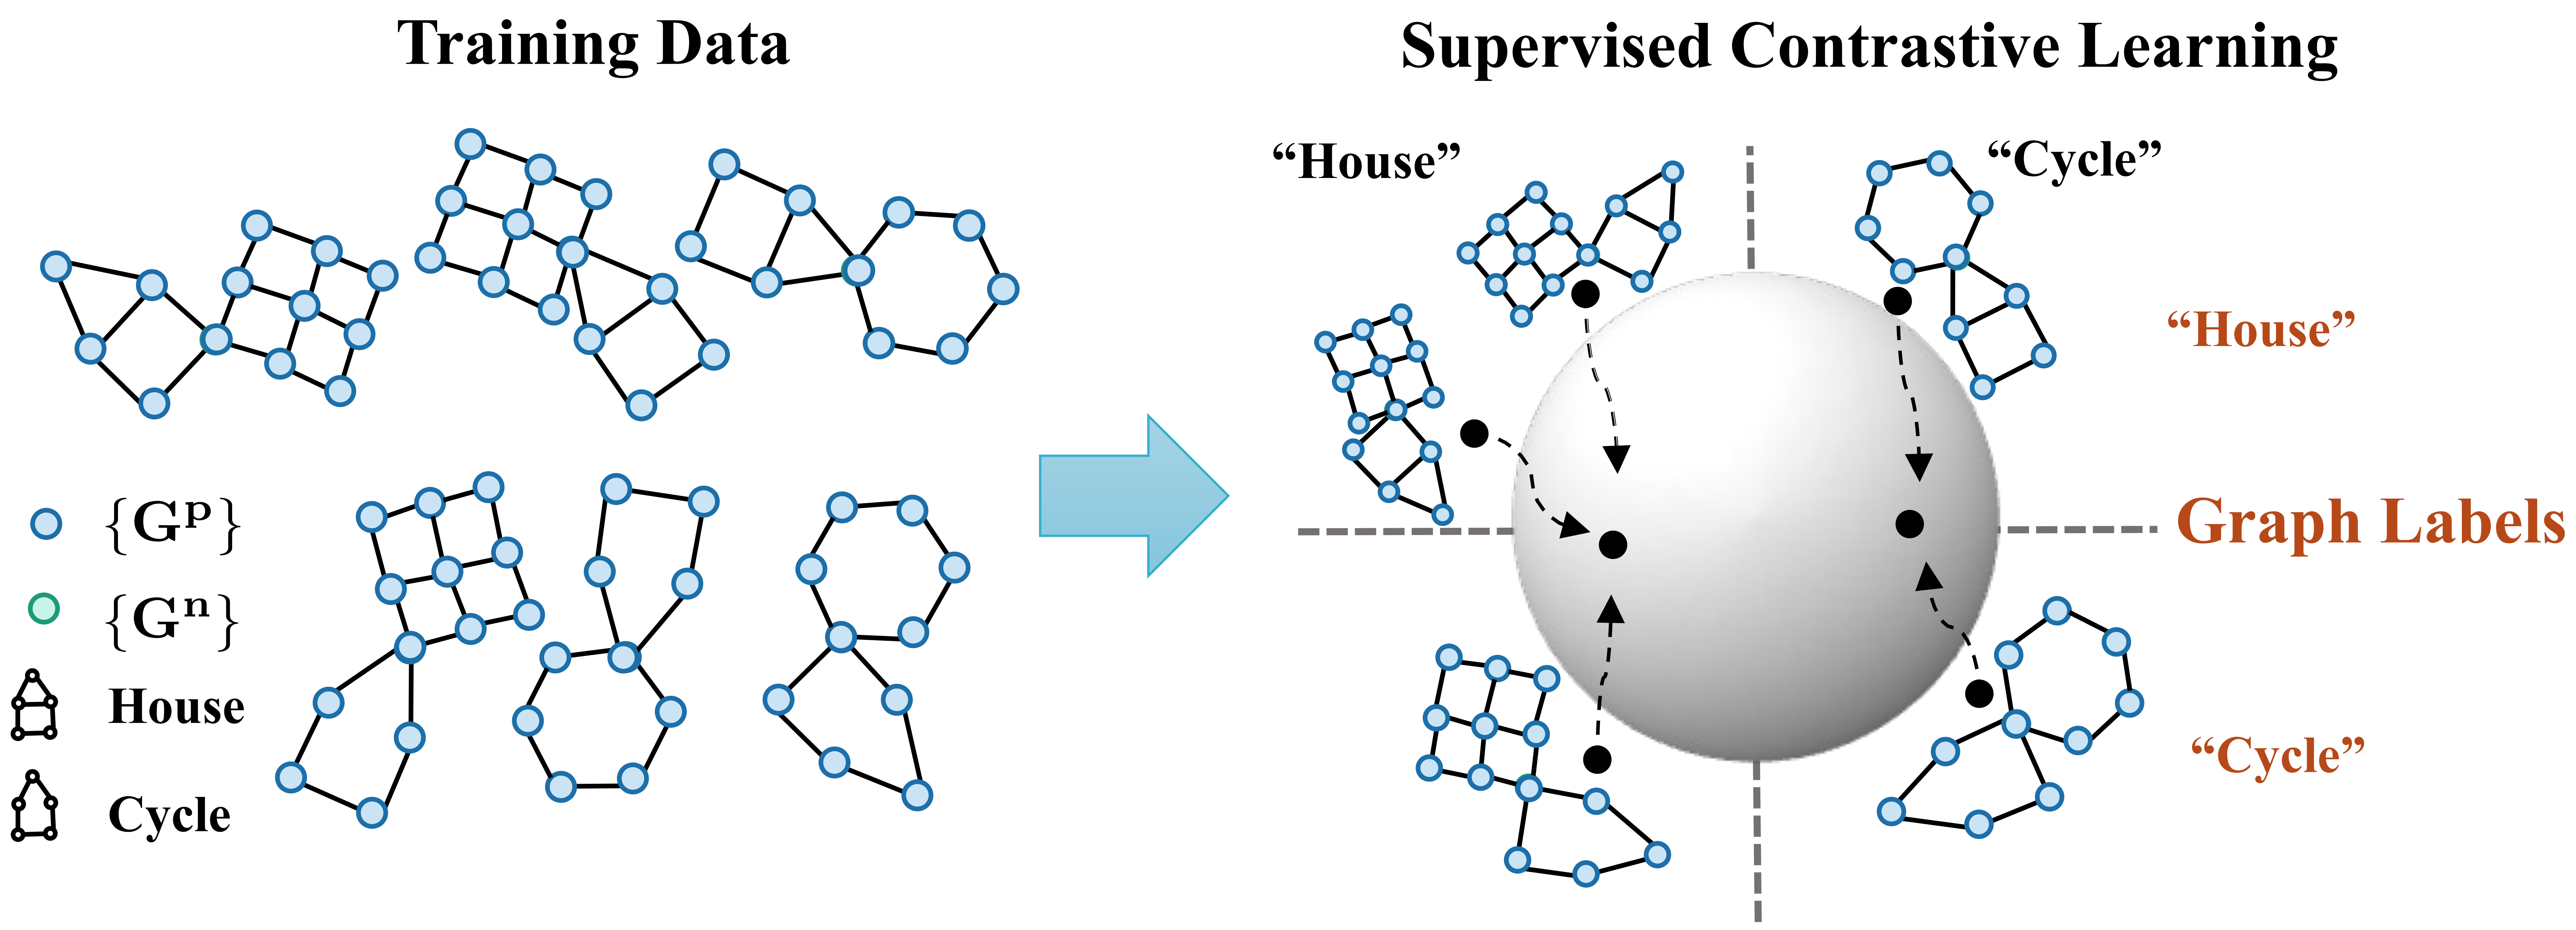

GALA: invariant GrAph Learning Assistant

Improving the contrastive invariant subgraph extraction via an **Environment Assistant**:

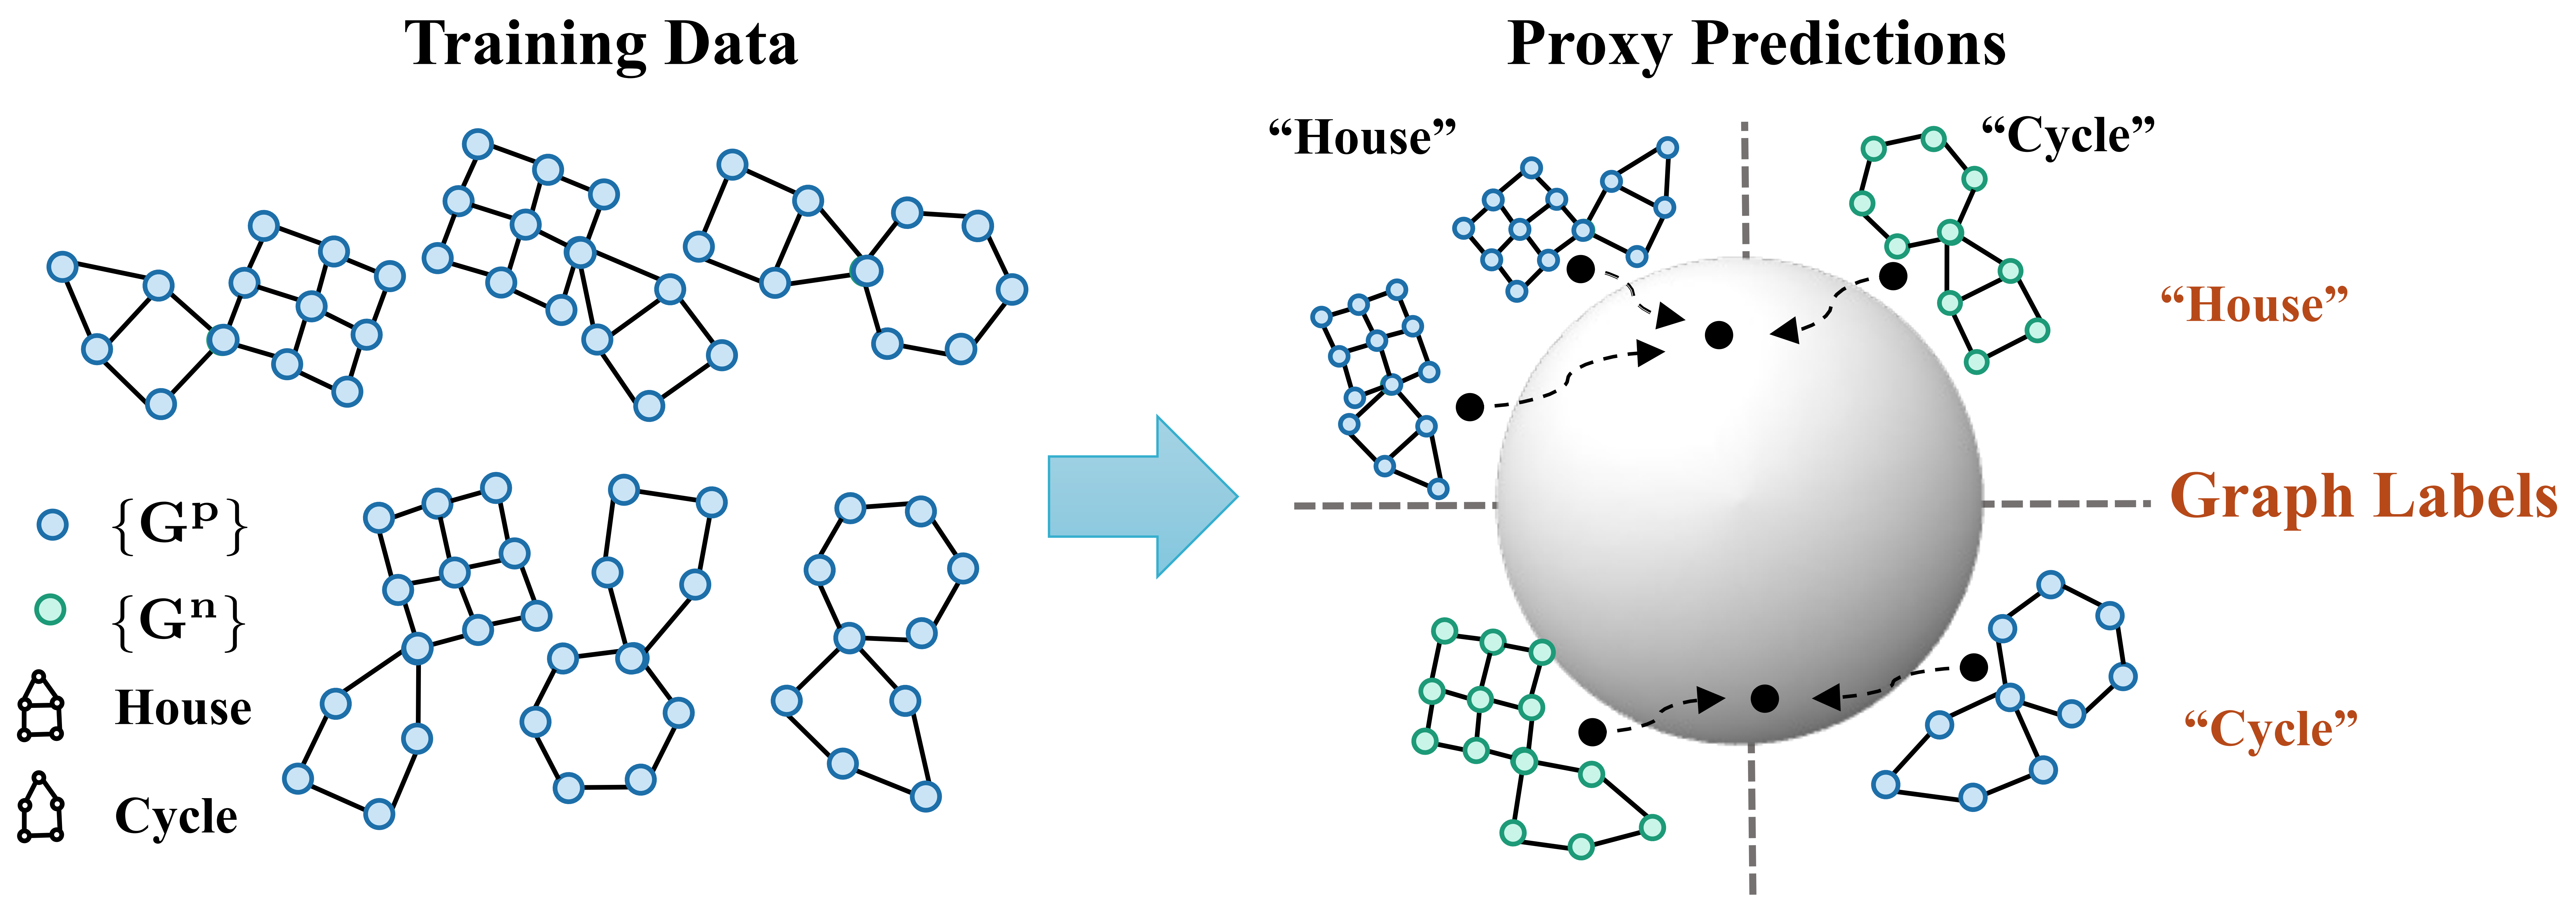

GALA: invariant GrAph Learning Assistant

Consider the following dataset dominated by spurious features:

Step 1. Obtain Environment Assistant Predictions

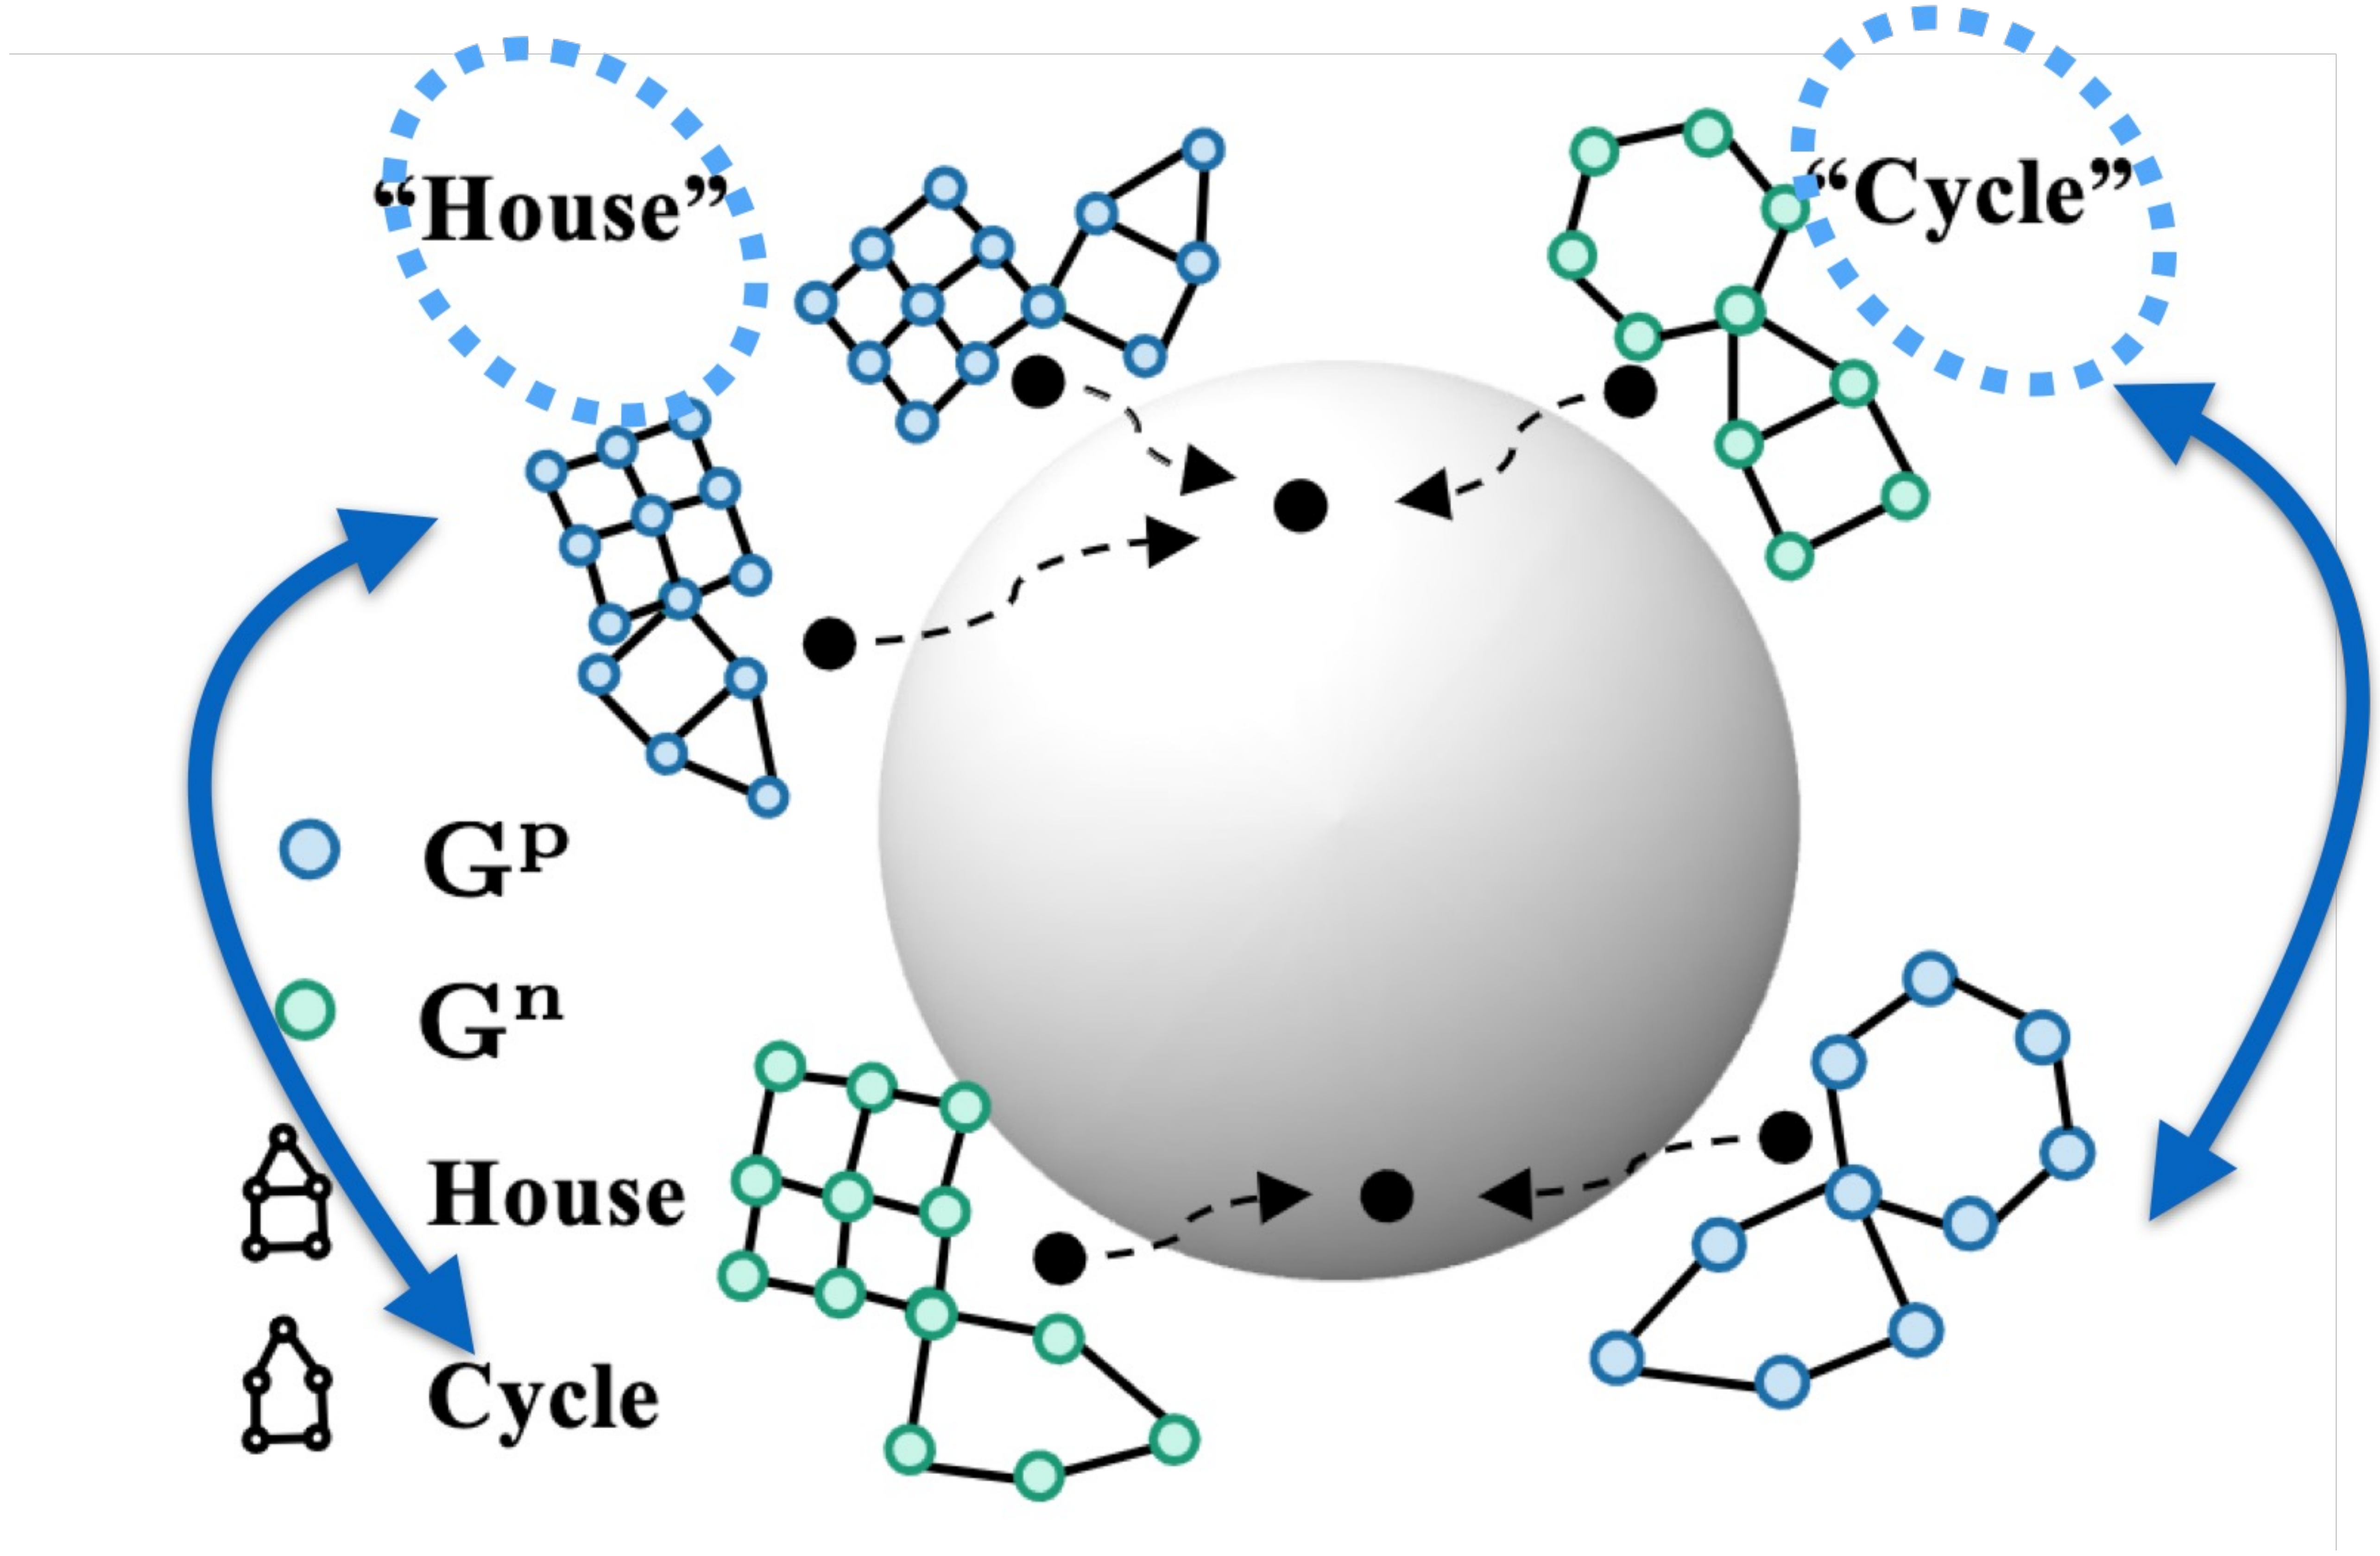

GALA: invariant GrAph Learning Assistant

Consider the following dataset dominated by spurious features:

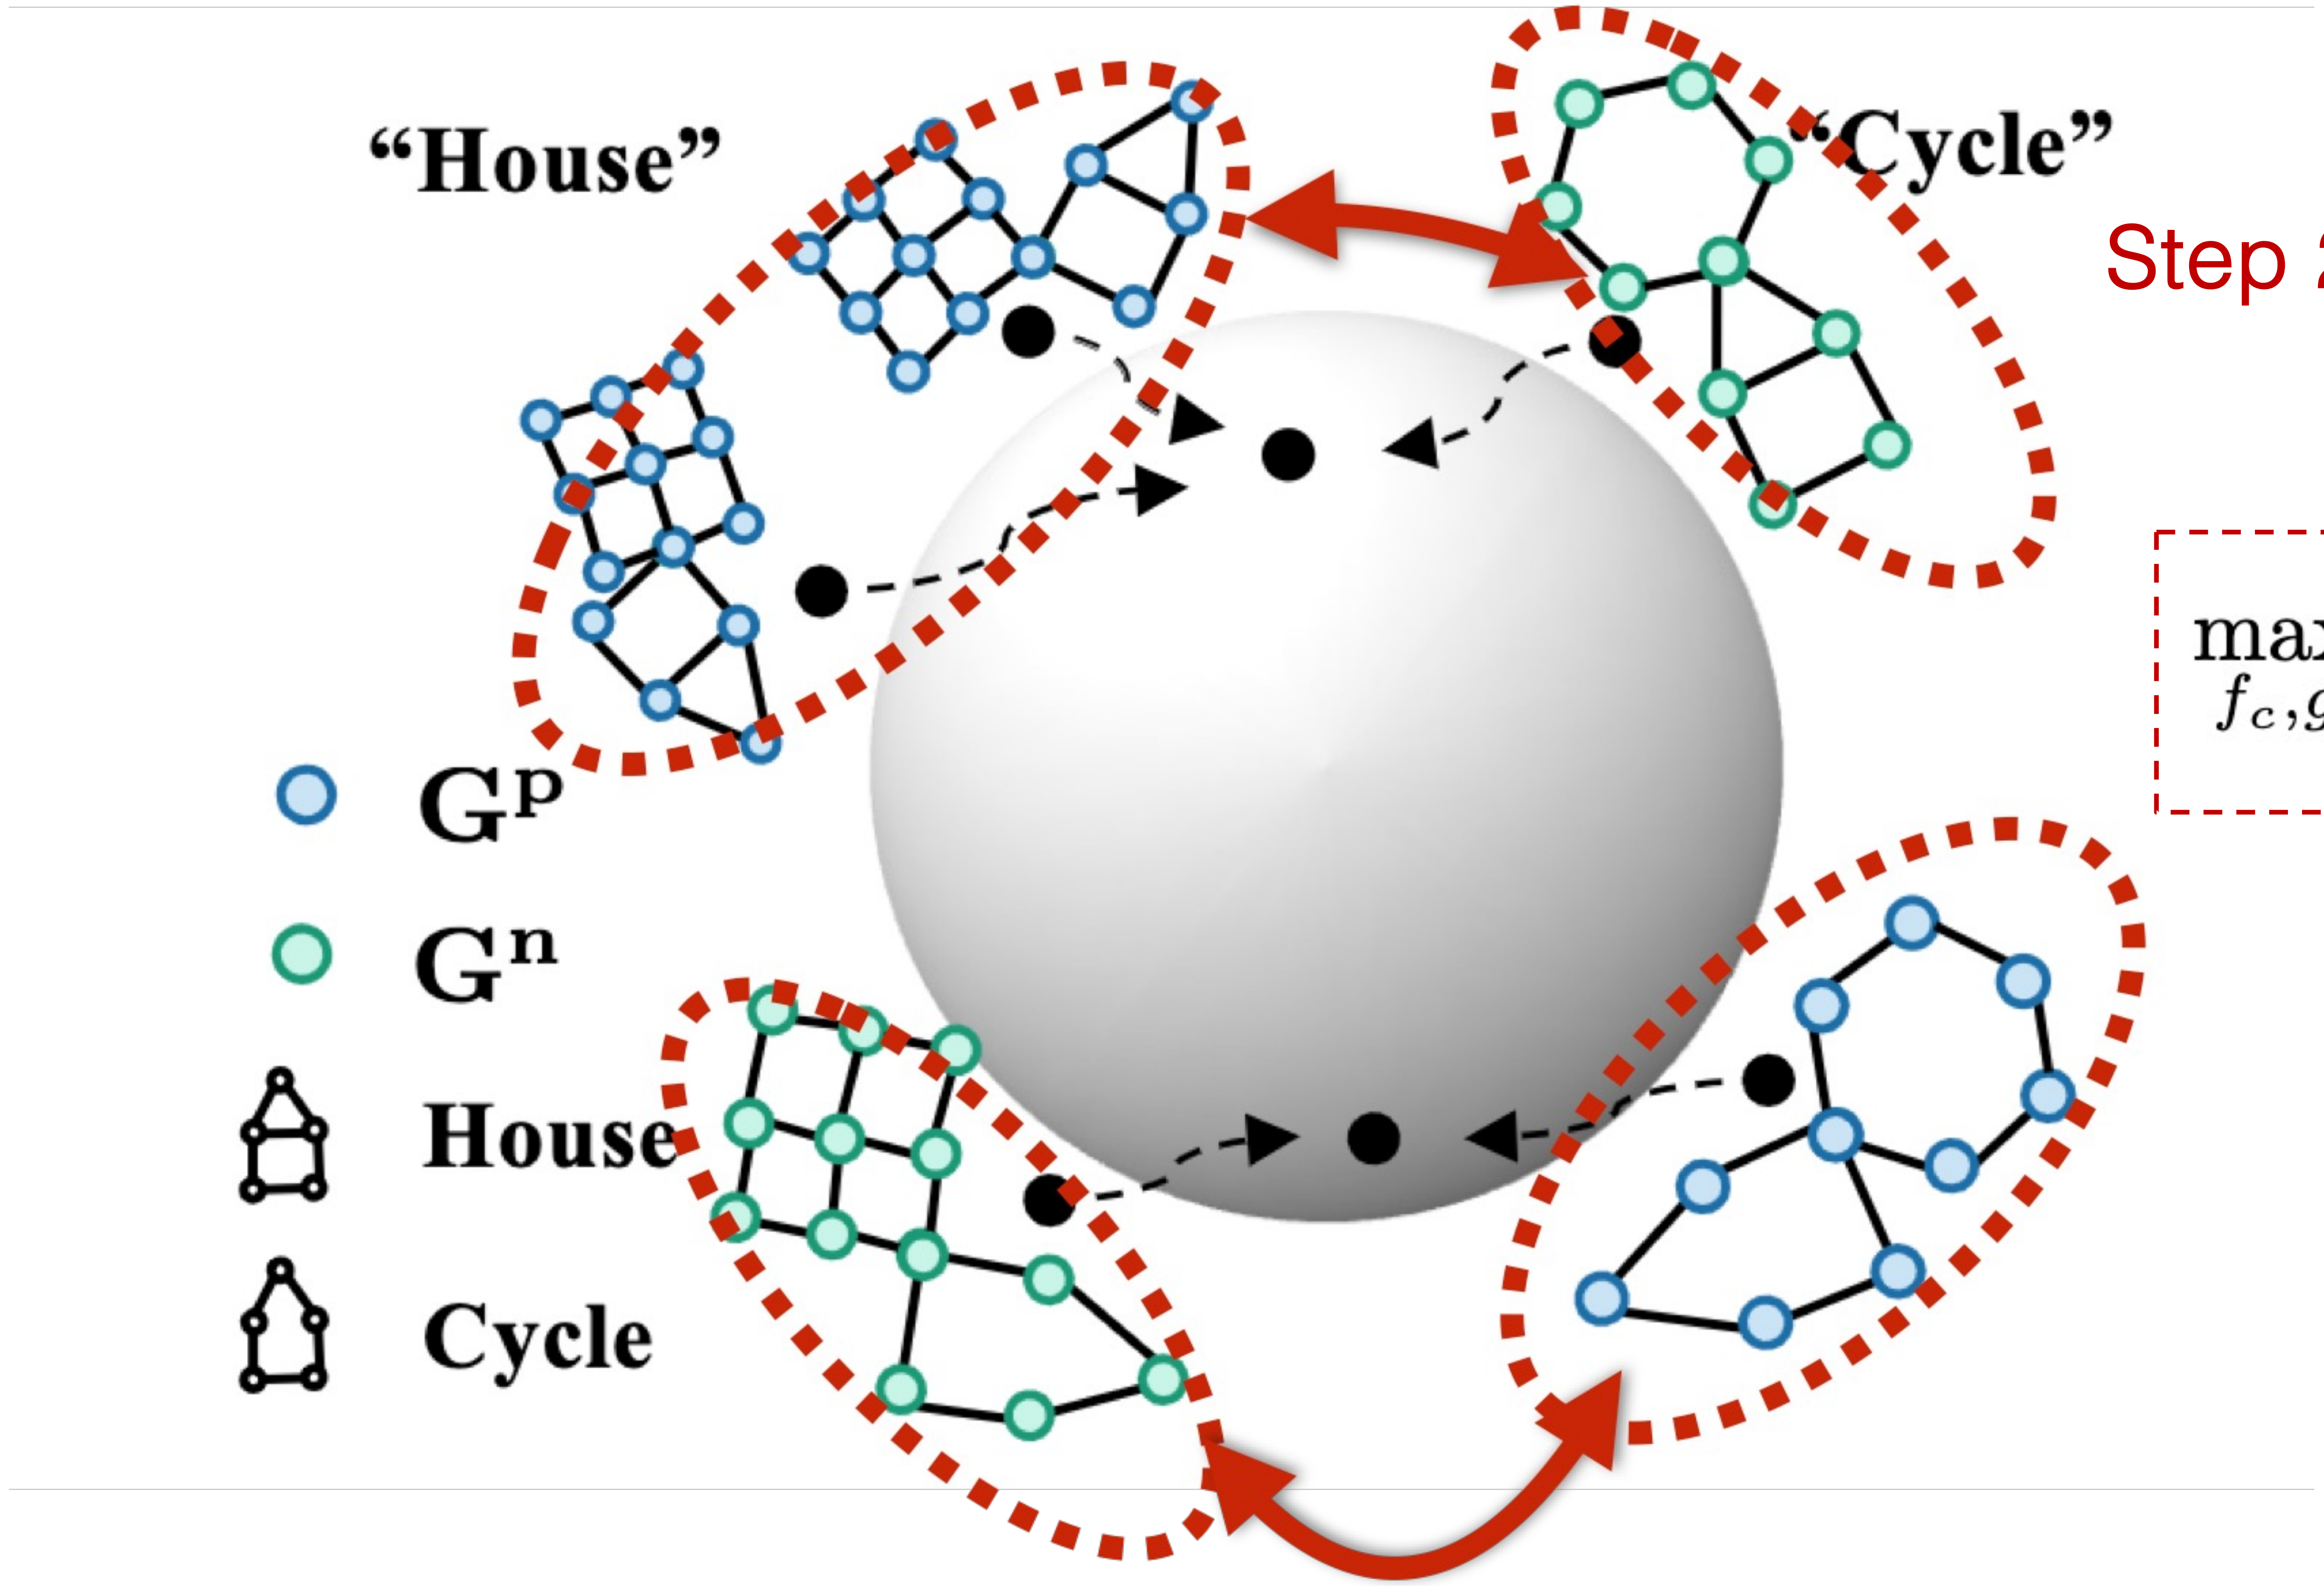

Step 2. Contrasting samples from the positive and negative graphs with the same class

$$\max_{f_c, g} I(\widehat{G}_c; Y), \ s.t. \ g \in \underset{\hat{g}, |\widehat{G}_c^p| \le s_c}{\arg\max} I(\widehat{G}_c^p; \widehat{G}_c^n | Y),$$

Proof-of-Concept Experiments

Theorem 1 (Informal)
Given the same data generation process, and the aforementioned variation sufficiency and variation consistency assumptions, when the environment assistant model learns properly distinguishes the variations of the spurious subgraphs, GALA provably identifies the invariant subgraph for OOD generalization.

Datasets	$\{0.8, 0.6\}$	$\{0.8, 0.7\}$	$\{0.8, 0.9\}$	$\{0.7, 0.9\}$	Avg.
ERM	77.33±0.47	75.65±1.62	51.37±1.20	42.73±3.82	61.77
IRM	78.32±0.70	75.13±0.77	50.76±2.56	41.32±2.50	61.38
V-Rex	77.69±0.38	74.96±1.40	49.47±3.36	41.65±2.78	60.94
IB-IRM	78.00±0.68	73.93±0.79	50.93±1.87	42.05±0.79	61.23
EIIL	76.98±1.24	74.25±1.74	51.45±4.92	39.71±2.64	60.60
XGNN	83.84±0.59	83.05±0.20	53.37±1.32	38.28±1.71	64.63
GREA	82.86±0.50	82.72±0.50	50.34±1.74	39.01±1.21	63.72
GSAT	80.54±0.88	78.11±1.23	48.63±2.18	36.62±0.87	63.32
CAL	76.98±6.03	62.95±8.58	51.57±6.33	46.23±3.93	59.43
MoleOOD	49.93±2.25	49.85±7.31	38.49±4.25	34.81±1.65	43.27
GIL	83.51±0.41	82.67±1.18	51.76±4.32	40.07±2.61	64.50
DisC	60.47±17.9	54.29±15.0	45.06±7.82	39.42±8.59	50.81
CIGA	84.03±0.53	83.21±0.30	57.87±3.38	43.62±3.20	67.18
GALA	**84.27**±0.34	**83.65**±0.44	**76.42**±3.53	**72.50**±1.06	**79.21**
Oracle	84.73±0.36	85.42±0.25	84.28±0.15	78.38±0.19	

Stronger invariant correlations

Stronger spurious correlations

Real-World Experiments

GALA consistently improves the OOD generalization performance under various real-world graph distribution shifts on a number of realistic graph benchmarks:

Datasets	EC50-Assay	EC50-Sca	EC50-Size	Ki-Assay	Ki-Sca	Ki-Size	CMNIST-sp	Graph-SST2	Avg.(Rank)†
ERM	76.42±1.59	64.56±1.25	61.61±1.52	74.61±2.28	69.38±1.65	76.63±1.34	21.56±5.38	81.54±1.13	65.79 (6.50)
IRM	77.14±2.55	64.32±0.42	62.33±0.86	75.10±3.38	69.32±1.84	76.25±0.73	20.25±3.12	82.52±0.79	65.91 (6.13)
V-Rex	75.57±2.17	64.73±0.53	62.80±0.89	74.16±1.46	71.40±2.77	76.68±1.35	30.71±11.8	81.11±1.37	67.15 (5.25)
IB-IRM	64.70±2.50	62.62±2.05	58.28±0.99	71.98±3.26	69.55±1.66	70.71±1.95	23.58±7.96	81.56±0.82	62.87 (10.6)
EIIL	64.20±5.40	62.88±2.75	59.58±0.96	74.24±2.48	69.63±1.46	76.56±1.37	23.55±7.68	82.46±1.48	64.14 (8.00)
XGNN	72.99±2.56	63.62±1.35	62.55±0.81	72.40±3.05	72.01±1.34	73.15±2.83	20.96±8.00	82.55±0.65	65.03 (7.13)
GREA	66.87±7.53	63.14±2.19	59.20±1.42	73.17±1.80	67.82±4.67	73.52±2.75	12.77±1.71	82.40±1.98	62.36 (10.1)
GSAT	76.07±1.95	63.58±1.36	61.12±0.66	72.26±1.76	70.16±0.80	75.78±2.60	15.24±3.72	80.57±0.88	64.35 (8.63)
CAL	75.10±2.71	64.79±1.58	63.38±0.88	75.22±1.73	71.08±4.83	72.93±1.71	23.68±4.68	82.38±1.01	66.07 (5.38)
DisC	61.94±7.76	54.10±5.69	57.64±1.57	54.12±8.53	55.35±10.5	50.83±9.30	50.26±0.40	76.51±2.17	56.59 (12.4)
MoleOOD	61.49±2.19	62.12±1.91	58.74±1.73	75.10±0.73	60.35±11.3	73.69±2.29	21.04±3.36	81.56±0.35	61.76 (10.0)
GIL	70.56±4.46	61.59±3.16	60.46±1.91	75.25±1.14	70.07±4.31	75.76±2.23	12.55±1.26	83.31±0.50	63.69 (8.00)
CIGA	75.03±2.47	65.41±1.16	64.10±1.08	73.95±2.50	71.87±3.32	74.46±2.32	15.83±2.56	82.93±0.63	65.45 (5.88)
GALA	**77.56**±2.88	**66.28**±0.45	**64.25**±1.21	**77.92**±2.48	**73.17**±0.88	**77.40**±2.04	**68.94**±0.56	**83.60**±0.66	**73.64 (1.00)**
Oracle	84.77±0.58	82.66±1.19	84.53±0.60	91.08±1.43	88.58±0.64	92.50±0.53	67.76±0.60	91.40±0.26	

† Averaged rank is also reported in the parentheses because of dataset heterogeneity. A lower rank is better.

A Short Summary of GALA

NEURAL INFORMATION PROCESSING SYSTEMS

We conducted a retrospective study on the faithfulness of the augmented environment information for OOD generalization on graphs.

By showing the impossibility results, we developed a set of minimal assumptions for feasible invariant graph learning.

we proposed a provable feasible approach GALA under the assumptions. Extensive experiments with 11 datasets verified the superiority of GALA.

Paper

Code

Thank you!

Contact: yqchen@cse.cuhk.edu.hk

To appear at NeurIPS 2023

Spotlight Presentation *at ICLR'23 DG workshop*