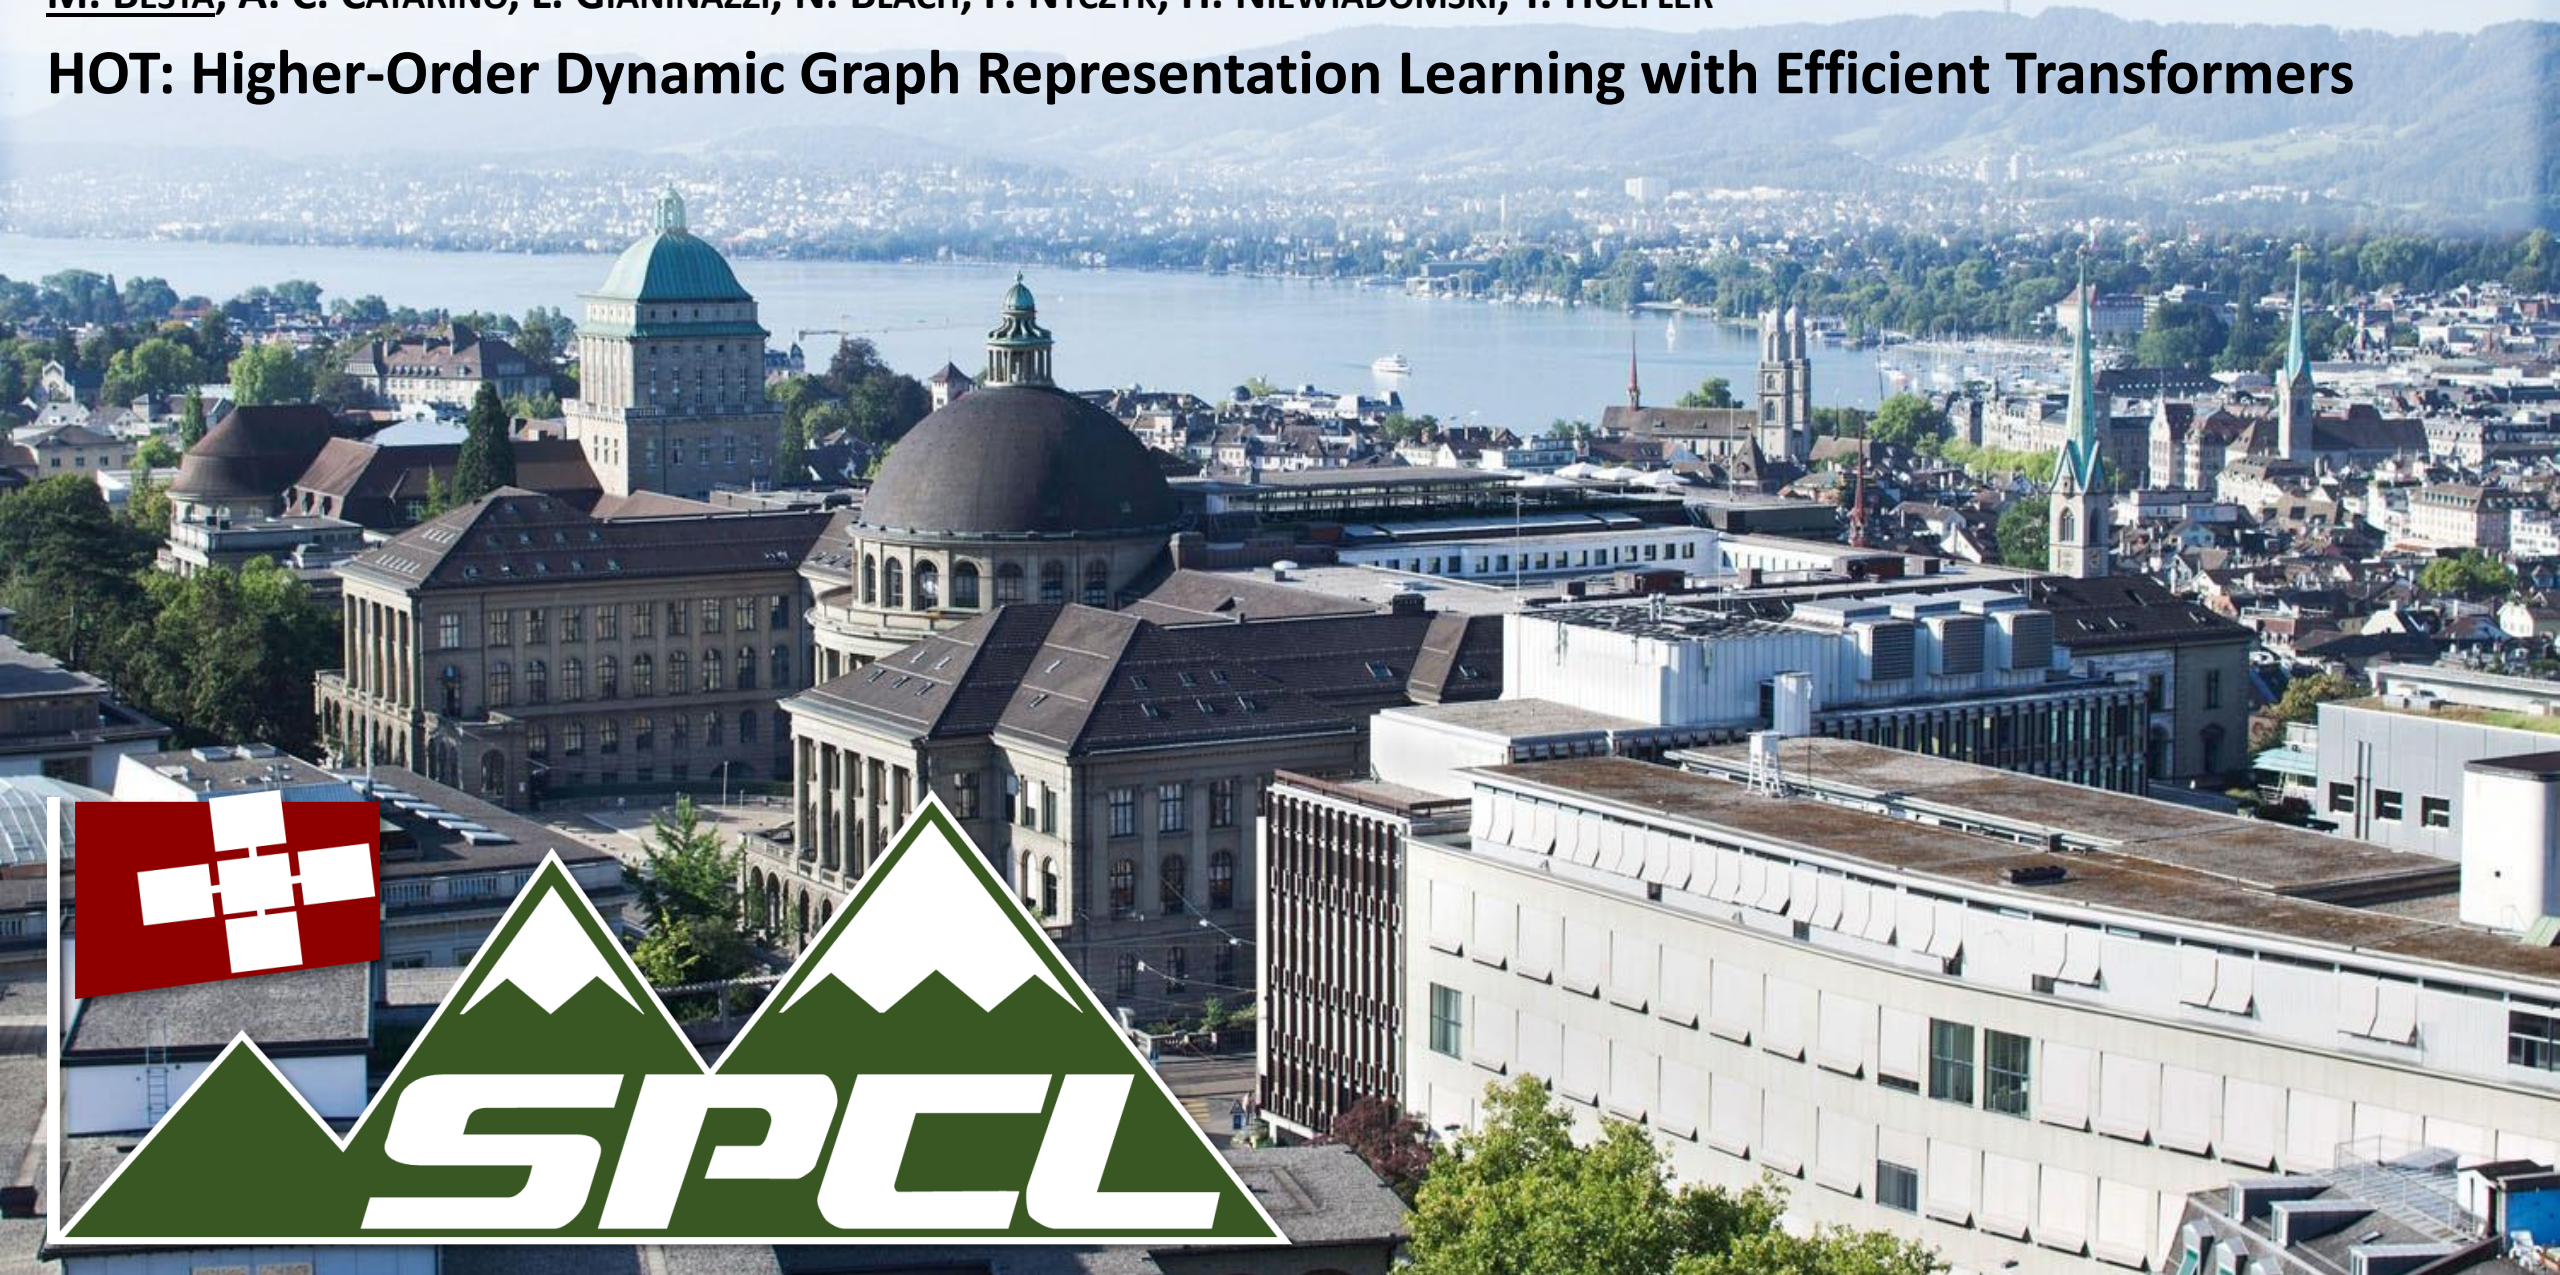


M. BESTA, A. C. CATARINO, L. GIANINAZZI, N. BLACH, P. NYCZYK, H. NIEWIADOMSKI, T. HOEFLER

HOT: Higher-Order Dynamic Graph Representation Learning with Efficient Transformers



M. BESTA, A. C. CATARINO, L. GIANINAZZI, N. BLACH, P. NYCZYK, H. NIEWIADOMSKI, T. HOEFLER

HOT: Higher-Order Dynamic Graph Representation Learning with Efficient Transformers

@ LOG'23

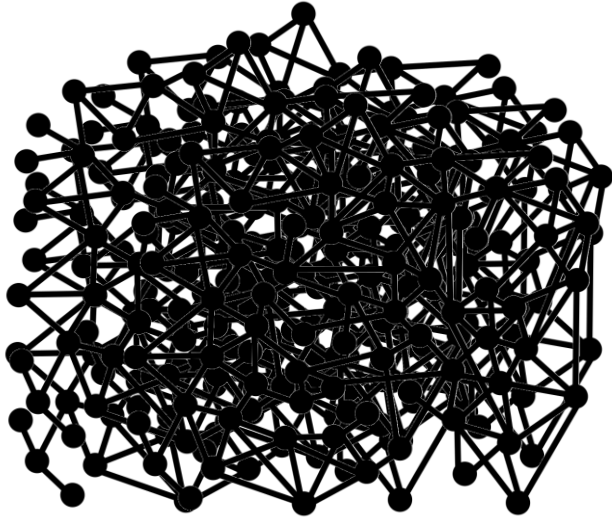
HOT: Higher-Order Dynamic Graph Representation Learning with Efficient Transformers

Maciej Besta^{1*} Afonso Claudino Catarino^{1*} Lukas Gianinazzi¹ Nils Blach¹
Piotr Nyczyk² Hubert Niewiadomski² Torsten Hoefler¹

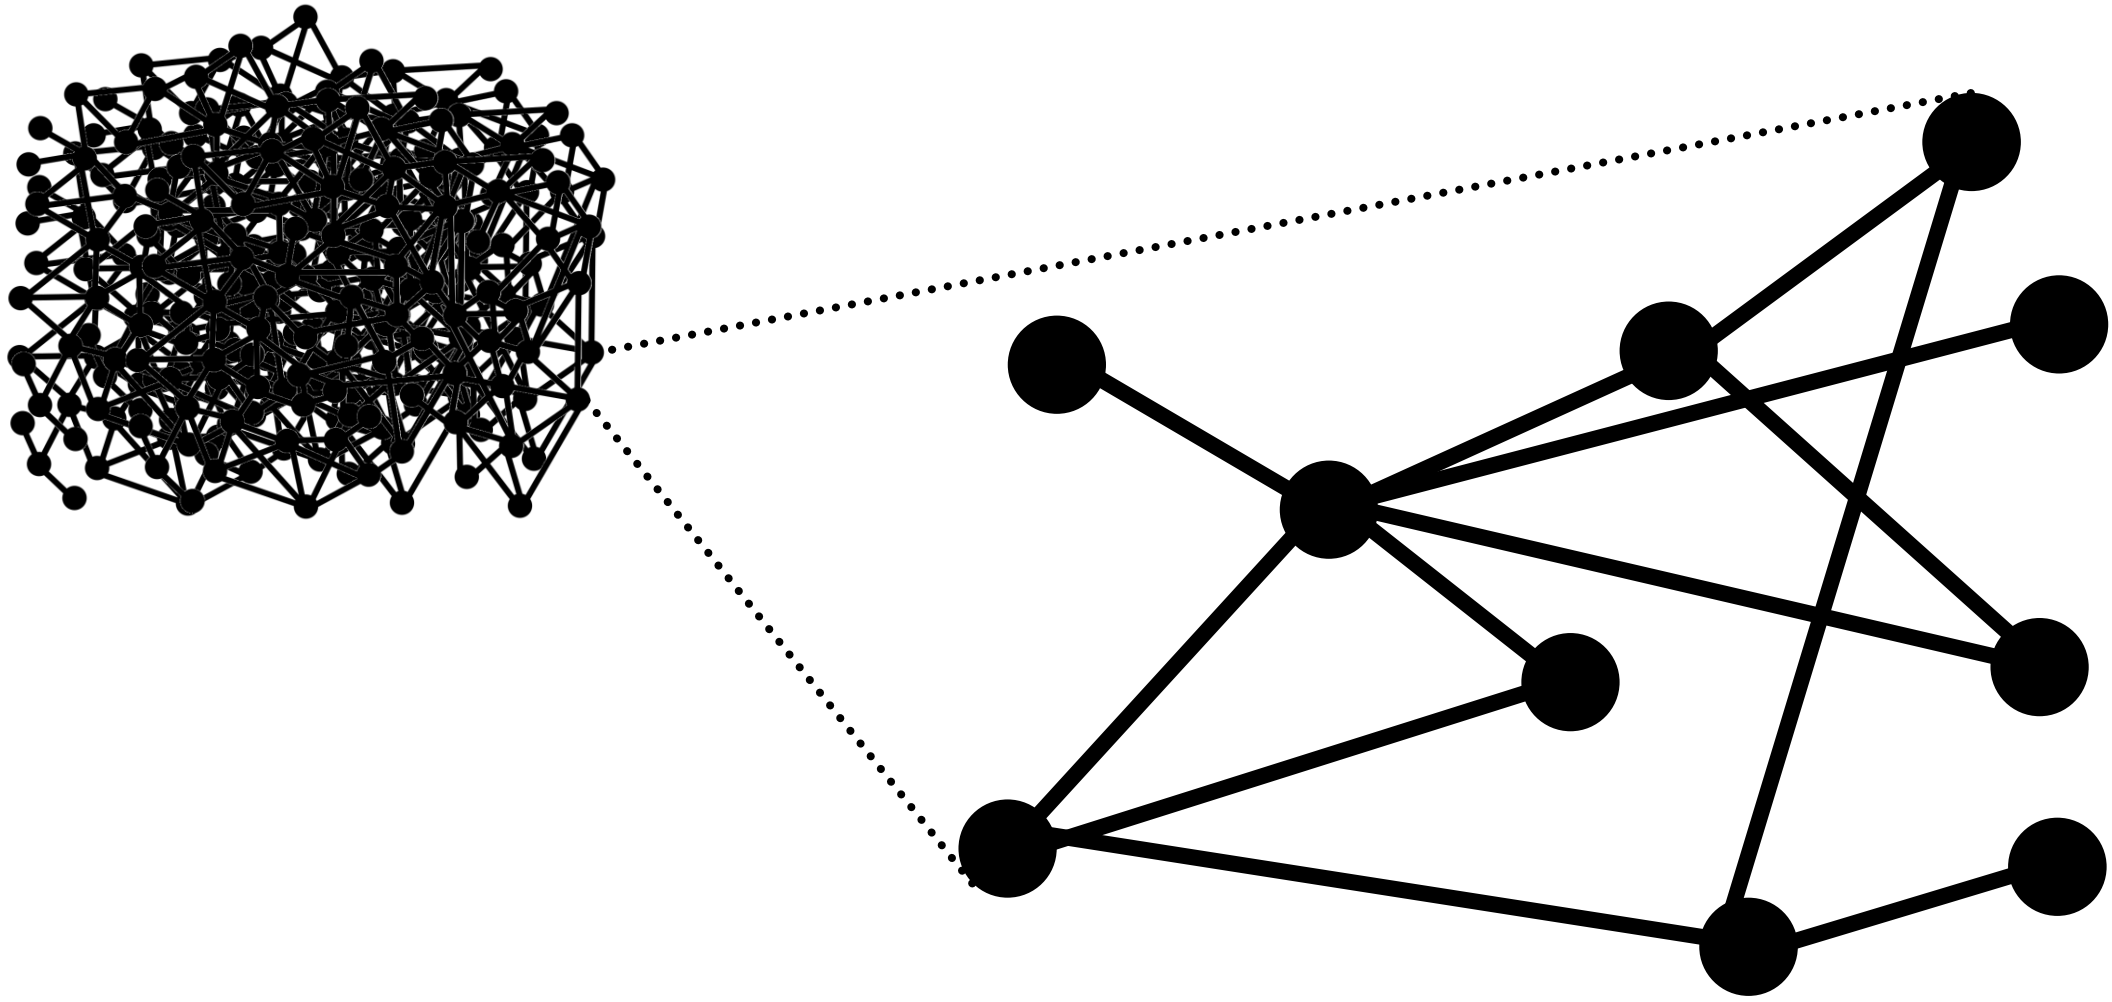
¹Department of Computer Science, ETH Zurich; ²Cledar

Link Prediction

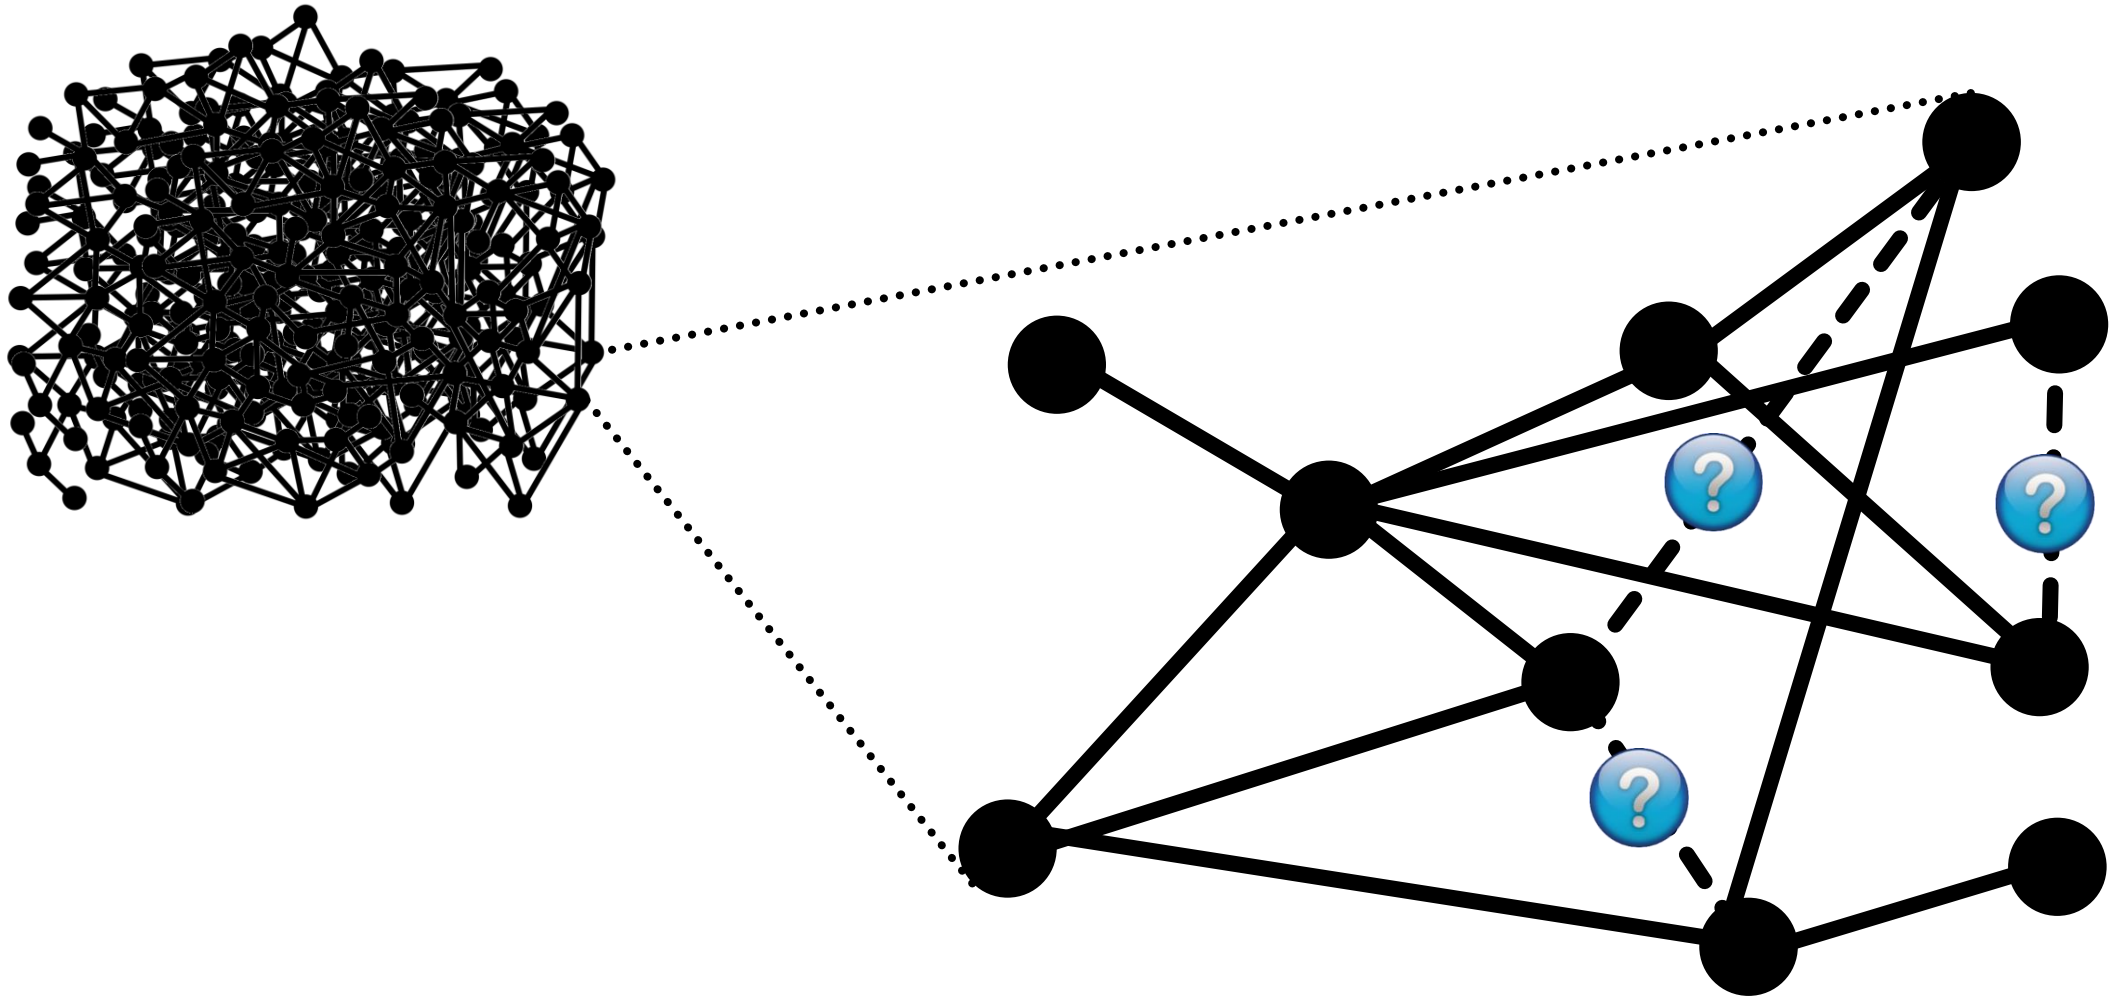
Link Prediction



Link Prediction

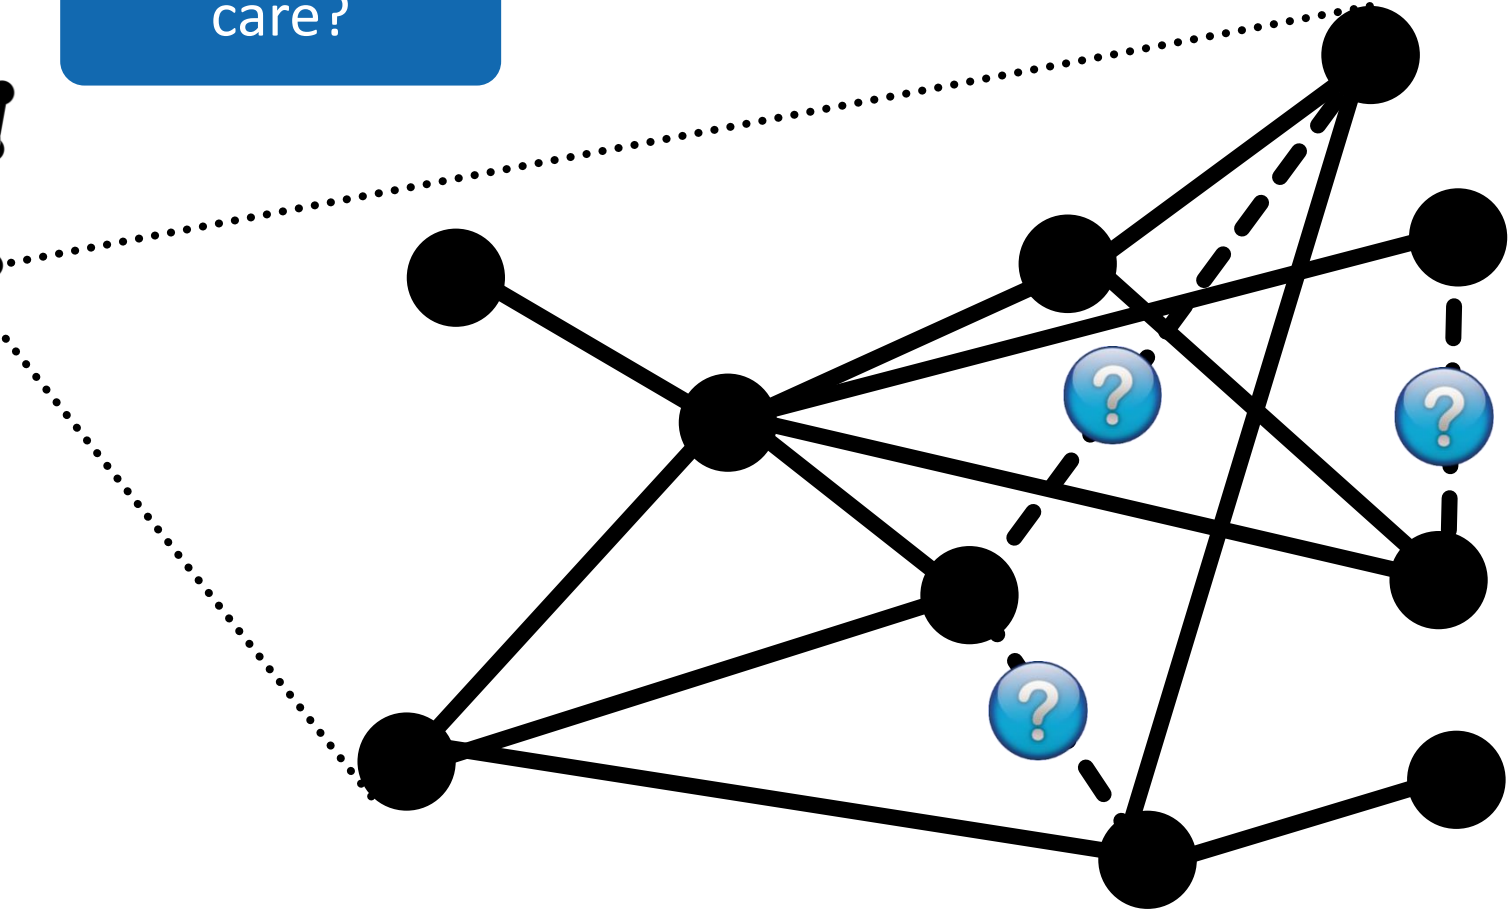
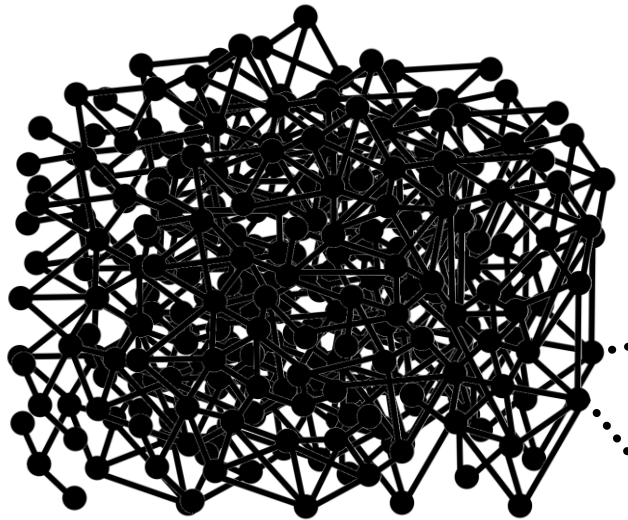


Link Prediction

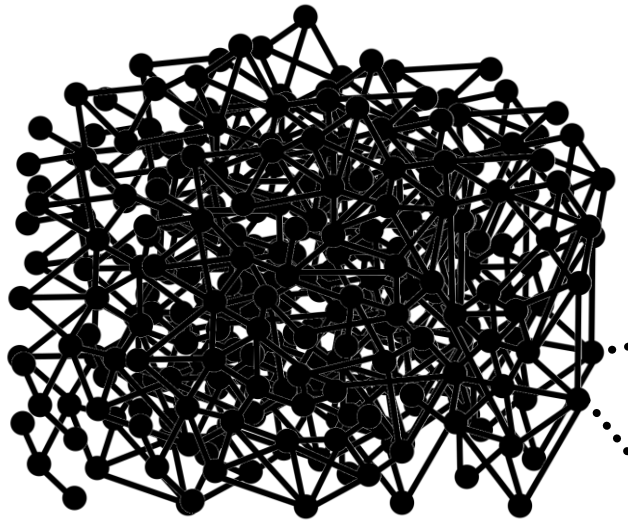


Link Prediction

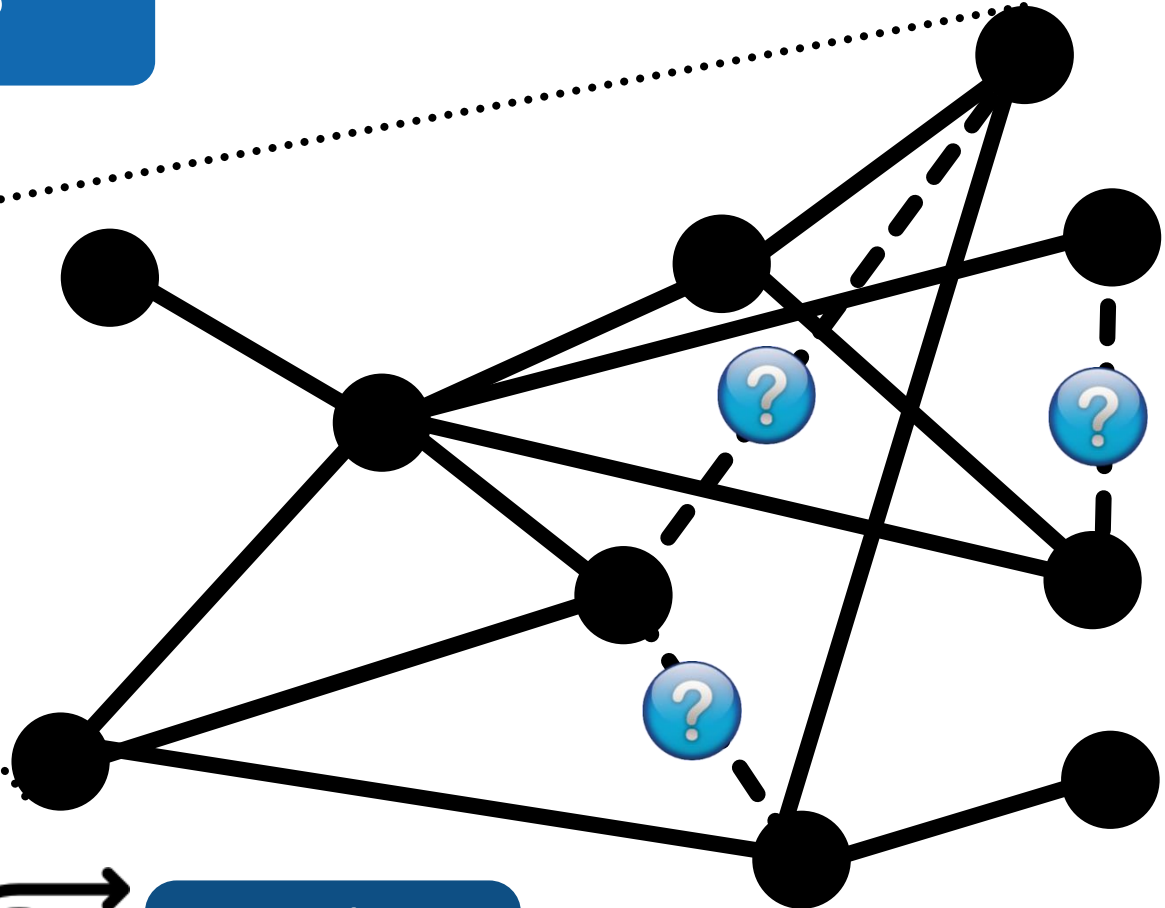
Why do we care?



Link Prediction



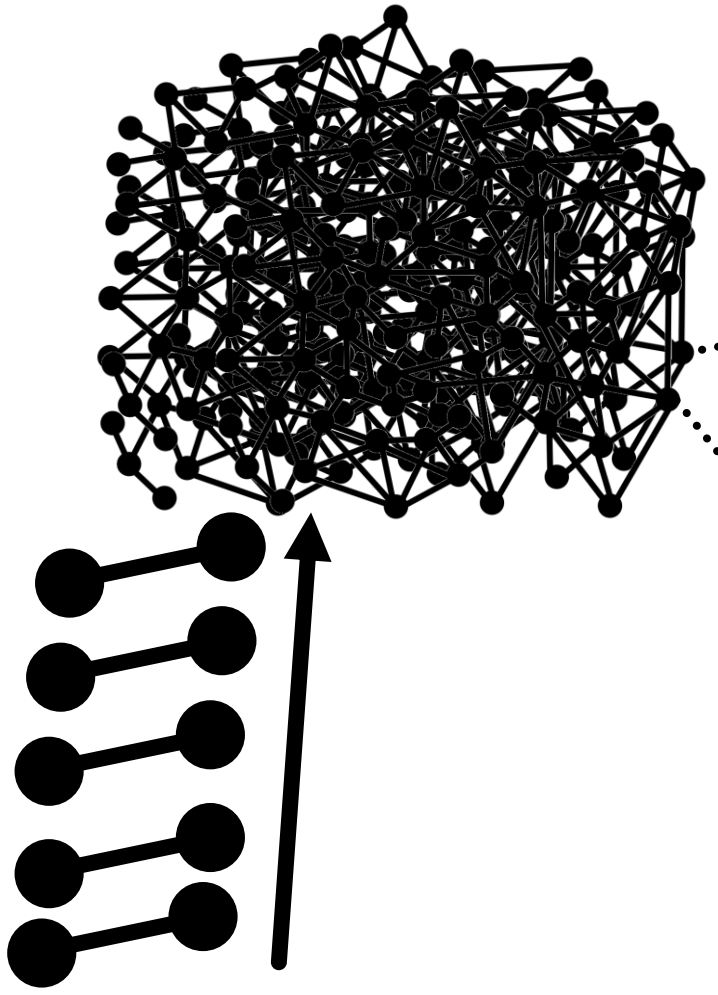
Why do we care?



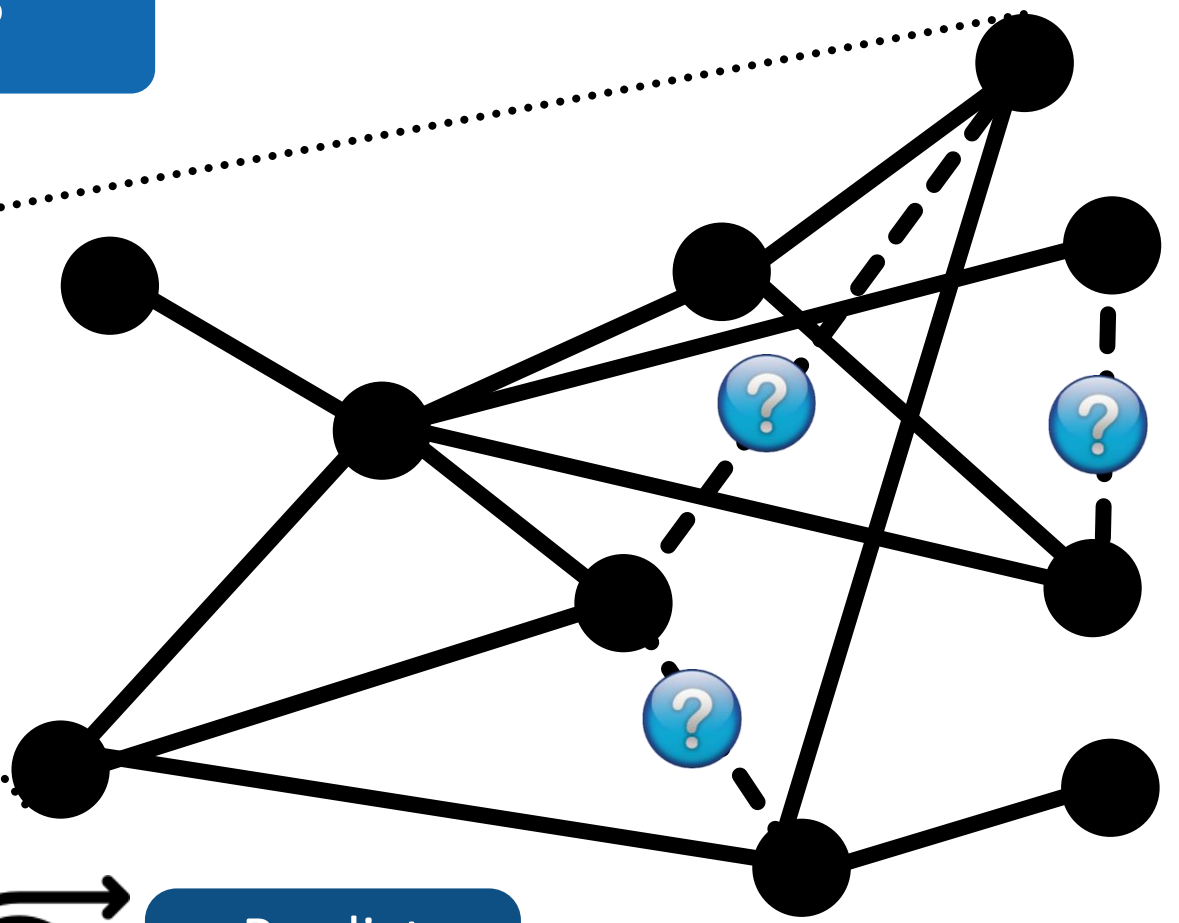
 Predict future data

Link Prediction

Why do we care?



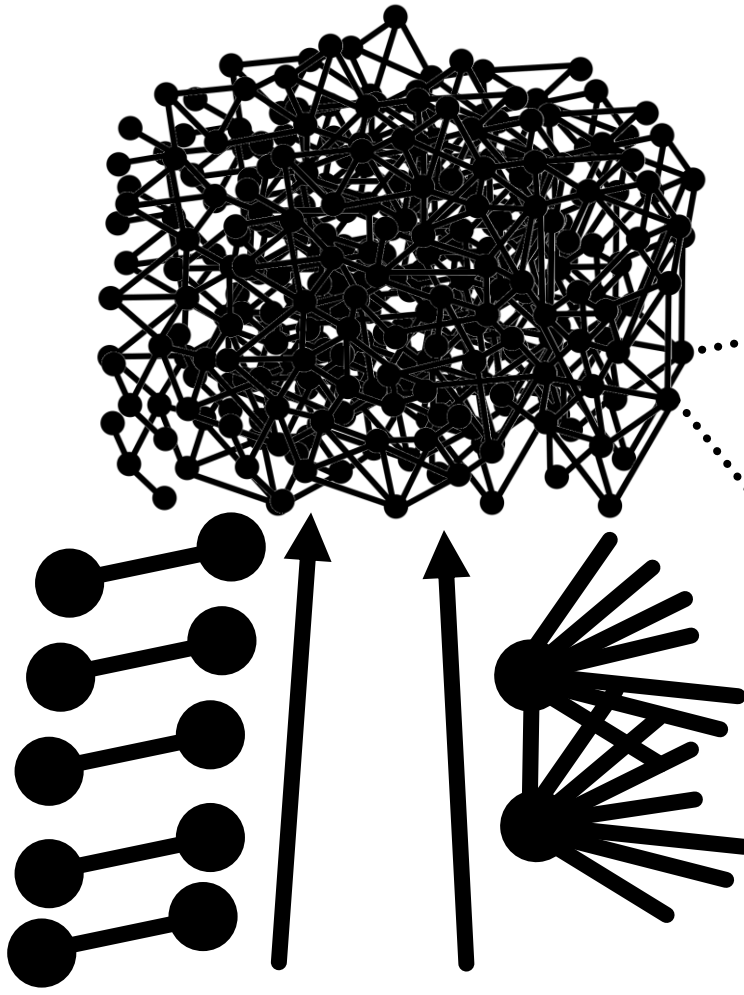
Predict future data



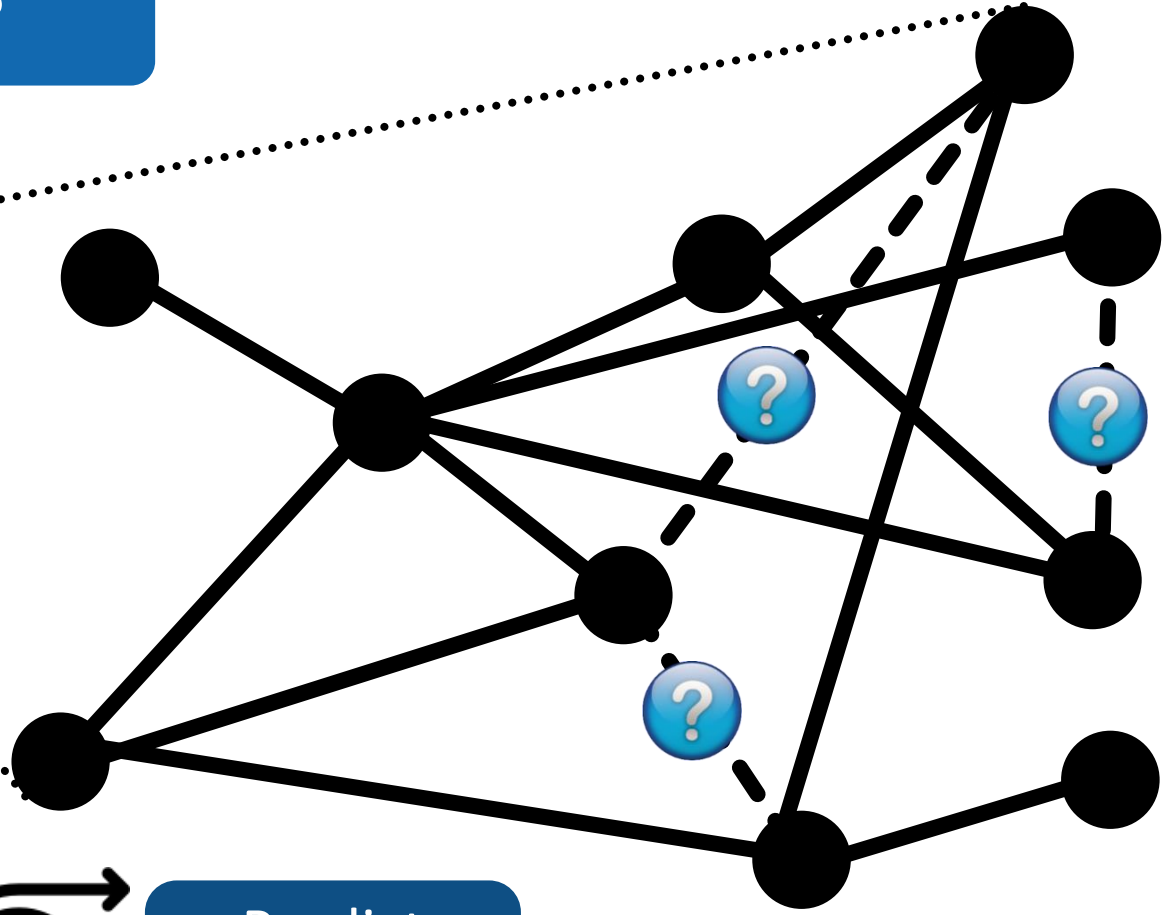
We have access to the history of past updates

Link Prediction

Why do we care?

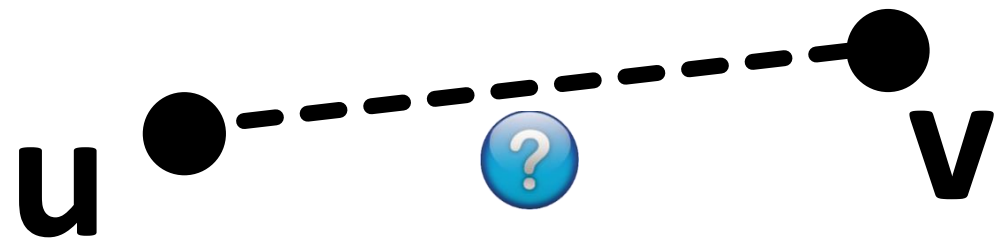


Predict future data

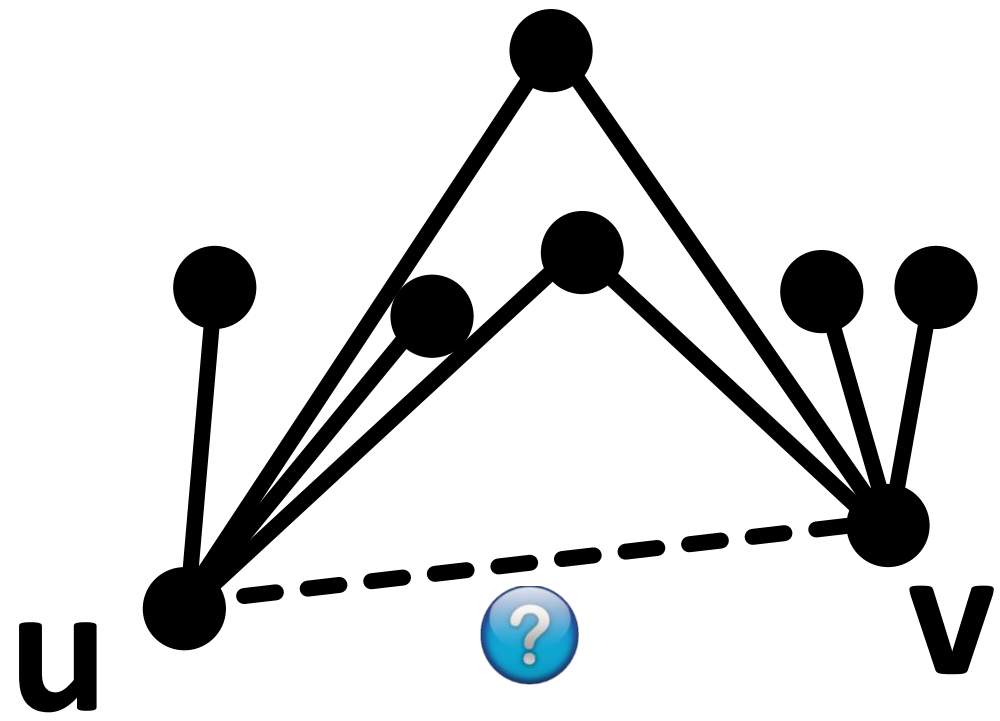


We have access to the history of past updates

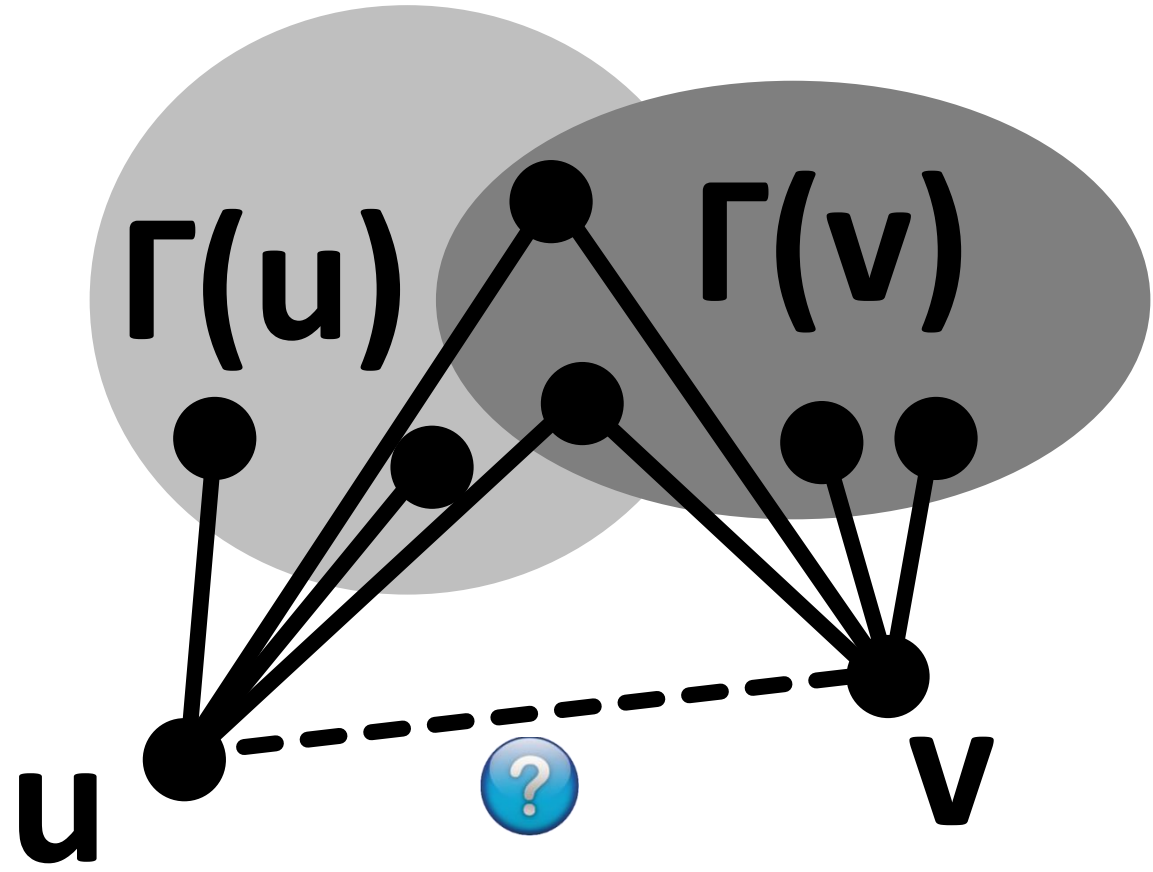
Link Prediction



Link Prediction



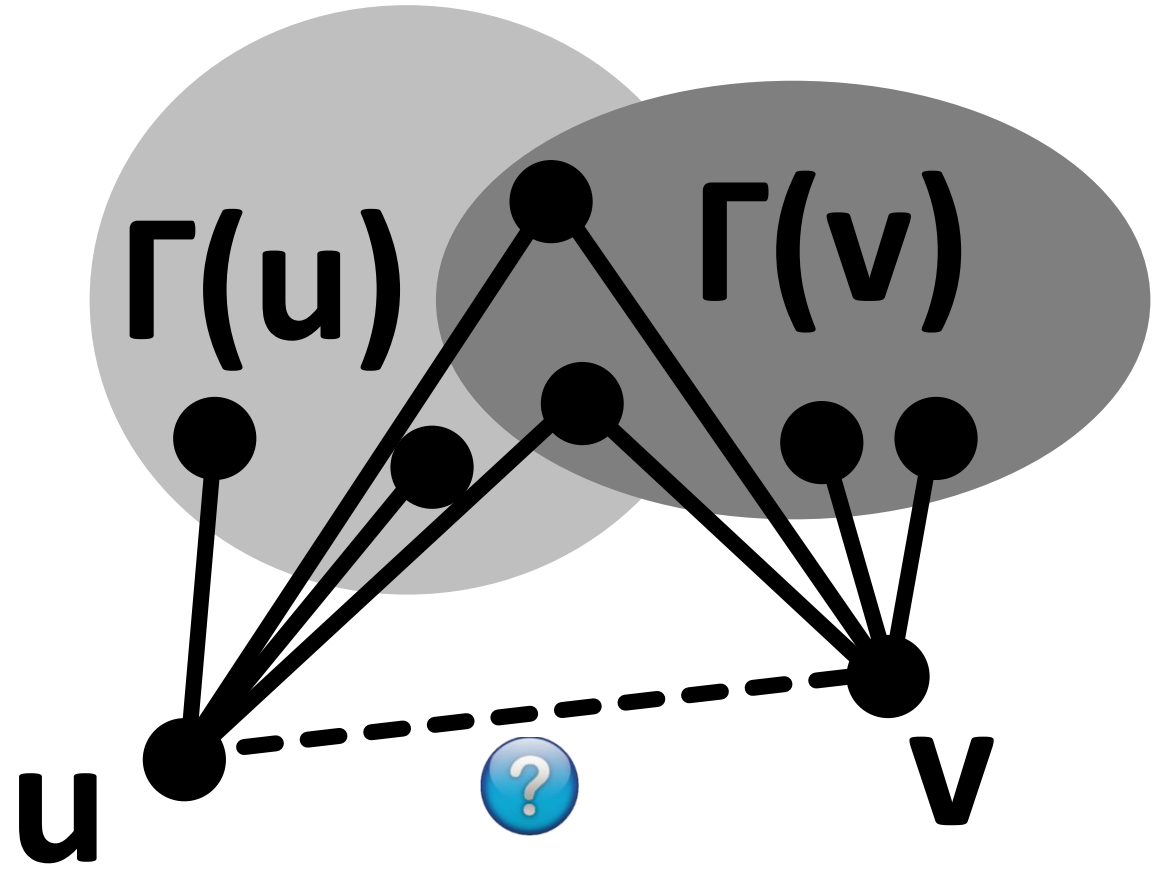
Link Prediction



Link Prediction

$$S_{u,v}^{Jaccard} = \frac{|\Gamma(u) \cap \Gamma(v)|}{|\Gamma(u) \cup \Gamma(v)|}$$

$$S_{u,v}^{CN} = |\Gamma(u) \cap \Gamma(v)|$$



Link Prediction

$$S_{u,v}^{Jaccard} = \frac{|\Gamma(u) \cap \Gamma(v)|}{|\Gamma(u) \cup \Gamma(v)|}$$

$$S_{u,v}^{RA} = \sum_{z \in \Gamma(u) \cap \Gamma(v)} \frac{1}{d_z}$$

$$S_{u,v}^{CN} = |\Gamma(u) \cap \Gamma(v)|$$

$$S_{u,v}^{HPI} = \frac{|\Gamma(u) \cap \Gamma(v)|}{\min\{d_u, d_v\}}$$

$$S_{u,v}^{Salton} = \frac{|\Gamma(u) \cap \Gamma(v)|}{\sqrt{d_u d_v}}$$

$$S_{u,v}^{HDI} = \frac{|\Gamma(u) \cap \Gamma(v)|}{\max\{d_u, d_v\}}$$

$$S_{u,v}^{Sorensen} = \frac{2|\Gamma(u) \cap \Gamma(v)|}{d_u + d_v}$$

$$S_{u,v}^{AA} = \sum_{z \in \Gamma(u) \cap \Gamma(v)} \frac{1}{\log d_z}$$

$$S_{u,v}^{PAI} = |\Gamma(u)| |\Gamma(v)| = d_u d_v$$

Link Prediction

$$S_{u,v}^{Jaccard} = \frac{|\Gamma(u) \cap \Gamma(v)|}{|\Gamma(u) \cup \Gamma(v)|}$$

$$S_{u,v}^{RA} = \sum_{z \in \Gamma(u) \cap \Gamma(v)} \frac{1}{d_z}$$

$$S_{u,v}^{CN} = |\Gamma(u) \cap \Gamma(v)|$$

Only „simple” (dyadic) edges are considered

$$S_{u,v}^{HPI} = \frac{|\Gamma(u) \cap \Gamma(v)|}{\min\{d_u, d_v\}}$$

$$\frac{|\Gamma(u) \cap \Gamma(v)|}{\sqrt{d_u d_v}}$$

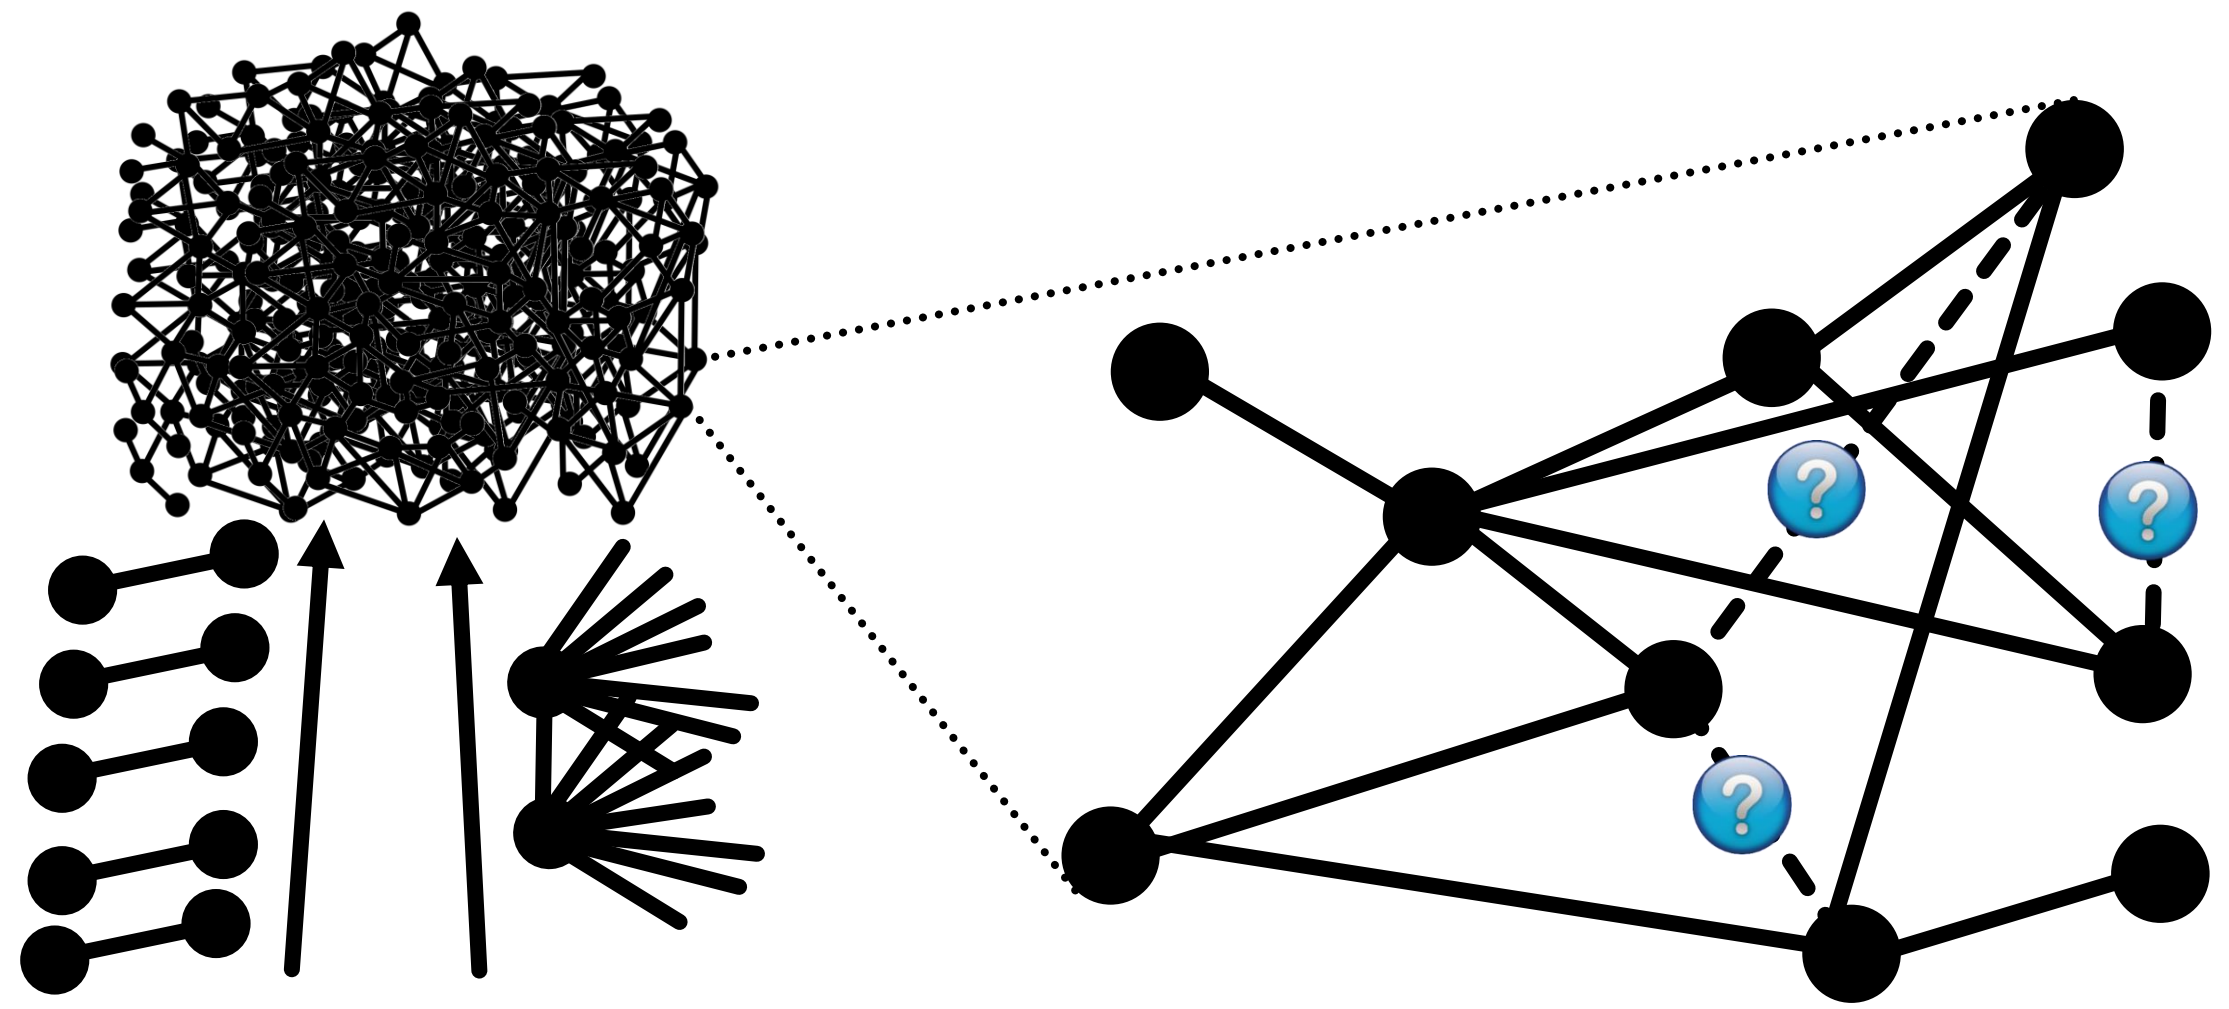
$$S_{u,v}^{HDI} = \frac{|\Gamma(u) \cap \Gamma(v)|}{\max\{d_u, d_v\}}$$

$$S_{u,v}^{Sorensen} = \frac{2|\Gamma(u) \cap \Gamma(v)|}{d_u + d_v}$$

$$S_{u,v}^{AA} = \sum_{z \in \Gamma(u) \cap \Gamma(v)} \frac{1}{\log d_z}$$

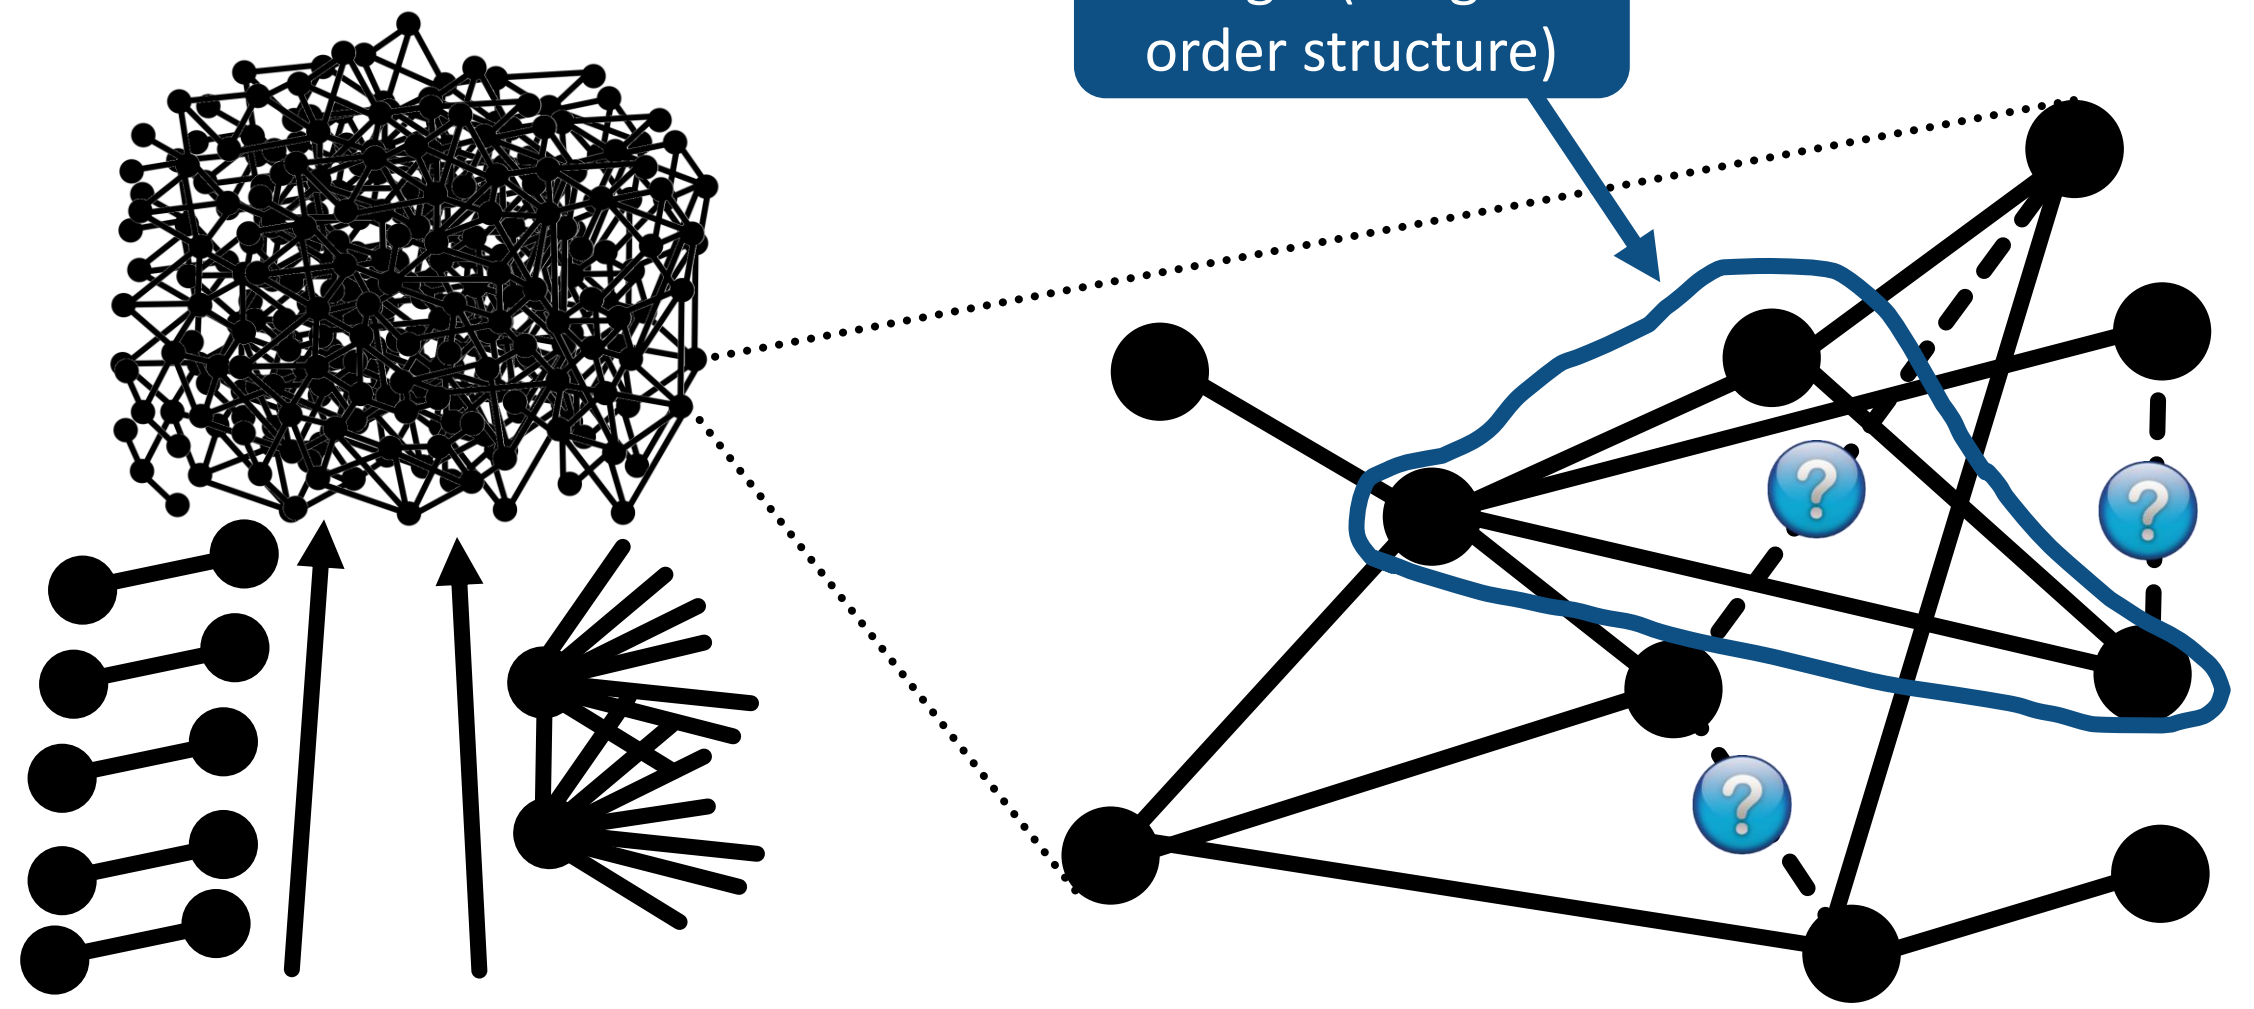
$$S_{u,v}^{PAI} = |\Gamma(u)| |\Gamma(v)| = d_u d_v$$

Link Prediction



We have access to the history of past updates

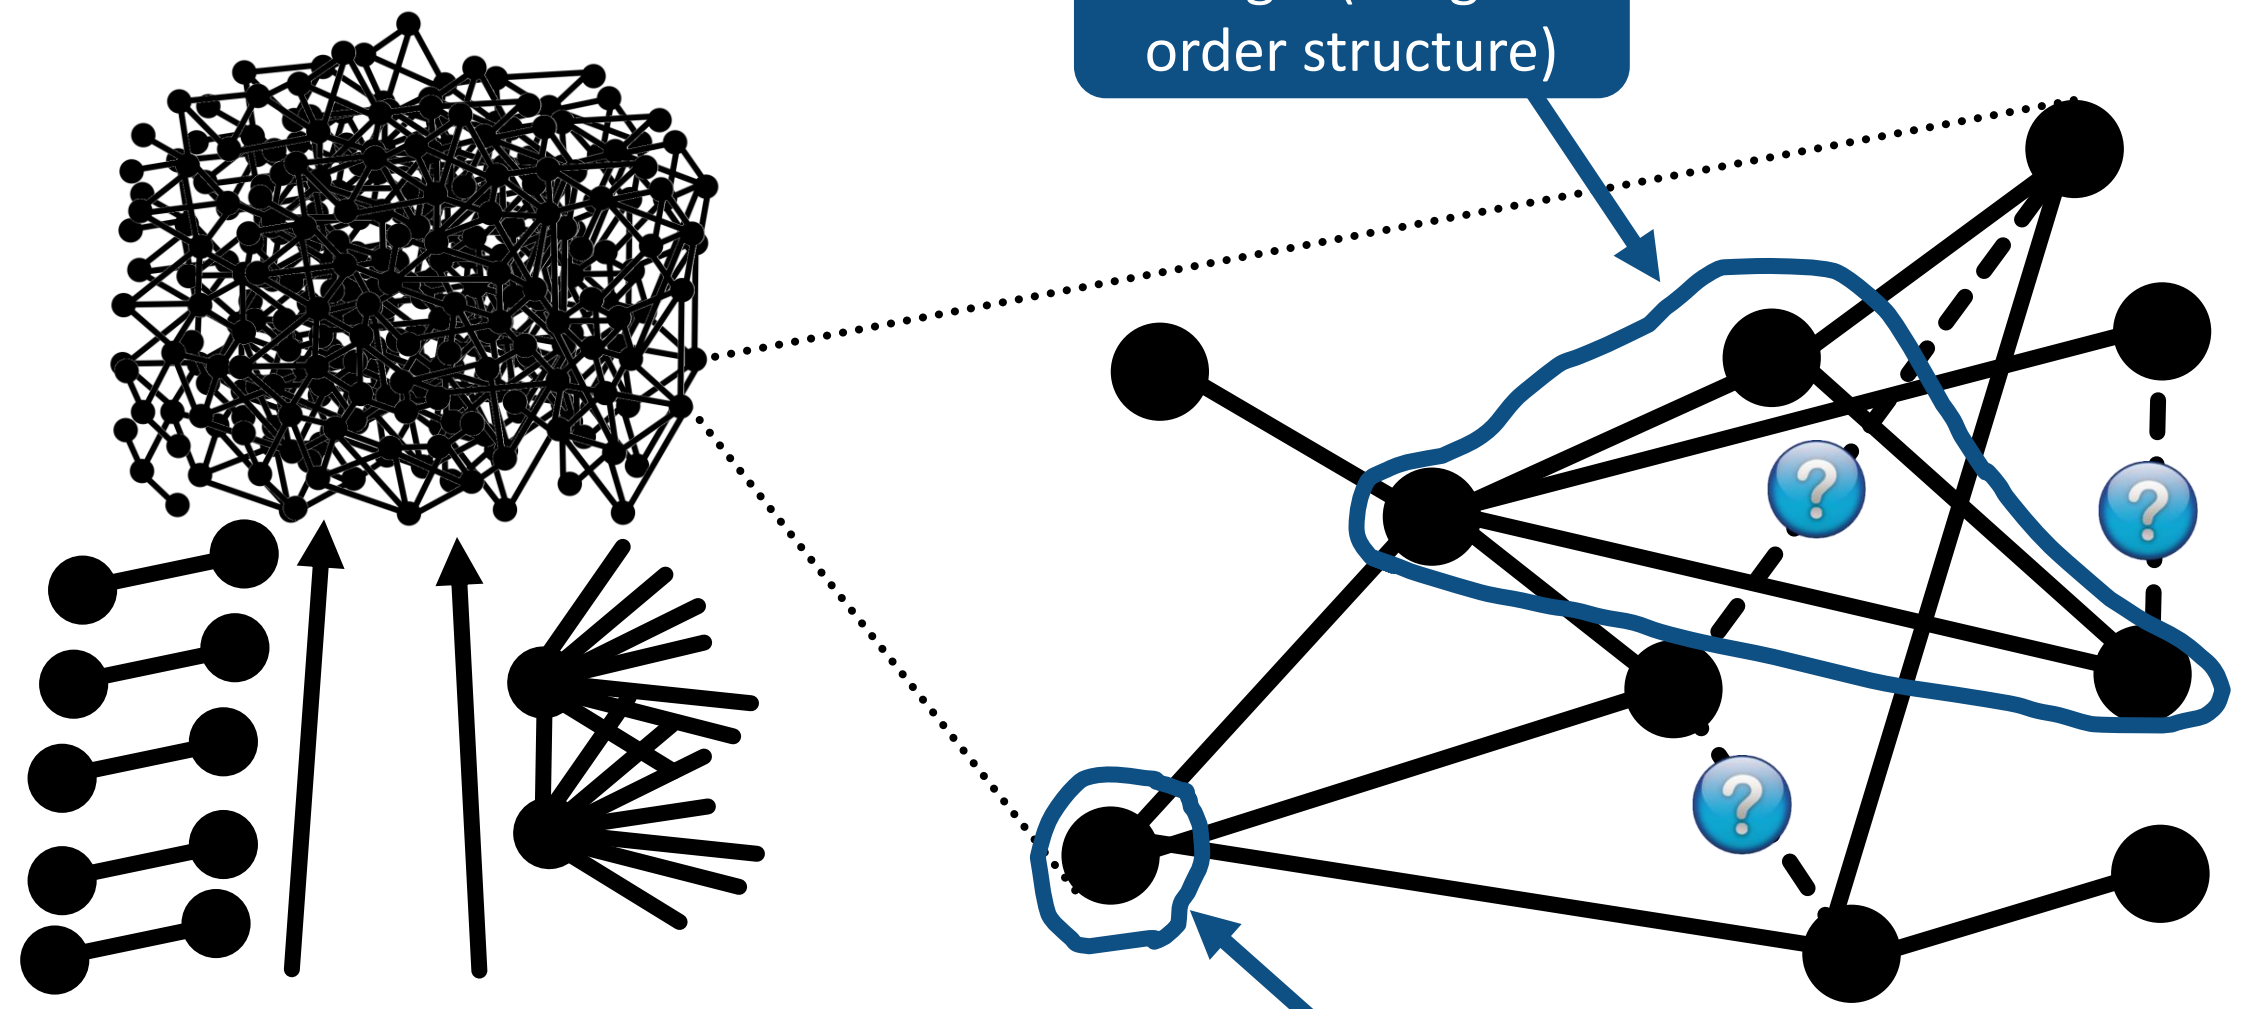
Link Prediction



Triangle (a higher-order structure)

We have access to the history of past updates

Link Prediction



We have access to the history of past updates

A higher order (2-hop) neighbor

Triangle (a higher-order structure)

Link Prediction

Higher Order (HO) - why do we even care?



We have access to the history of past updates

Triangle (a higher-order structure)



A higher order (2-hop) neighbor

Link Prediction

Higher Order (HO) - why do we even care?



Incorporating HO graph structures results in fundamentally more powerful predictions for many workloads [1, ...]



We have access to the history of past updates

A higher order (2-hop) neighbor

[1] C. Morris et al. Weisfeiler and leman go neural: Higher-order graph neural networks. AAI 2019.

Link Prediction

Higher Order (HO) - why do we even care?



Incorporating HO graph structures results in fundamentally more powerful predictions for many workloads [1, ...]

Before we go on...
what Higher Order exactly is?



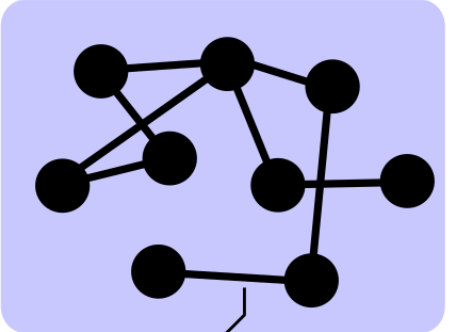
We have access to the history of past updates

A higher order (2-hop) neighbor

Higher-Order Graph Structures: Summary & Scope

HO: „Anything that goes beyond a simple dyadic edge”

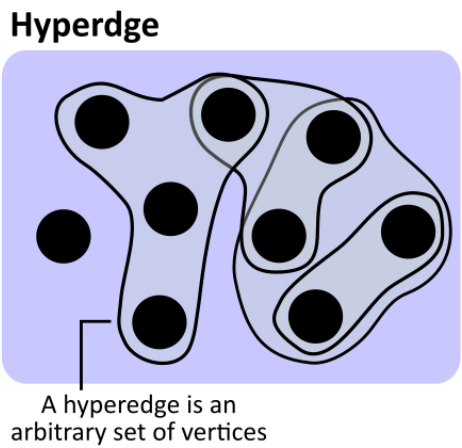
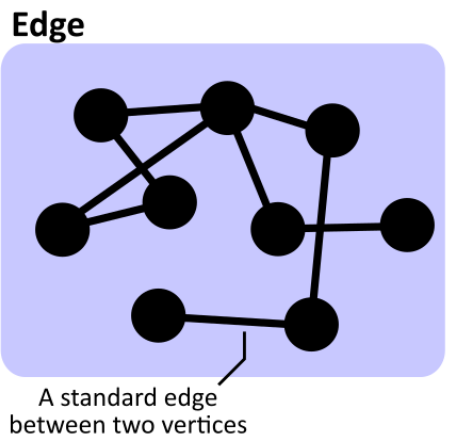
Edge



A standard edge between two vertices

Higher-Order Graph Structures: Summary & Scope

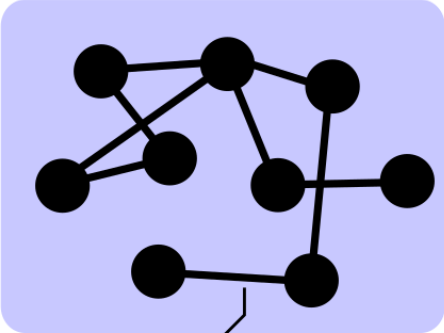
HO: „Anything that goes beyond a simple dyadic edge”



Higher-Order Graph Structures: Summary & Scope

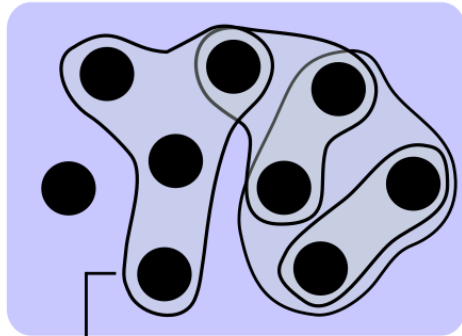
HO: „Anything that goes beyond a simple dyadic edge”

Edge



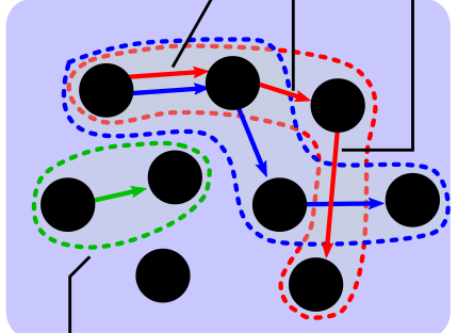
A standard edge between two vertices

Hyperedge



A hyperedge is an arbitrary set of vertices

Node-tuple

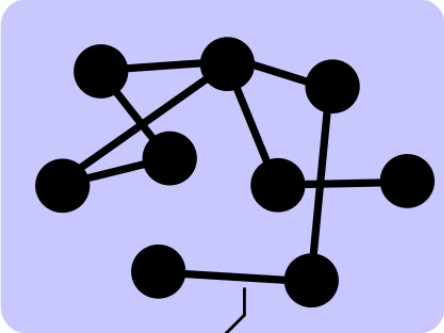


Node-tuple is an ordered tuple of vertices. Tuples can overlap

Higher-Order Graph Structures: Summary & Scope

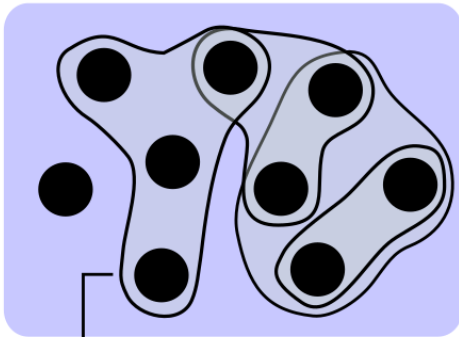
HO: „Anything that goes beyond a simple dyadic edge”

Edge



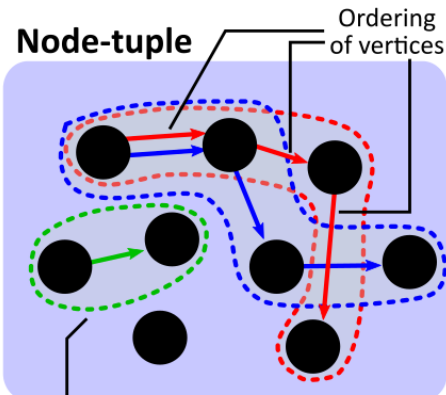
A standard edge between two vertices

Hyperedge



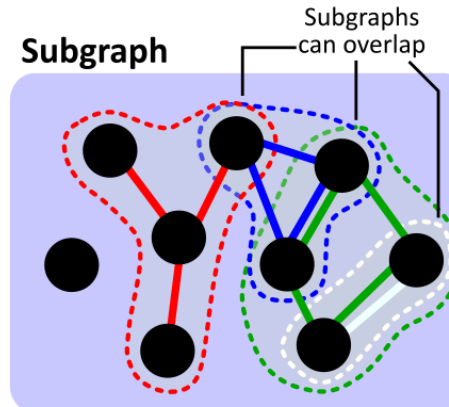
A hyperedge is an arbitrary set of vertices

Node-tuple



Node-tuple is an ordered tuple of vertices. Tuples can overlap

Subgraph

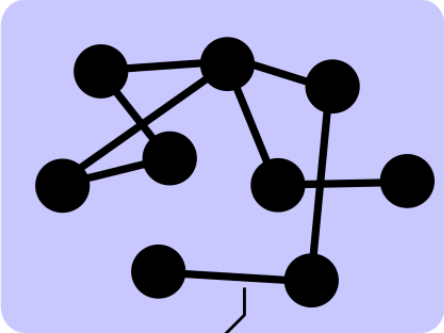


A subgraph can be used as an HO link - it is a subset of vertices and edges

Higher-Order Graph Structures: Summary & Scope

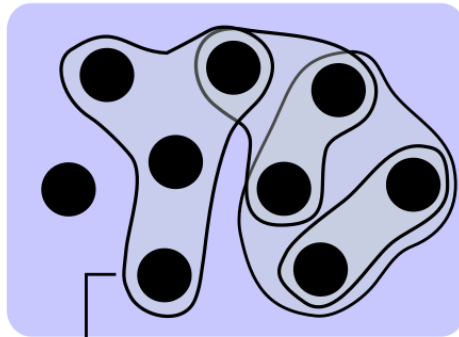
HO: „Anything that goes beyond a simple dyadic edge”

Edge



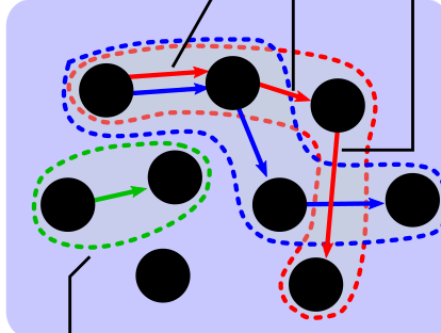
A standard edge between two vertices

Hyperedge



A hyperedge is an arbitrary set of vertices

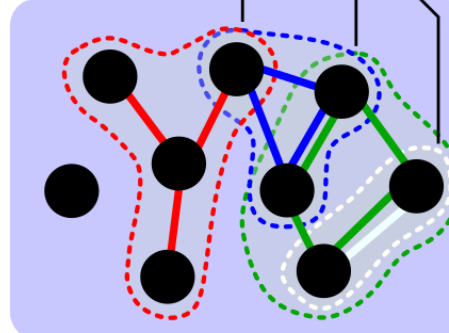
Node-tuple



Node-tuple is an ordered tuple of vertices. Tuples can overlap

Ordering of vertices

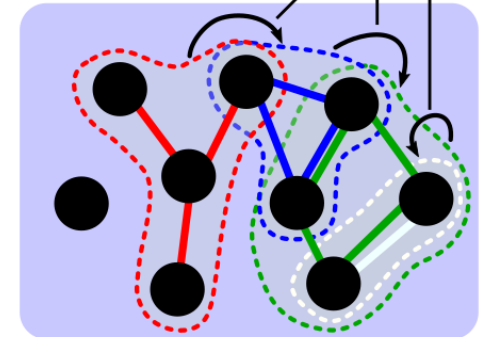
Subgraph



A subgraph can be used as an HO link - it is a subset of vertices and edges

Subgraphs can overlap

Subgraph-tuple



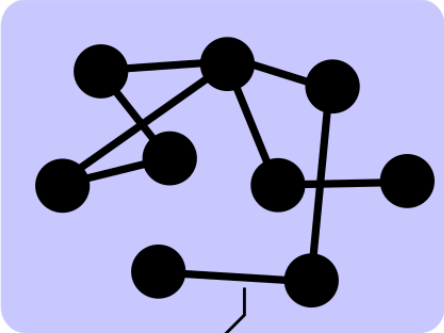
Subgraph-tuple is an ordered tuple of subgraphs used as HO links

Ordering of subgraphs

Higher-Order Graph Structures: Summary & Scope

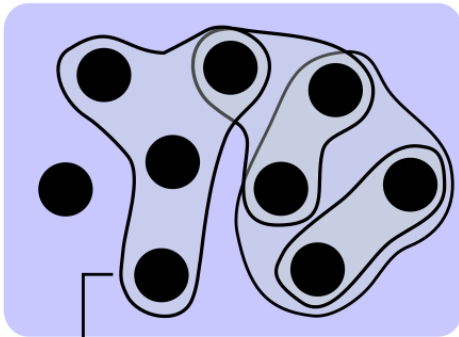
HO: „Anything that goes beyond a simple dyadic edge”

Edge



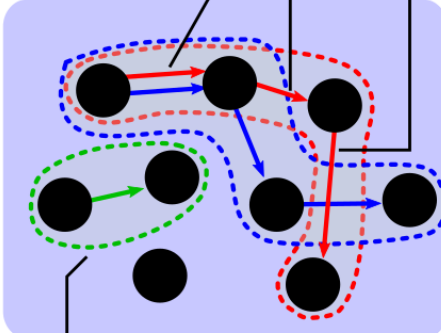
A standard edge between two vertices

Hyperedge



A hyperedge is an arbitrary set of vertices

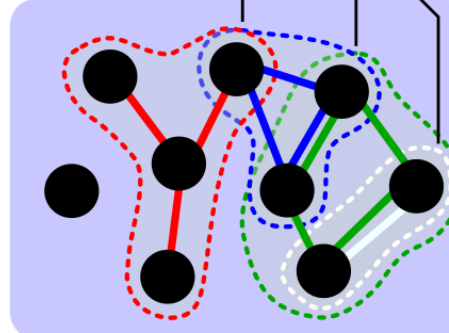
Node-tuple



Node-tuple is an ordered tuple of vertices. Tuples can overlap

Ordering of vertices

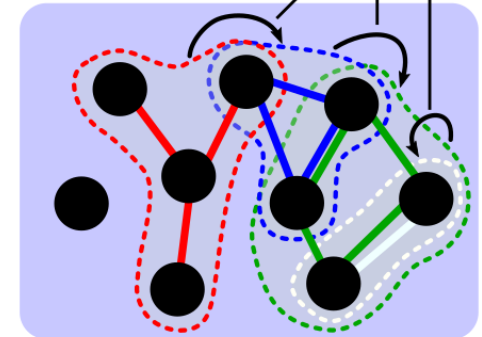
Subgraph



A subgraph can be used as an HO link - it is a subset of vertices and edges

Subgraphs can overlap

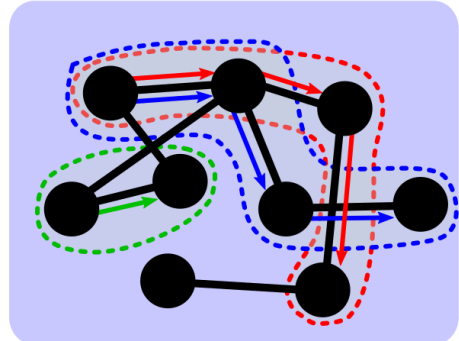
Subgraph-tuple



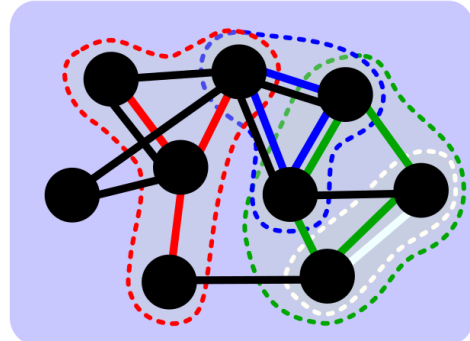
Subgraph-tuple is an ordered tuple of subgraphs used as HO links

Ordering of subgraphs

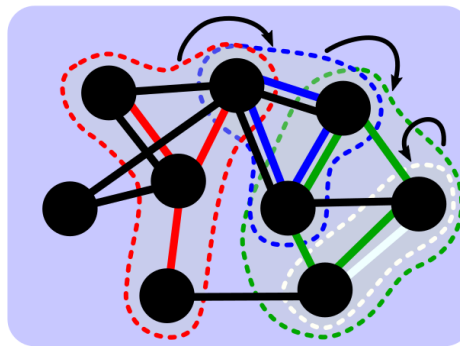
Node-tuples + edges



Subgraphs + edges



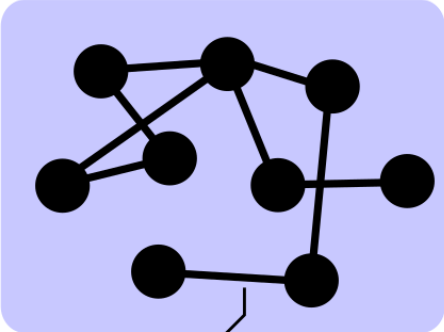
Subgraph-tuples + edges



Higher-Order Graph Structures: Summary & Scope

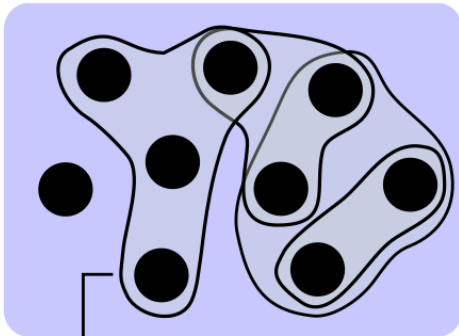
HO: „Anything that goes beyond a simple dyadic edge”

Edge



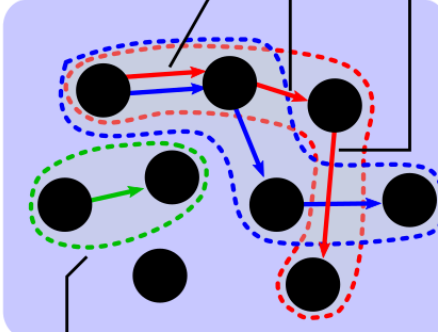
A standard edge between two vertices

Hyperedge



A hyperedge is an arbitrary set of vertices

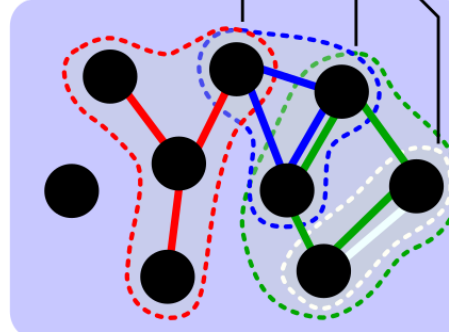
Node-tuple



Node-tuple is an ordered tuple of vertices. Tuples can overlap

Ordering of vertices

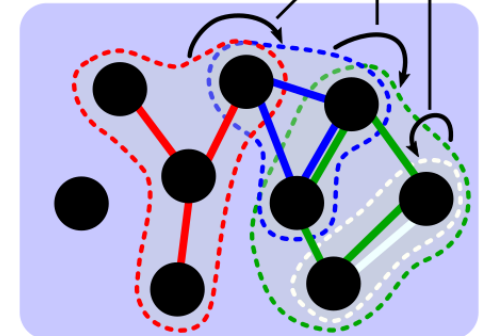
Subgraph



A subgraph can be used as an HO link - it is a subset of vertices and edges

Subgraphs can overlap

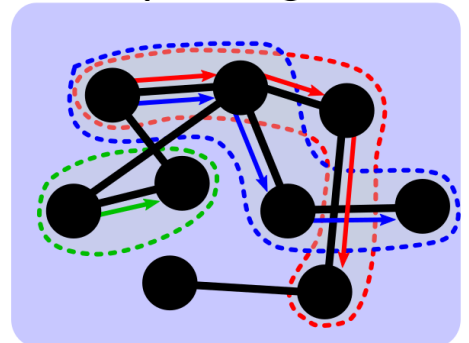
Subgraph-tuple



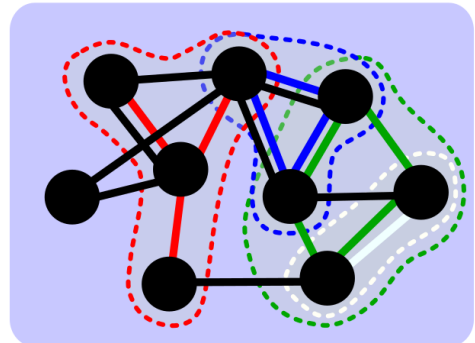
Subgraph-tuple is an ordered tuple of subgraphs used as HO links

Ordering of subgraphs

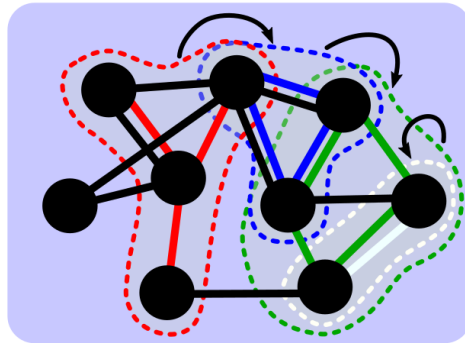
Node-tuples + edges



Subgraphs + edges



Subgraph-tuples + edges



Higher-Order Graph Structures: Summary & Scope

HO: „Anything that goes beyond a simple dyadic edge”

Edge

Hyperedge

Node-tuple

Ordering of vertices

Subgraph

Subgraphs can overlap

Subgraph-tuple

Ordering of subgraphs

Demystifying Higher-Order Graph Neural Networks

Maciej Besta^{1†}, Florian Scheidl¹, Lukas Gianinazzi¹, Shachar Klaiman², Jürgen Müller², Torsten Hoefler¹

¹ETH Zurich ²BASF SE

Node-tuple

Higher-Order Graph Structures: Summary & Scope

HO: „Anything that goes beyond a simple dyadic edge“

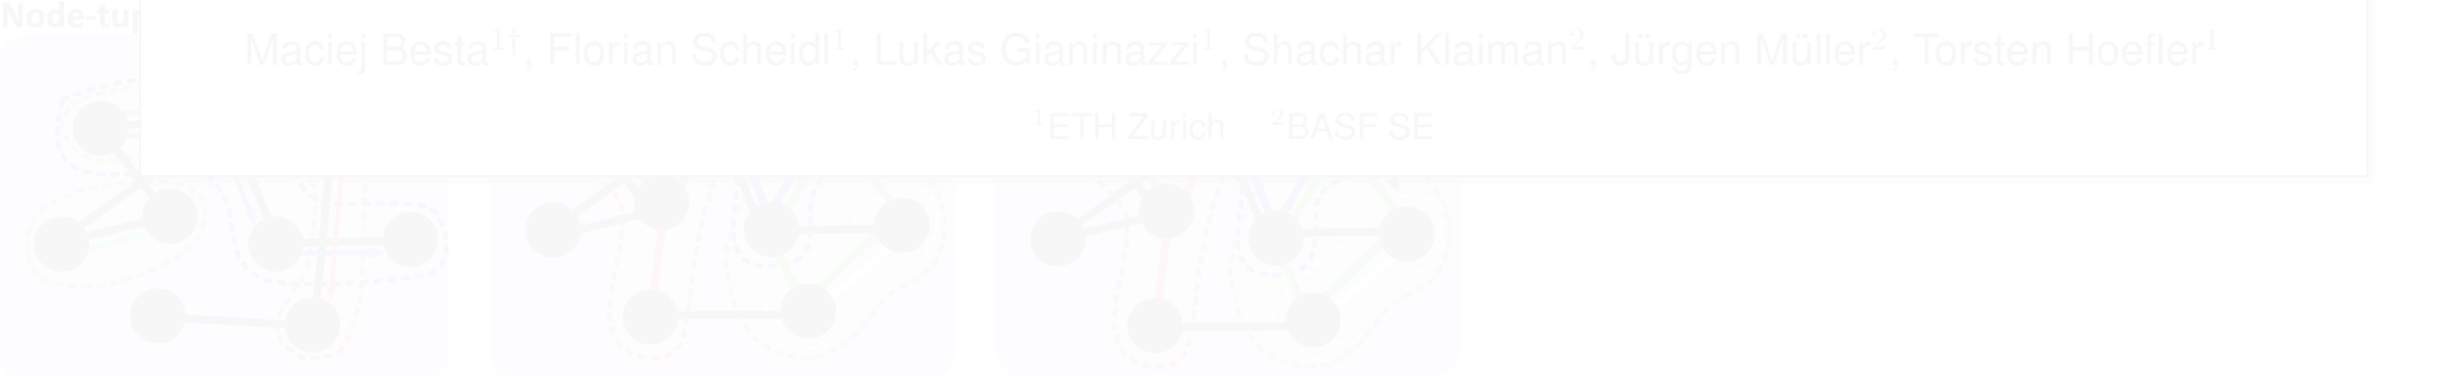


So we want to harness HO...
But what's the whole pipeline architecture?

Derivational Neural Networks

Maciej Besta^{1†}, Florian Scheidl¹, Lukas Gianinazzi¹, Shachar Klaiman², Jürgen Müller², Torsten Hoefler¹

¹ETH Zurich ²BASF SE



HO-Enhanced Pipeline for Dynamic Link Prediction

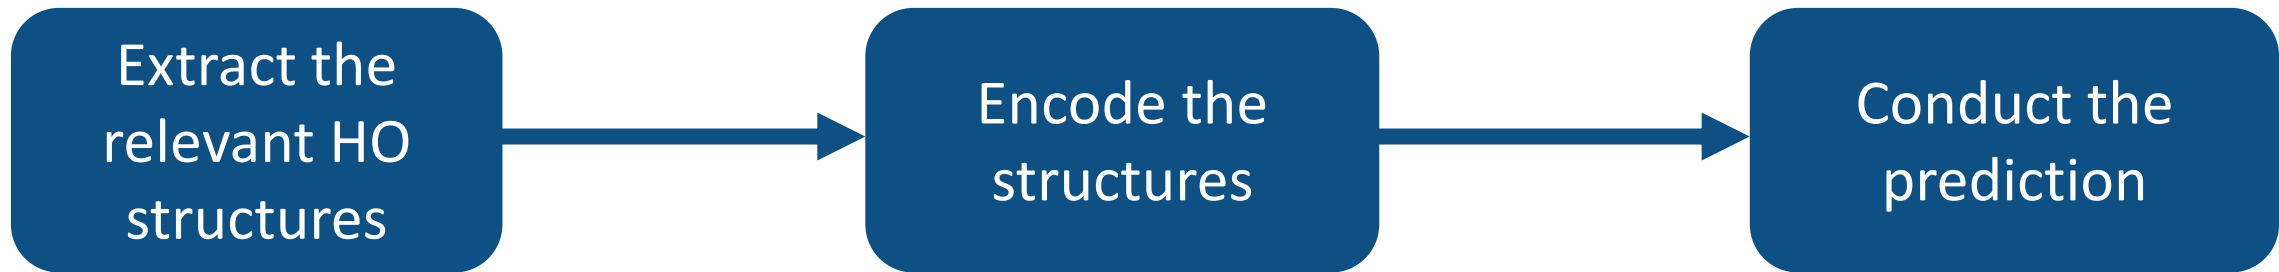
HO-Enhanced Pipeline for Dynamic Link Prediction

Extract the
relevant HO
structures

HO-Enhanced Pipeline for Dynamic Link Prediction



HO-Enhanced Pipeline for Dynamic Link Prediction



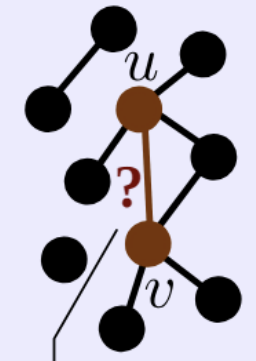
HO-Enhanced Pipeline for Dynamic Link Prediction



Temporal HO Structures

Temporal Higher-Order (HO) Example

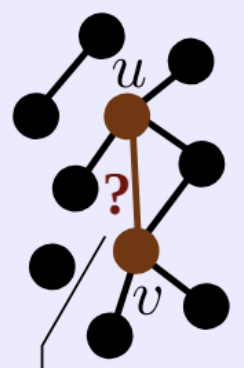
time



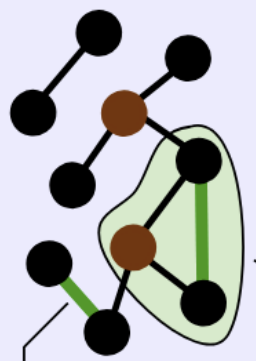
Will vertices
u and v be
connected?

Temporal HO Structures

Temporal Higher-Order (HO) Example



Will vertices u and v be connected?




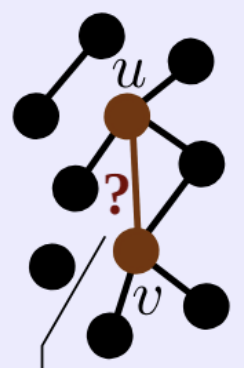
green:
added edge

A **triangle**
appearing at a
certain time

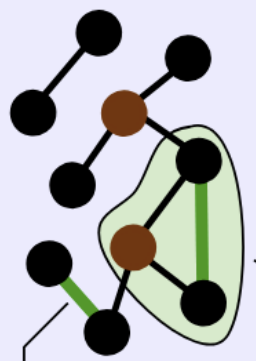
Temporal HO Structures

Temporal Higher-Order (HO) Example

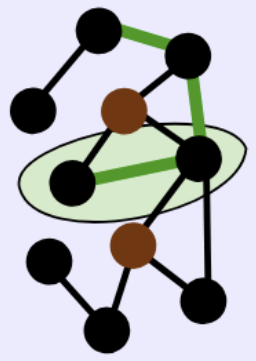
time 



Will vertices u and v be connected?



green:
added edge

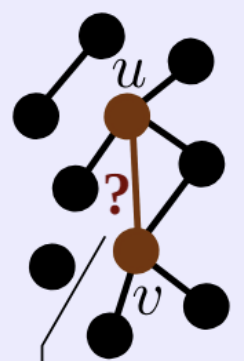


A **triangle**
appearing at a
certain time

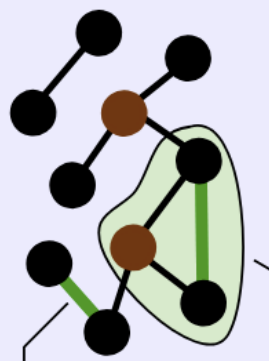
Temporal HO Structures

Temporal Higher-Order (HO) Example

time →

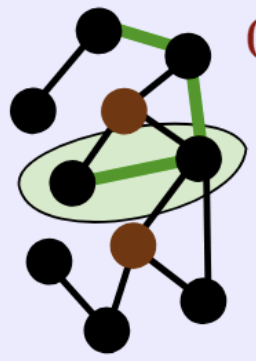


Will vertices u and v be connected?

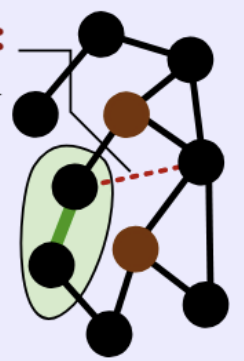


green: added edge

A triangle appearing at a certain time

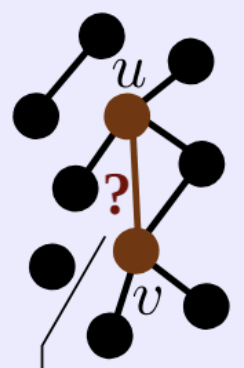
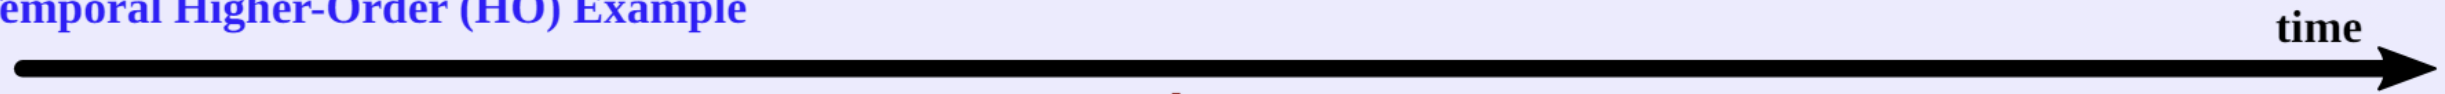


red (dashed): removed edge

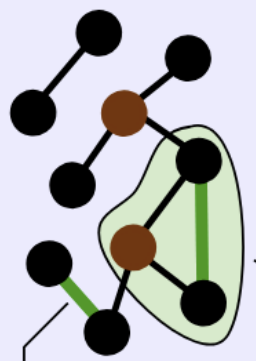


Temporal HO Structures

Temporal Higher-Order (HO) Example

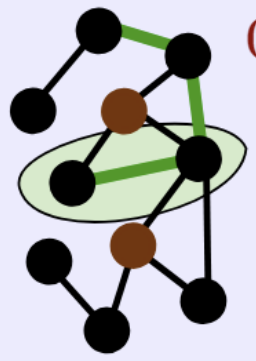


Will vertices u and v be connected?

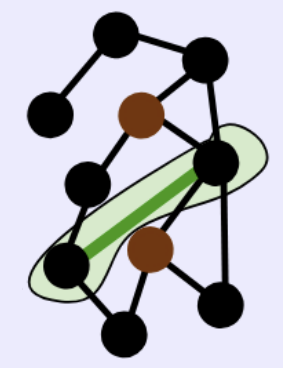
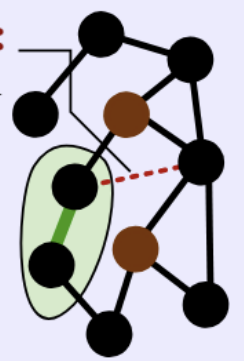


green:
added edge

A **triangle**
appearing at a
certain time

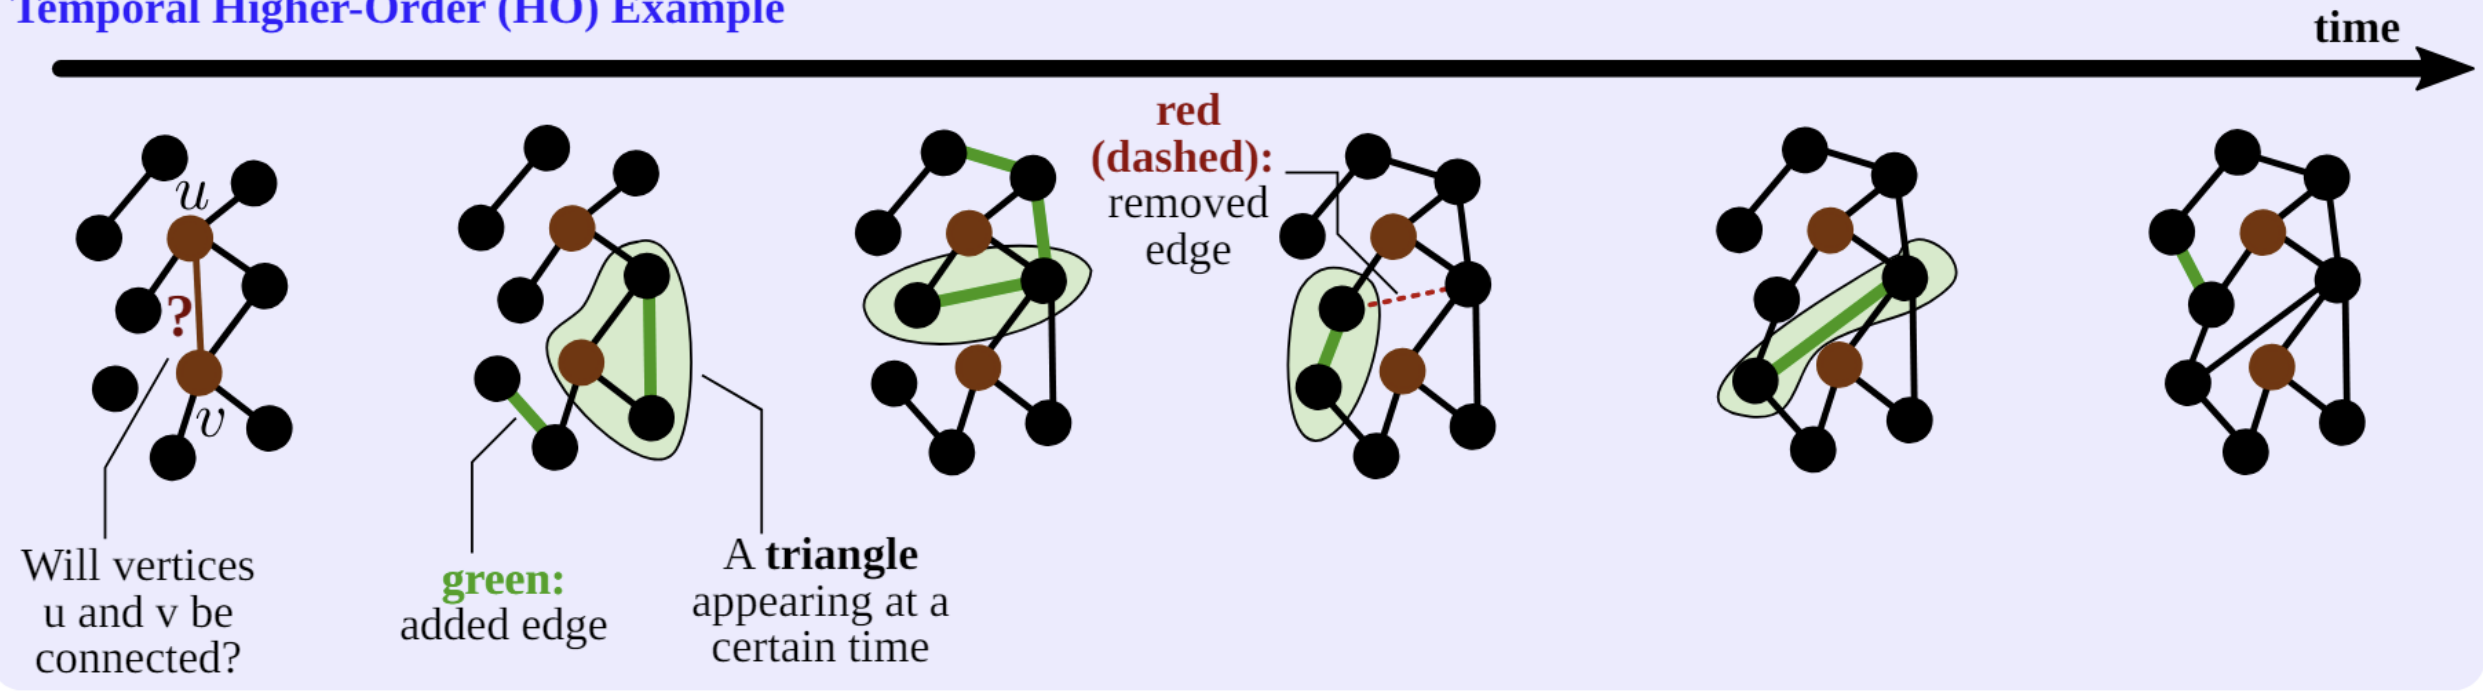


**red
(dashed):**
removed
edge



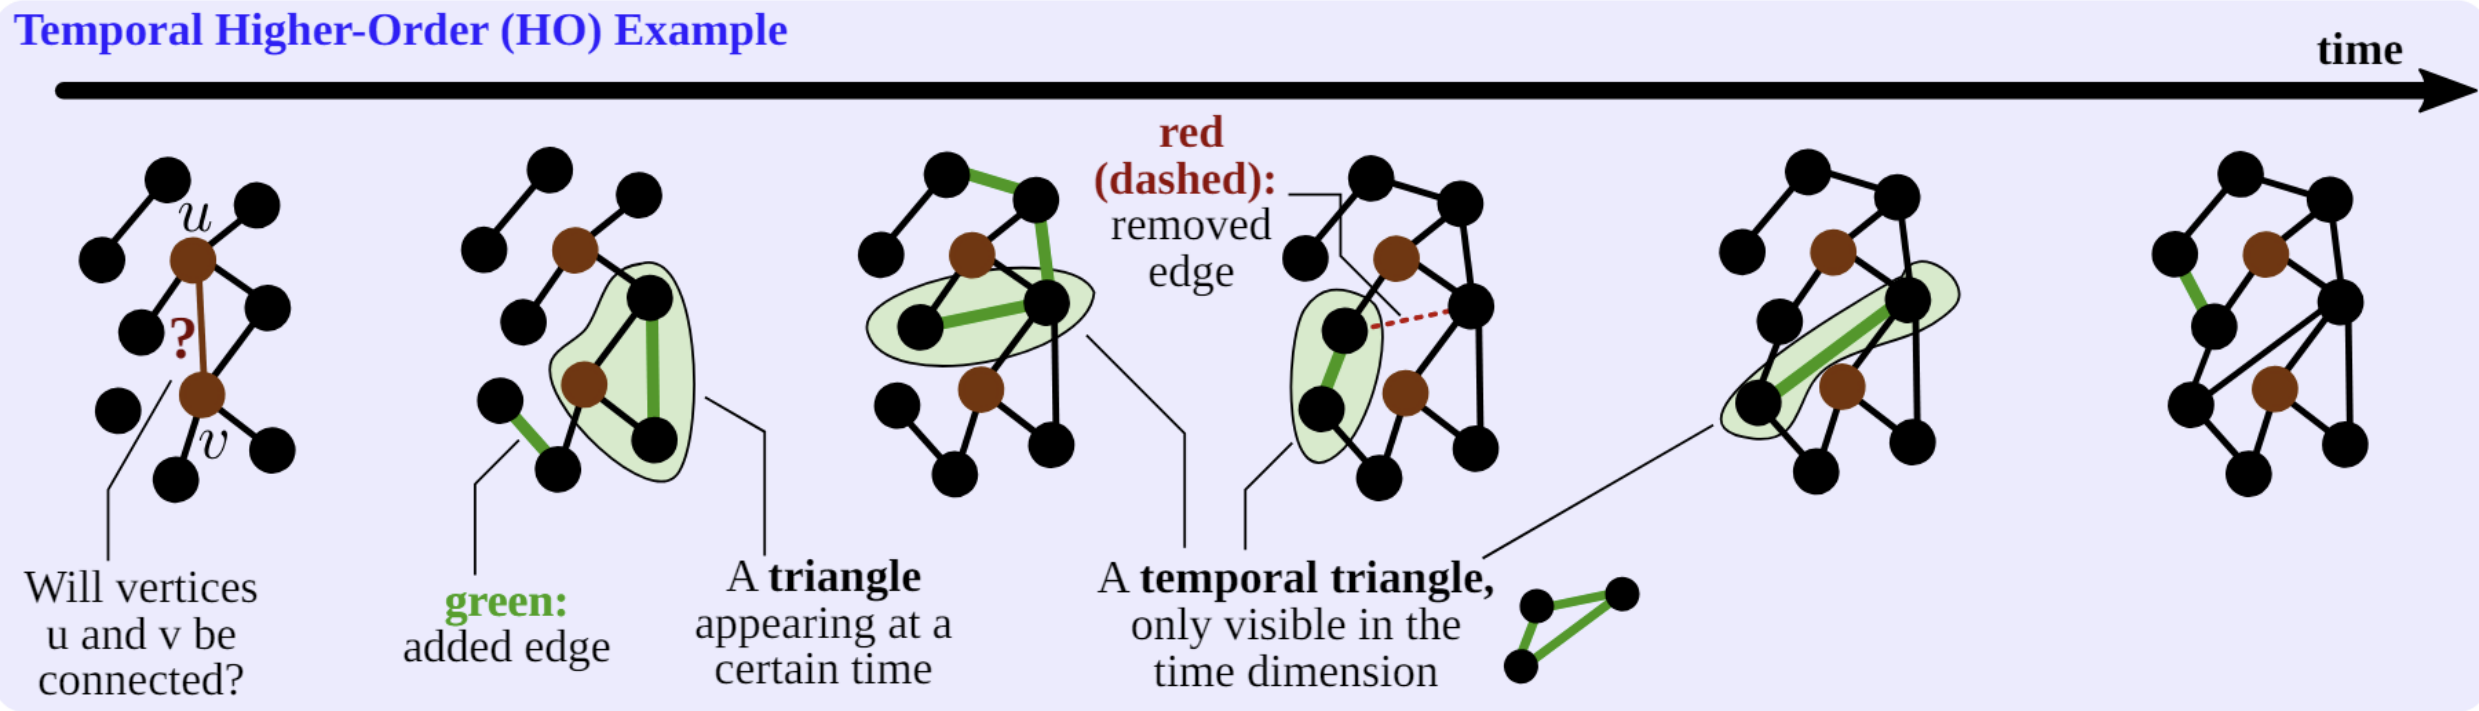
Temporal HO Structures

Temporal Higher-Order (HO) Example



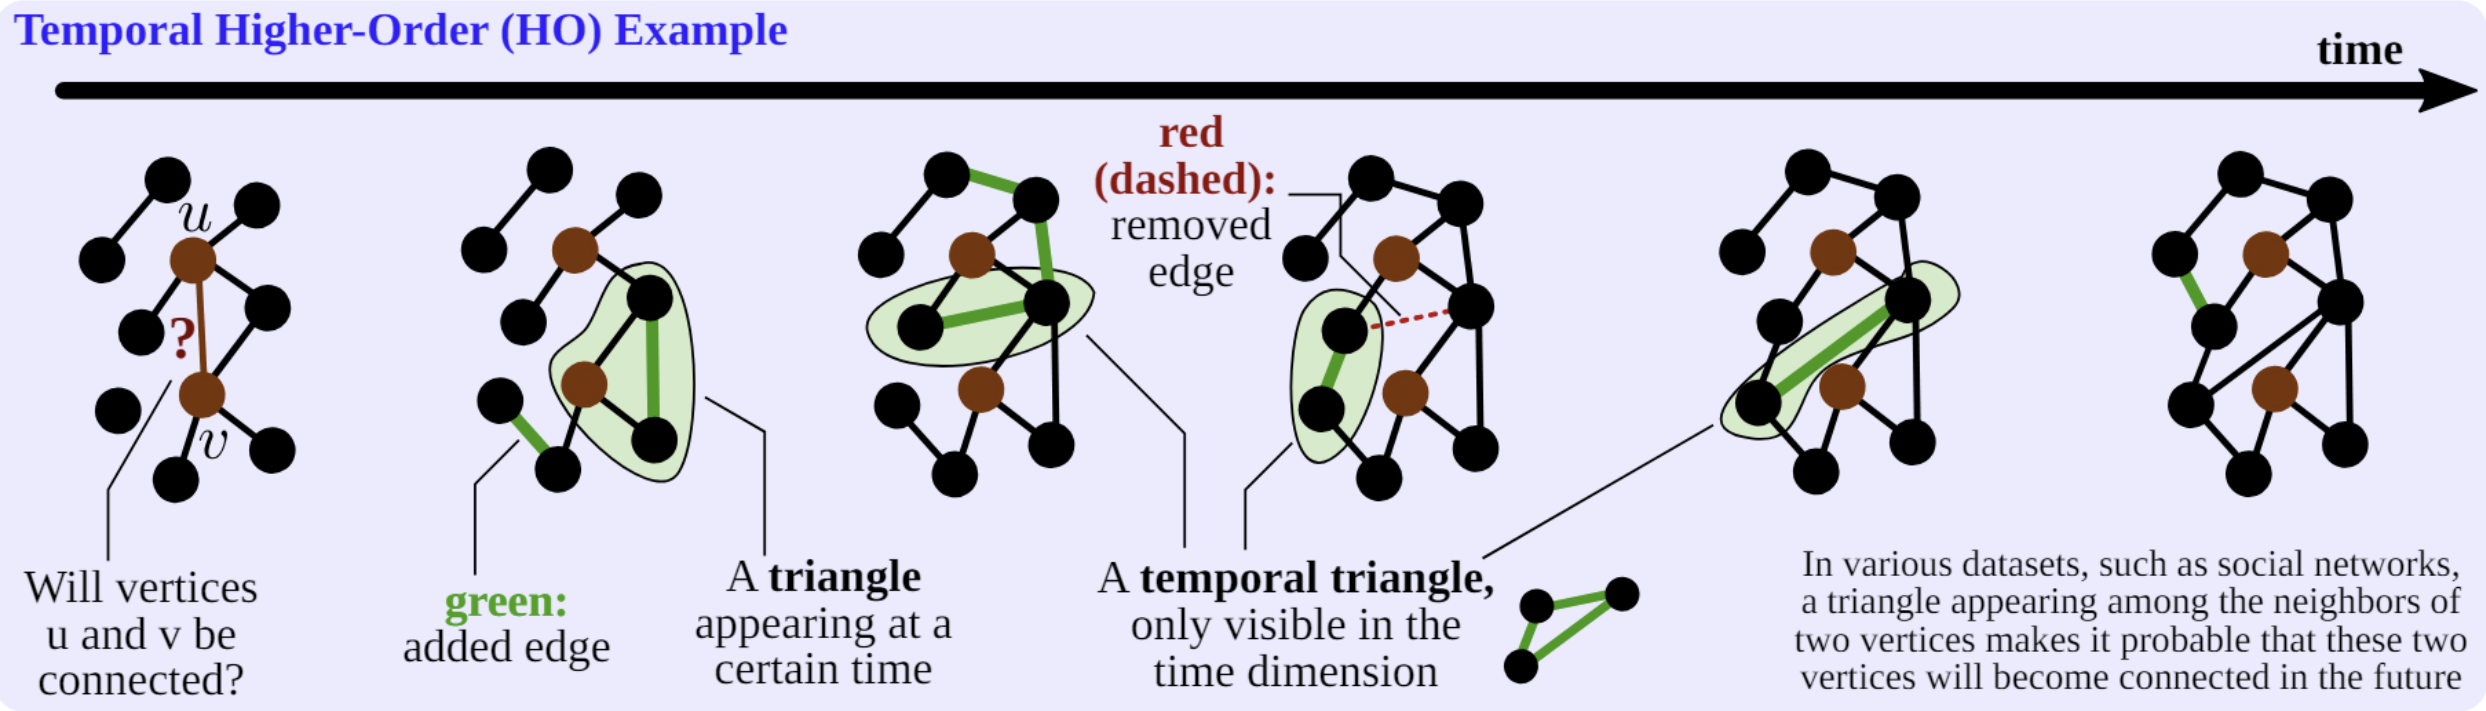
Temporal HO Structures

Temporal Higher-Order (HO) Example



Temporal HO Structures

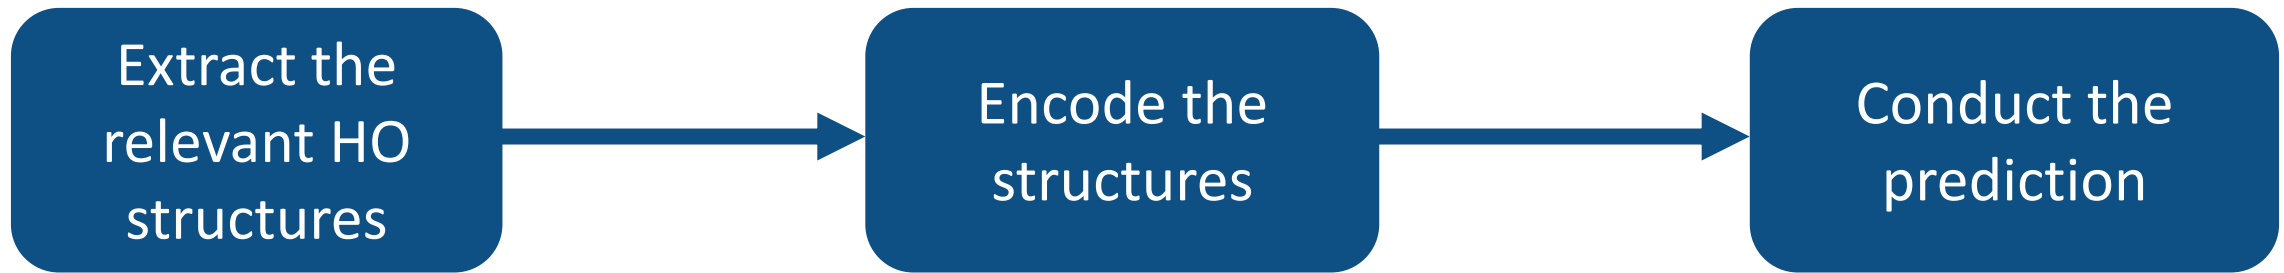
Temporal Higher-Order (HO) Example



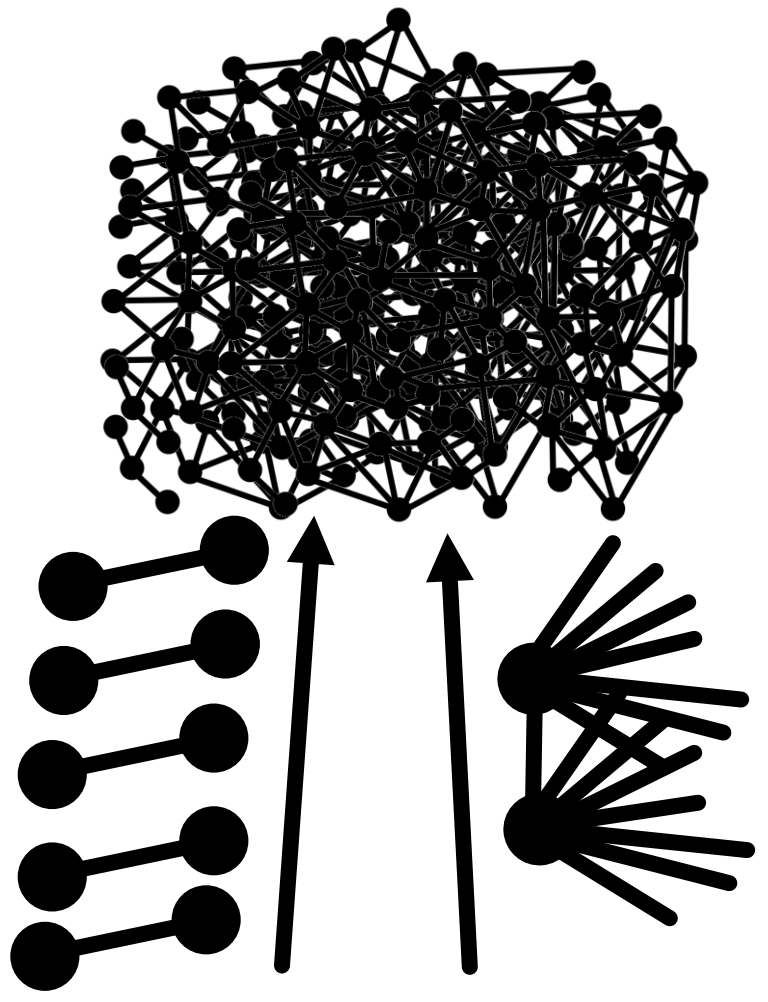
HO-Enhanced Pipeline for Dynamic Link Prediction



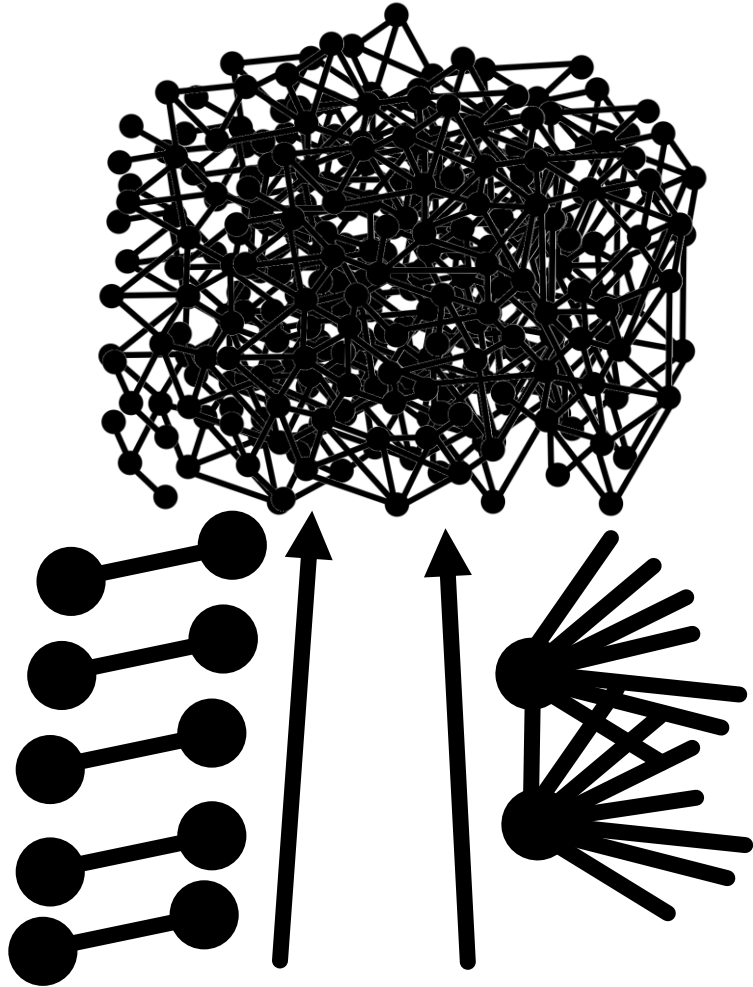
HO-Enhanced Pipeline for Dynamic Link Prediction



Formal Setting of Dynamic Link Prediction

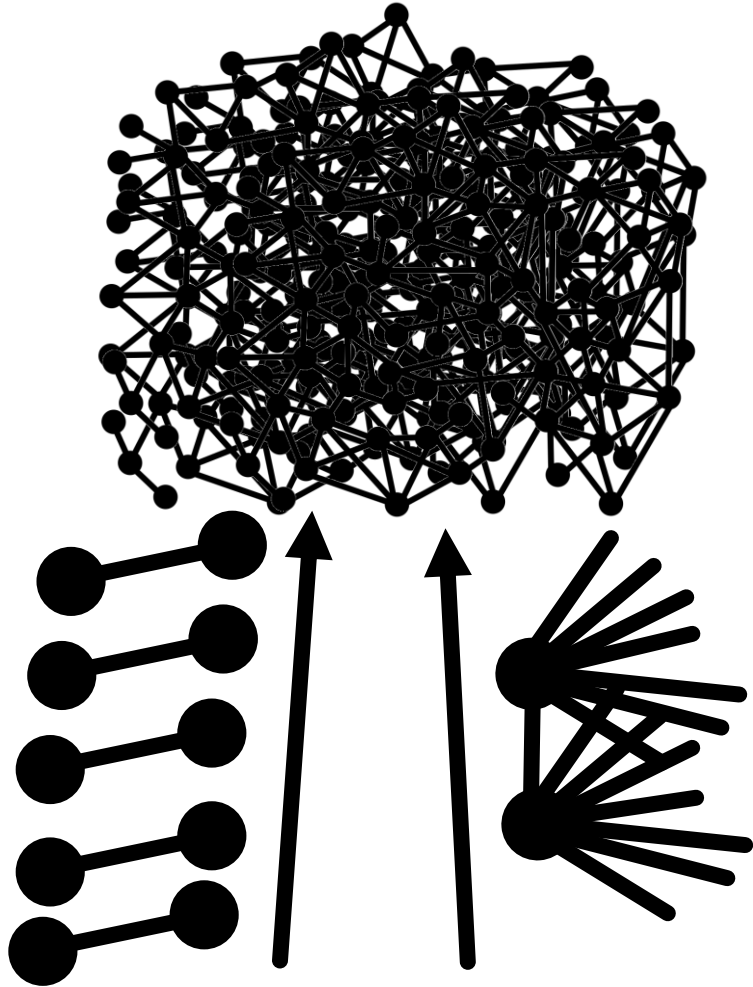


Formal Setting of Dynamic Link Prediction



Continuous-Time Dynamic Graph (CTDG) representation is a tuple $(G(0), T)$, where...

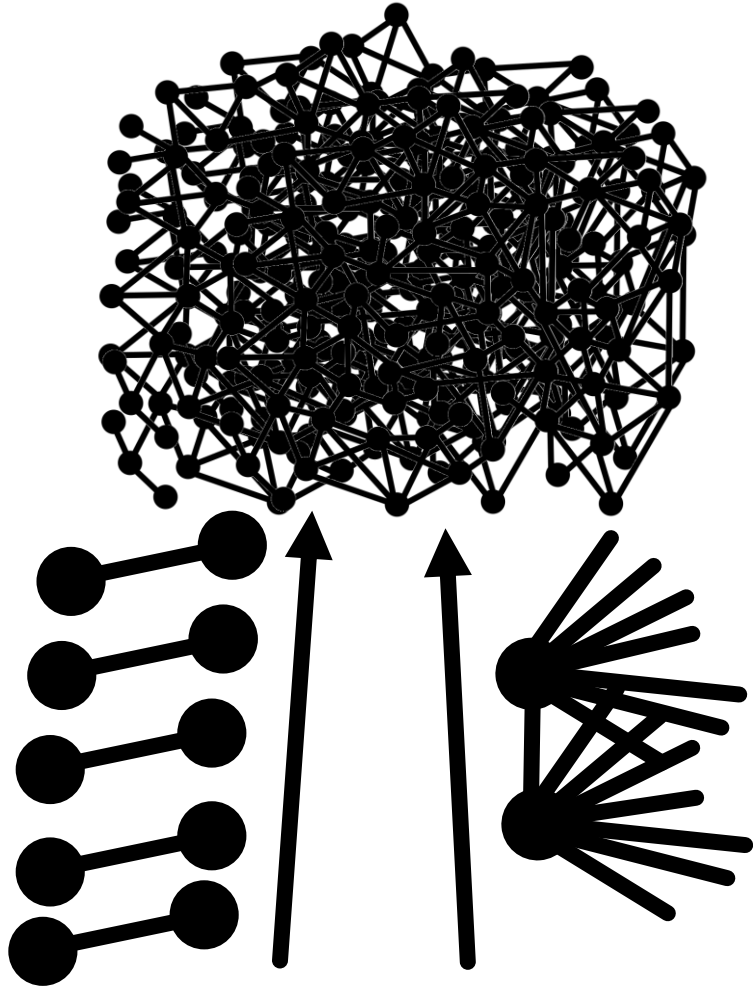
Formal Setting of Dynamic Link Prediction



Continuous-Time Dynamic Graph (CTDG) representation is a tuple $(G(0), T)$, where...

... $G(0) = (V(0), E(0), f(0), w(0))$ represents the initial state of the graph

Formal Setting of Dynamic Link Prediction

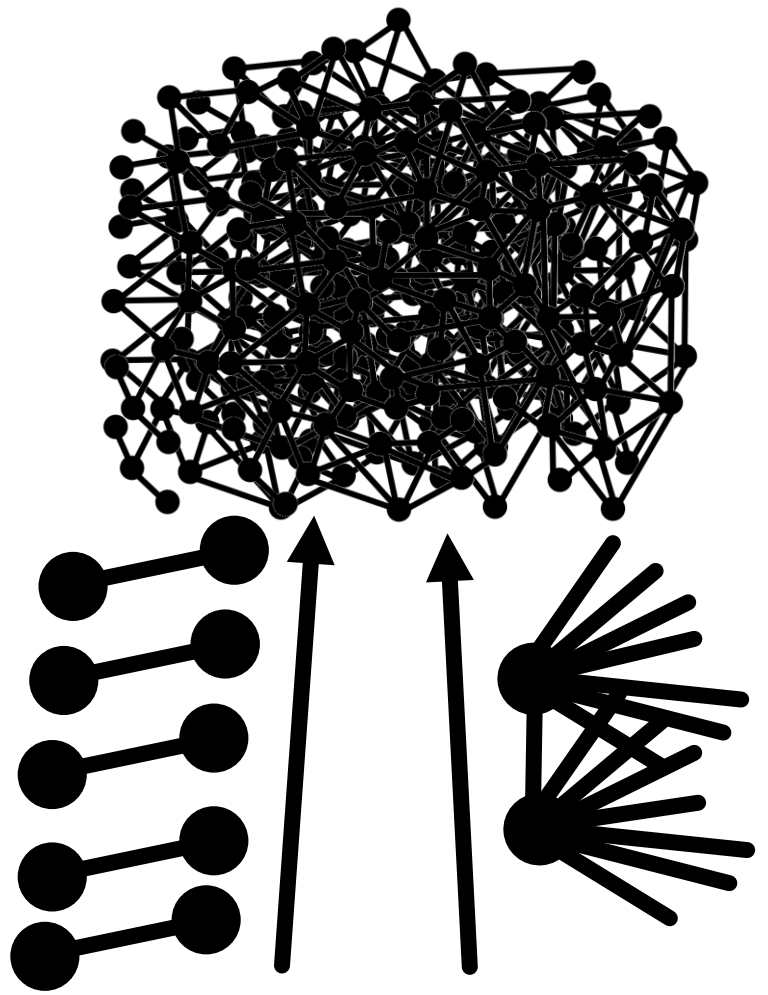


Continuous-Time Dynamic Graph (CTDG) representation is a tuple $(G(0), T)$, where...

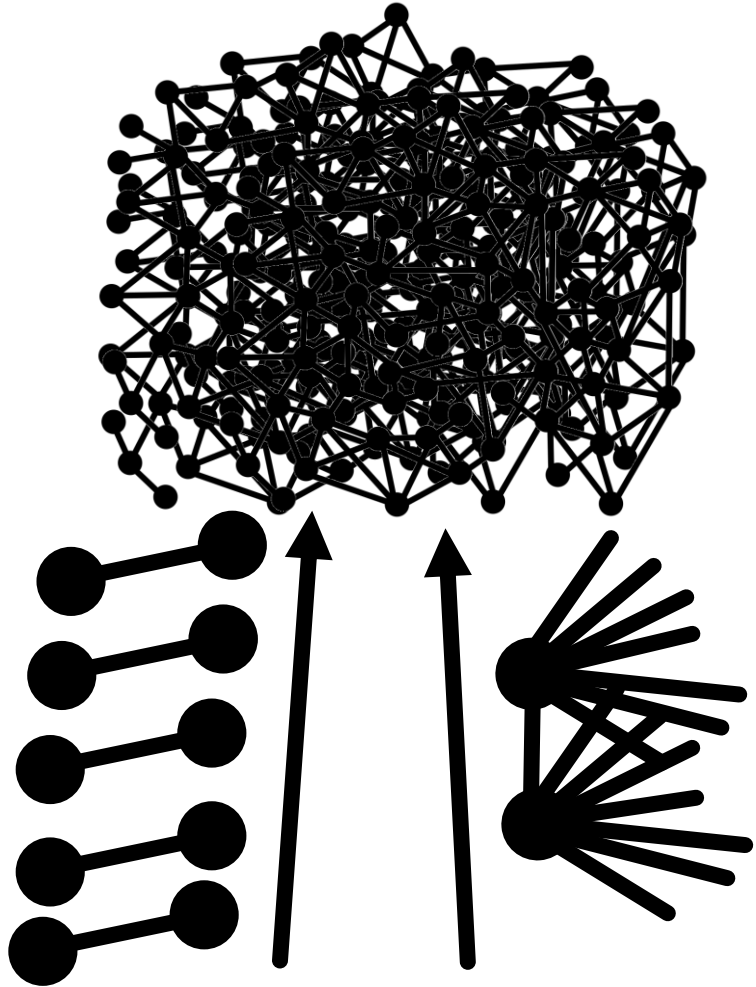
... $G(0) = (V(0), E(0), f(0), w(0))$ represents the initial state of the graph

... T is a set of tuples of the form (timestamp, event) representing events to be applied to the graph at given timestamps.

Formal Setting of Dynamic Link Prediction



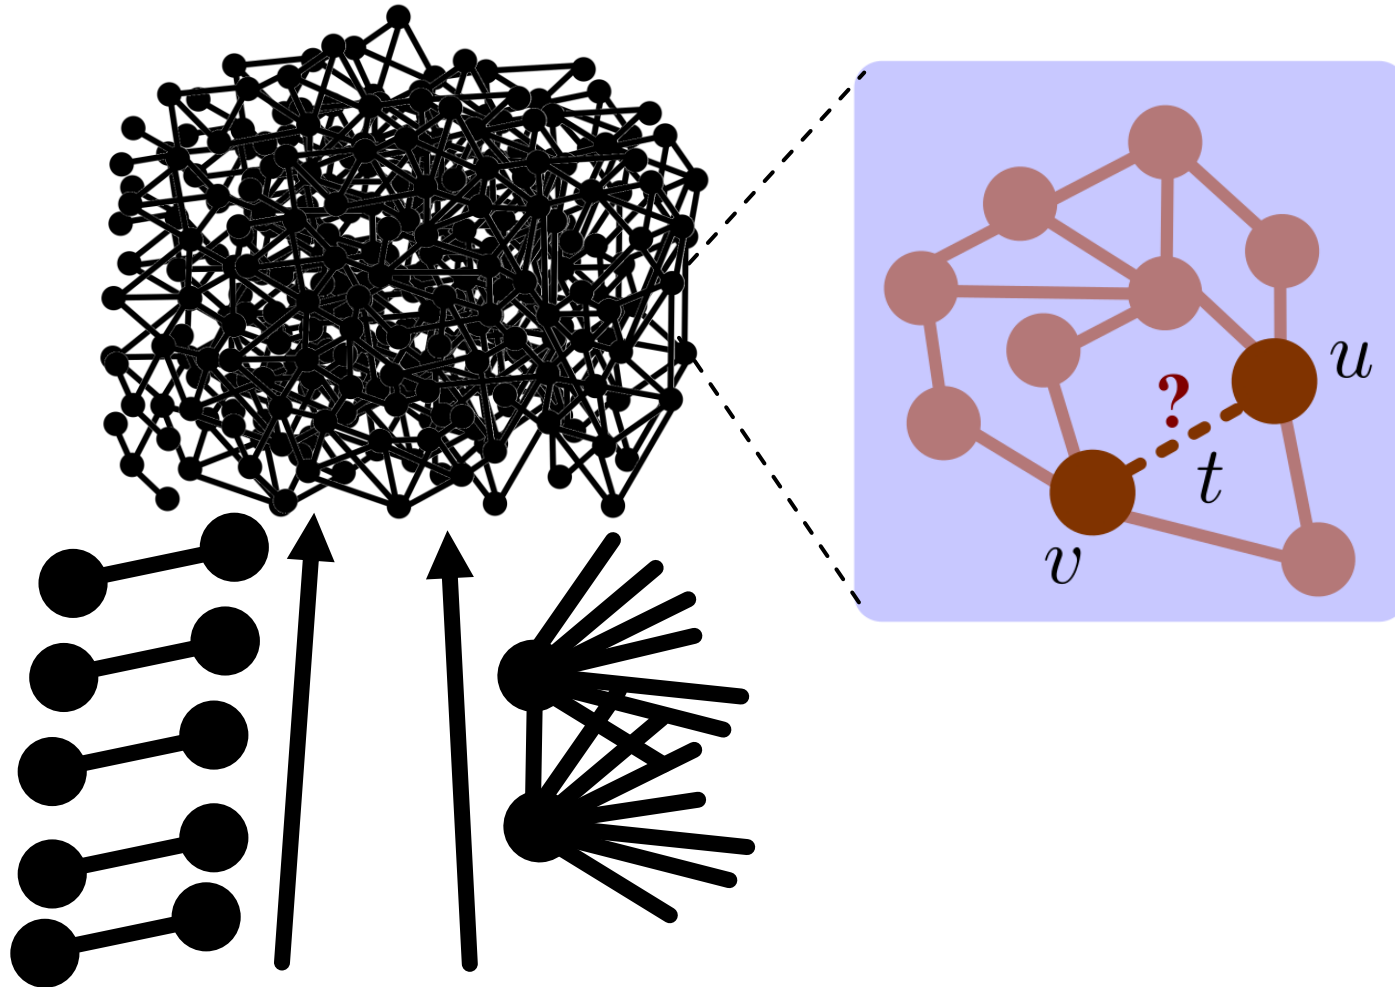
Formal Setting of Dynamic Link Prediction



Assume:

- a CTDG $(G(0), T)$,
- a timestamp $t \in \mathbb{N}$,
- an edge (u, v) .

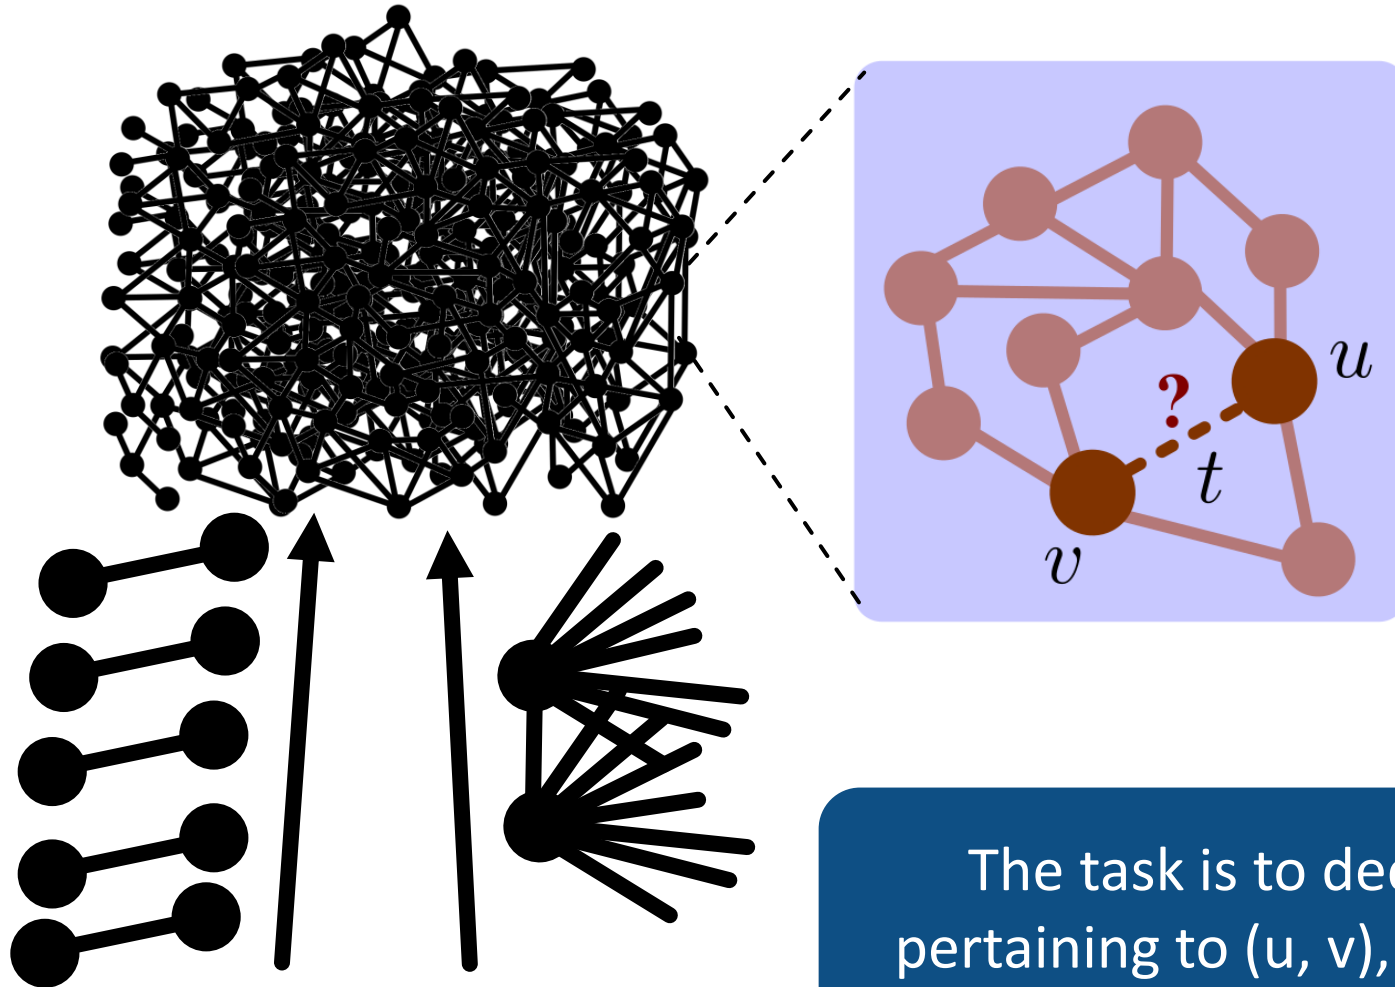
Formal Setting of Dynamic Link Prediction



Assume:

- a CTDG $(G(0), T)$,
- a timestamp $t \in \mathbb{N}$,
- an edge (u, v) .

Formal Setting of Dynamic Link Prediction

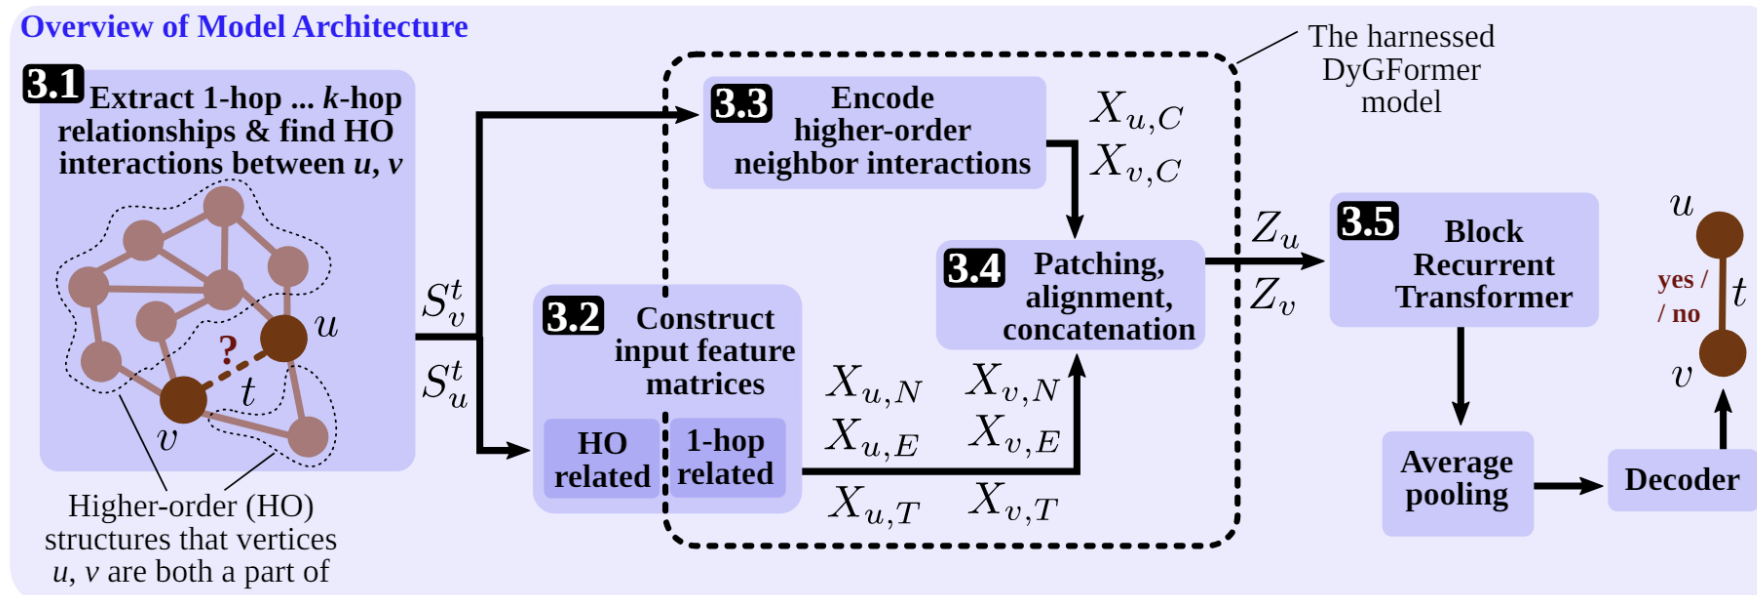


Assume:

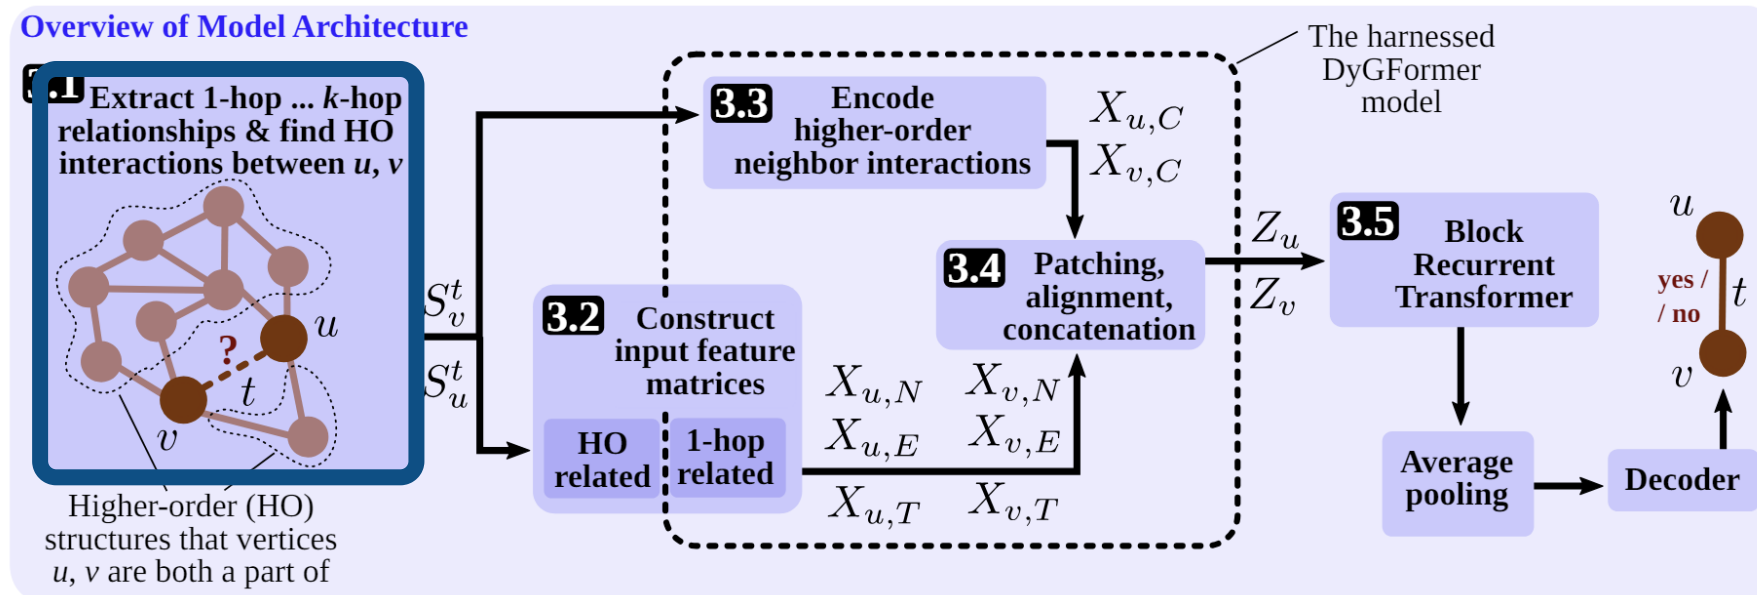
- a CTDG $(G(0), T)$,
- a timestamp $t \in \mathbb{N}$,
- an edge (u, v) .

The task is to decide, whether there is some event e pertaining to (u, v) , such that $(t, e) \in \{(t', e) \in T \mid t' = t\}$, while only considering the CTDG $(G(0), \{(t', e) \in T \mid t' < t\})$.

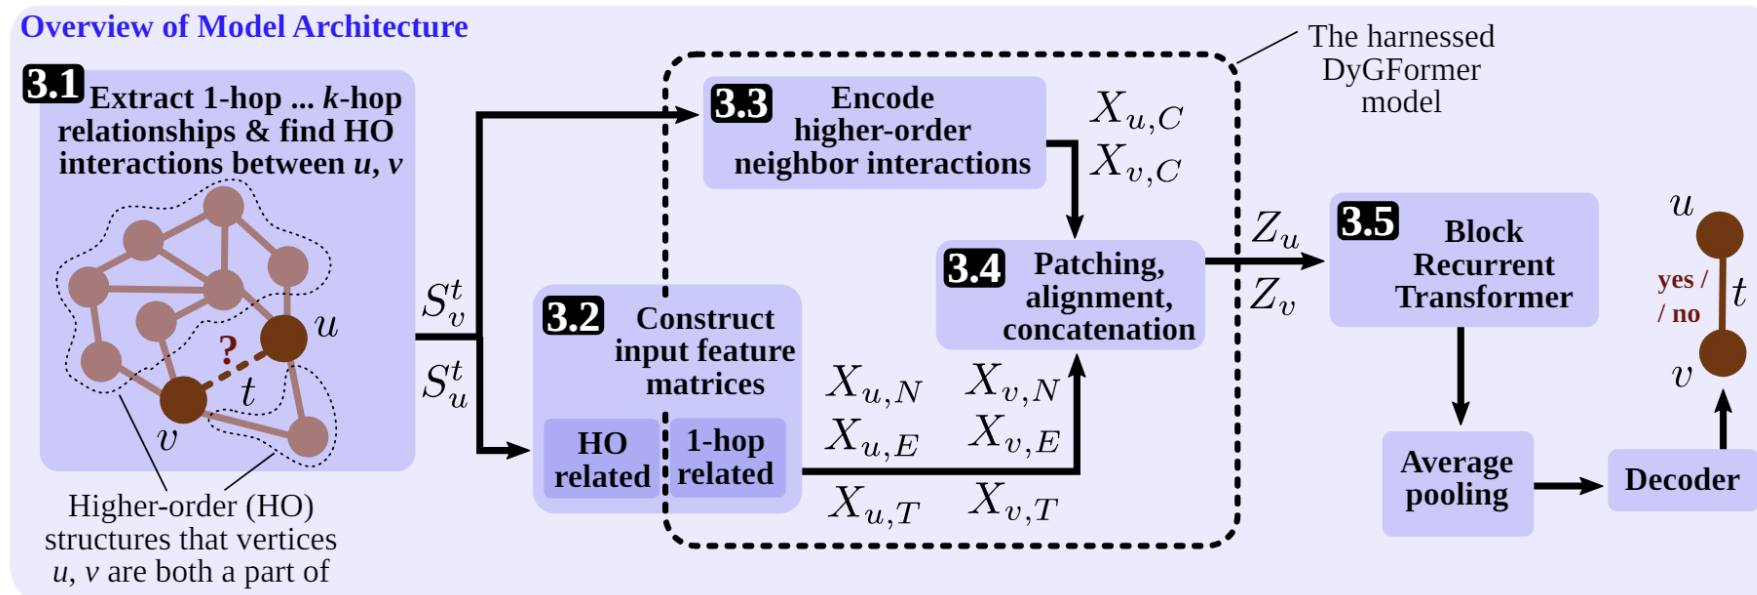
Encoding HO Structures



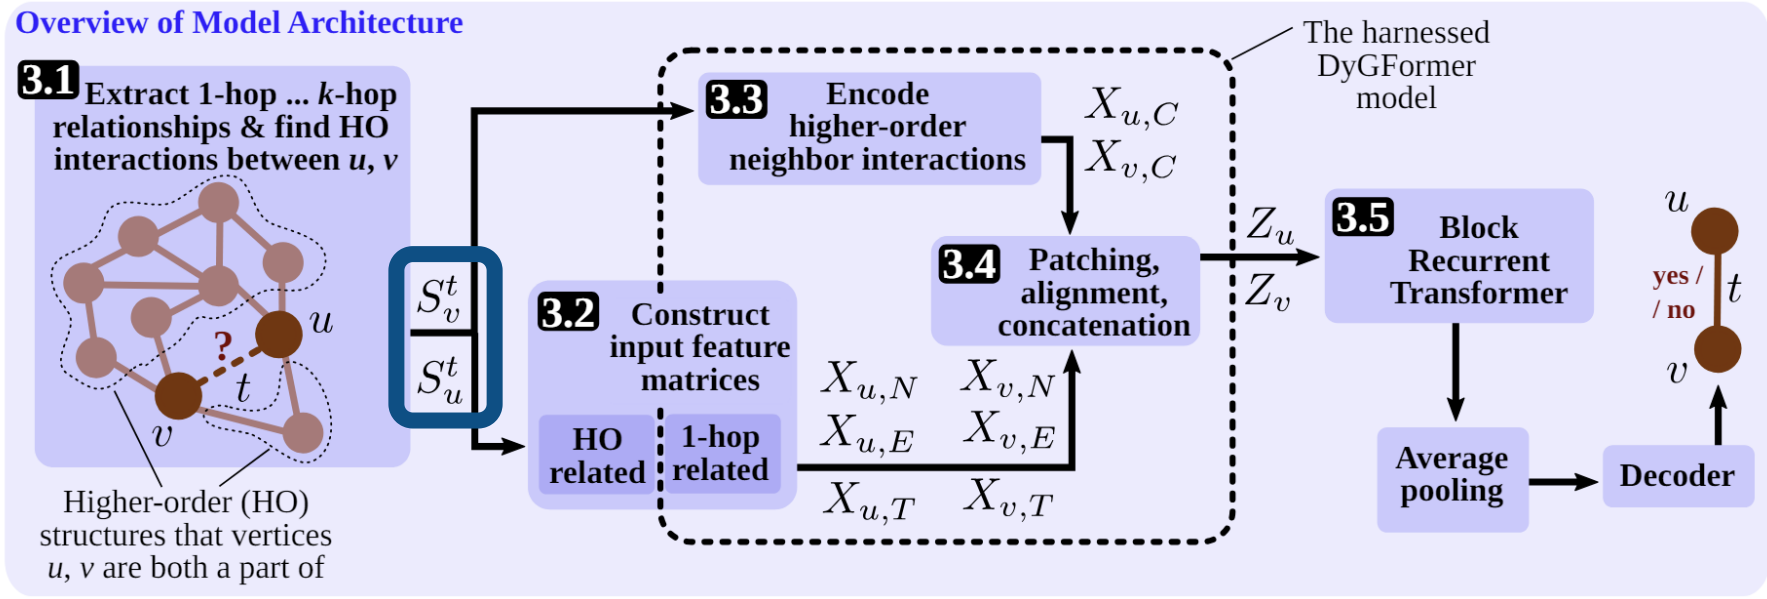
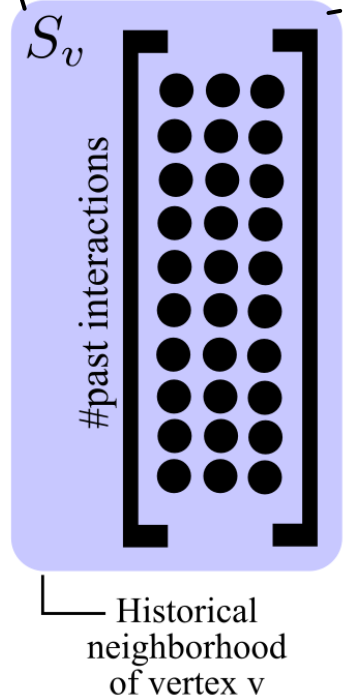
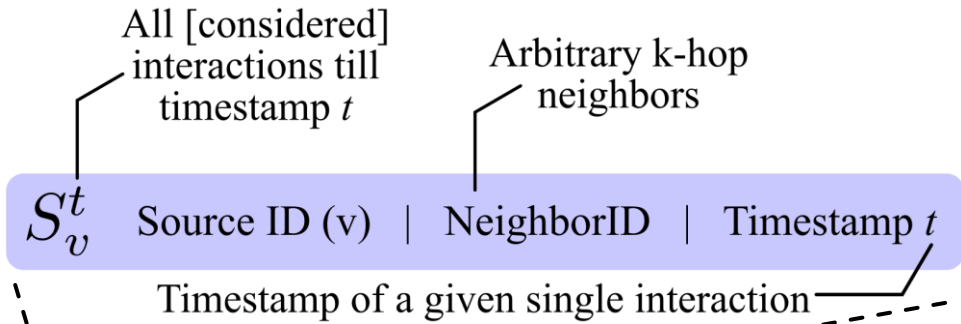
Encoding HO Structures



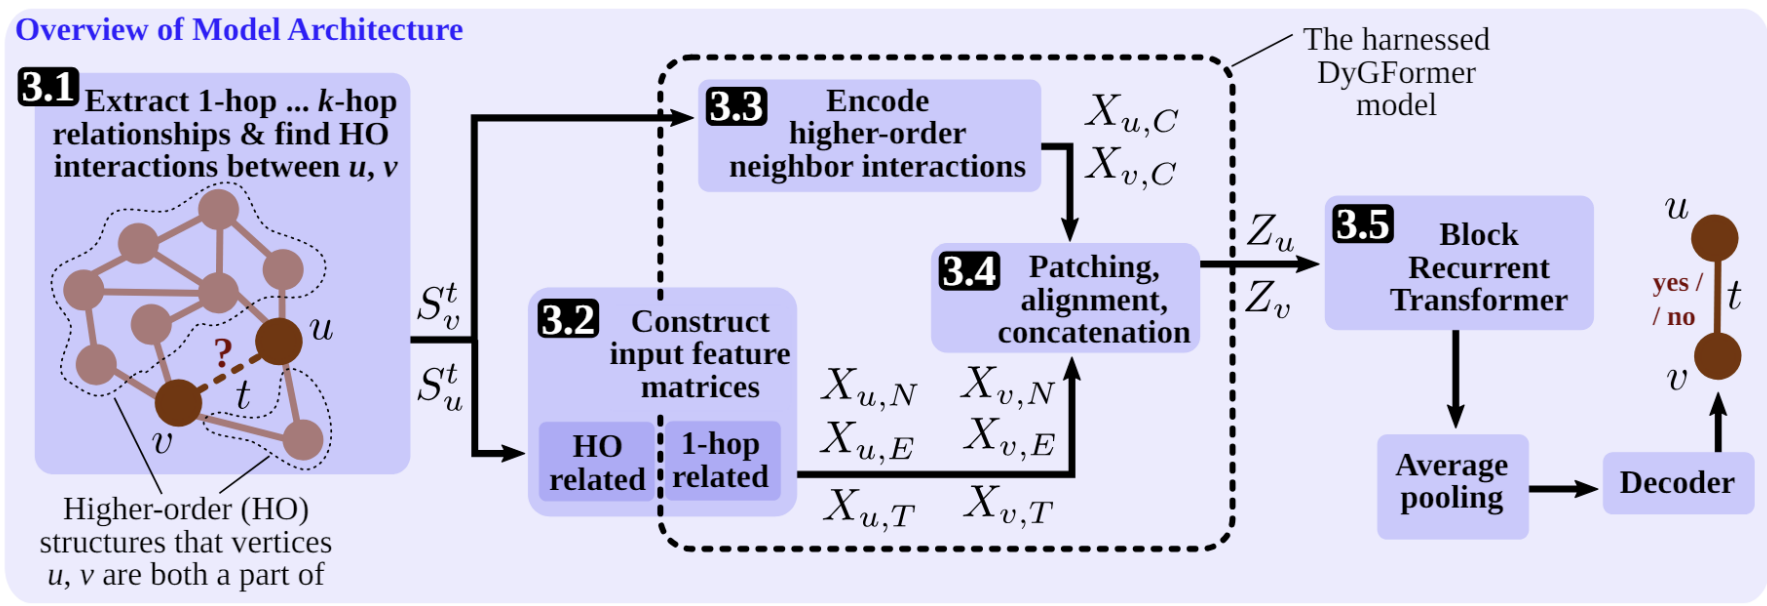
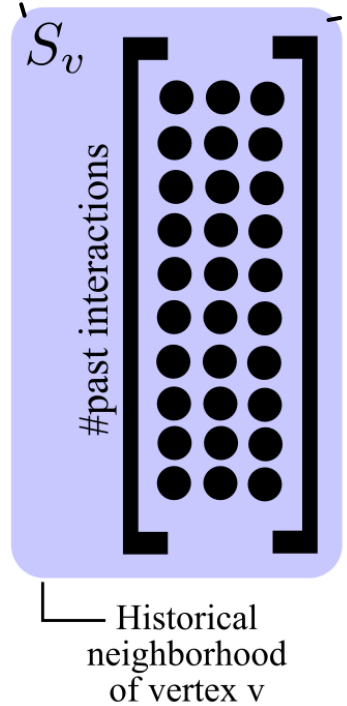
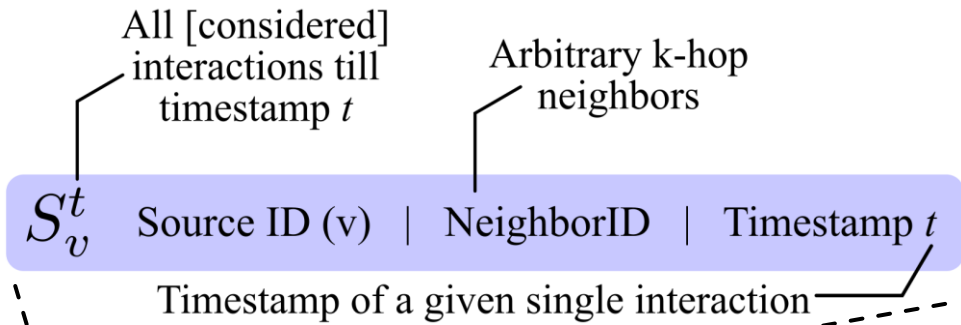
Encoding HO Structures



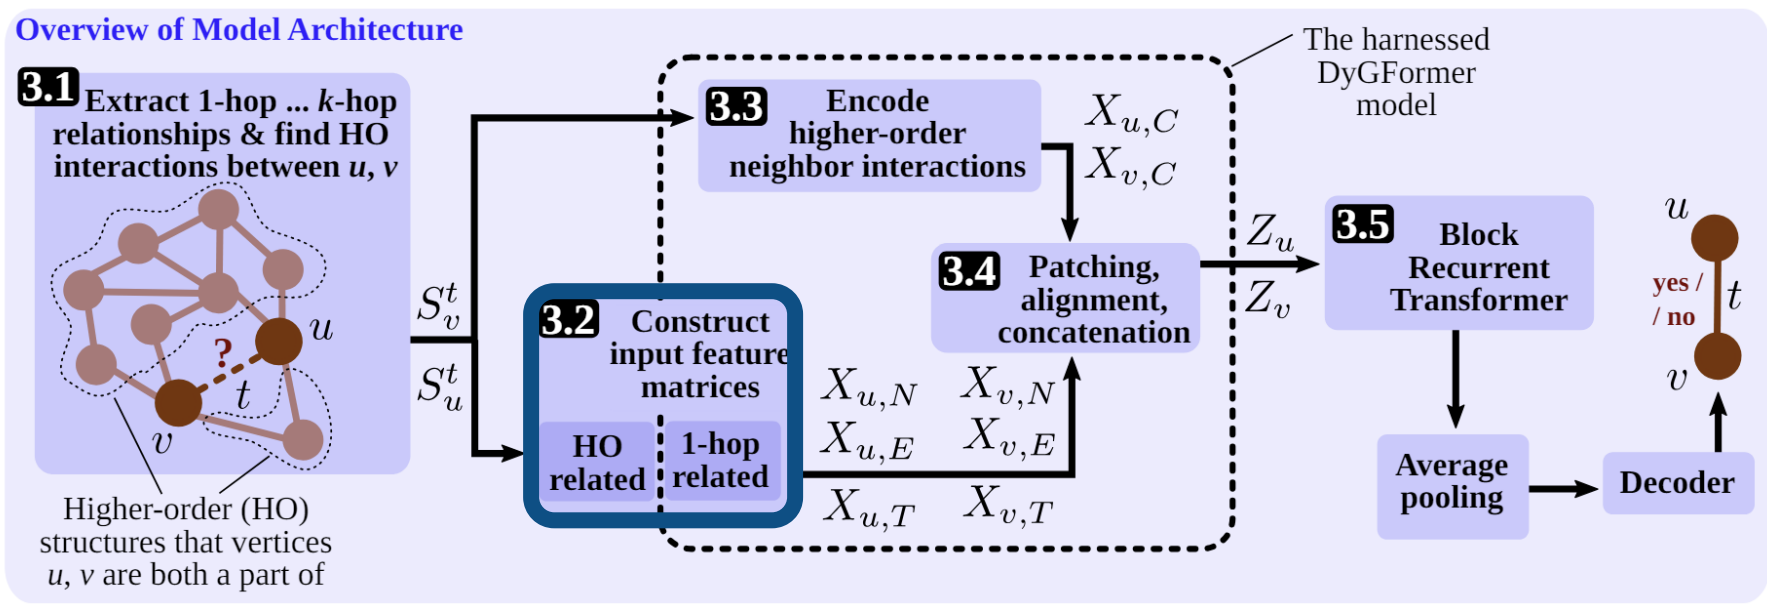
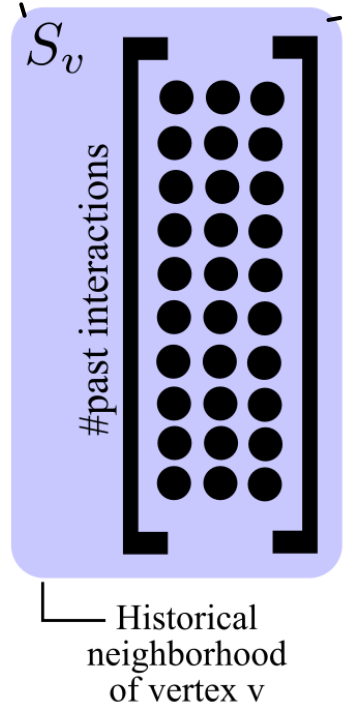
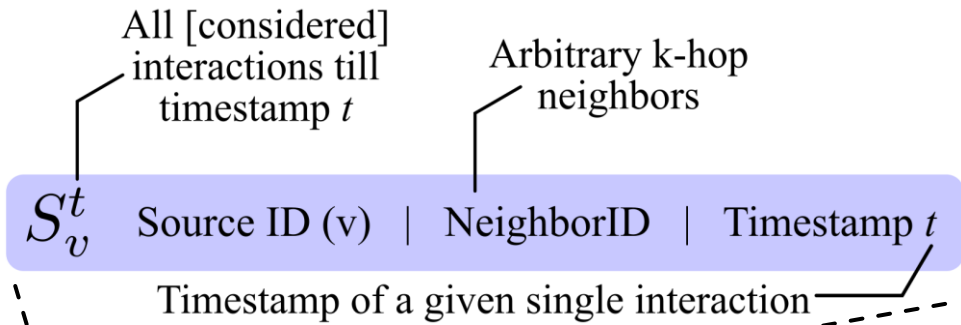
Encoding HO Structures



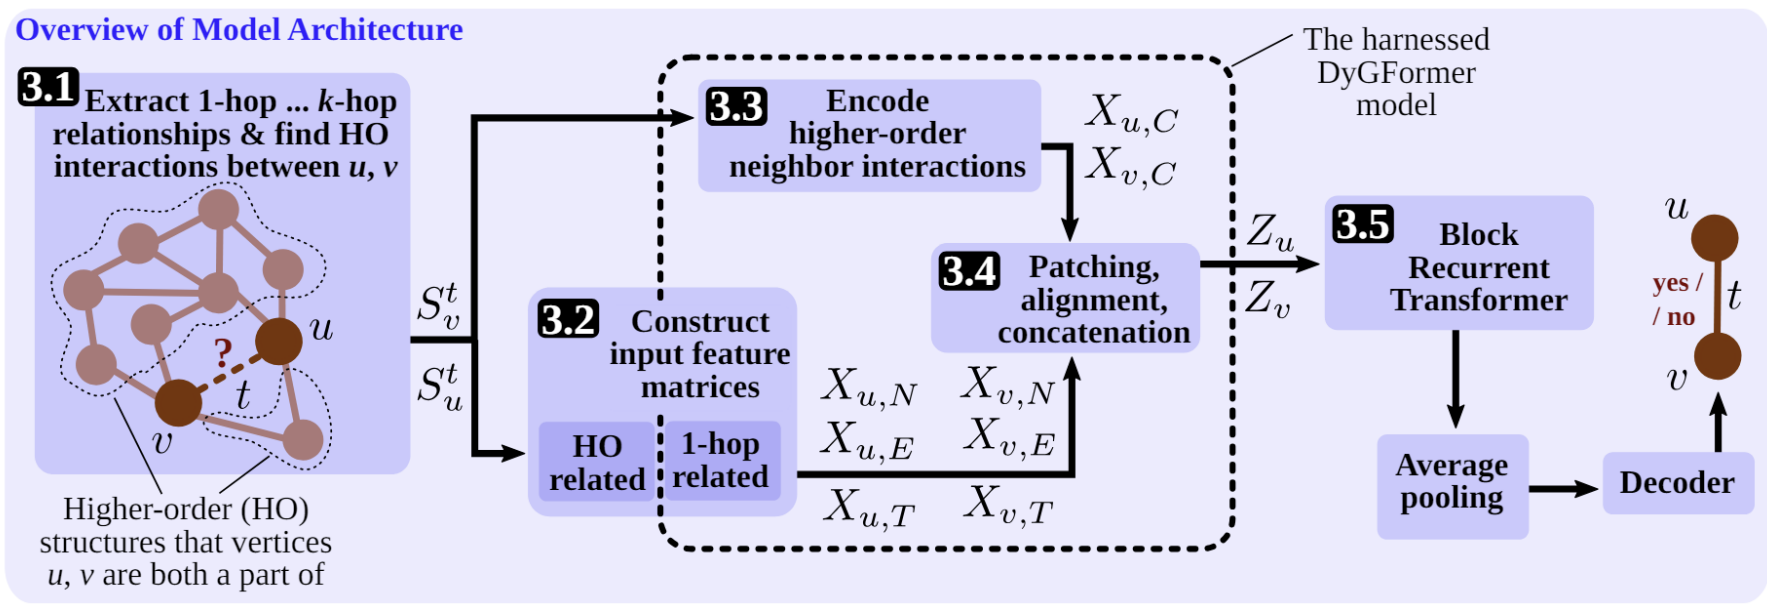
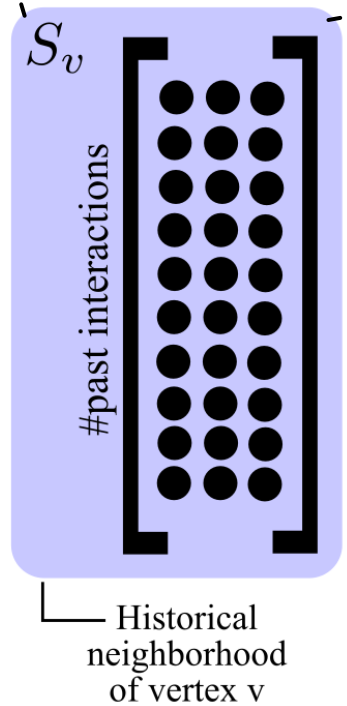
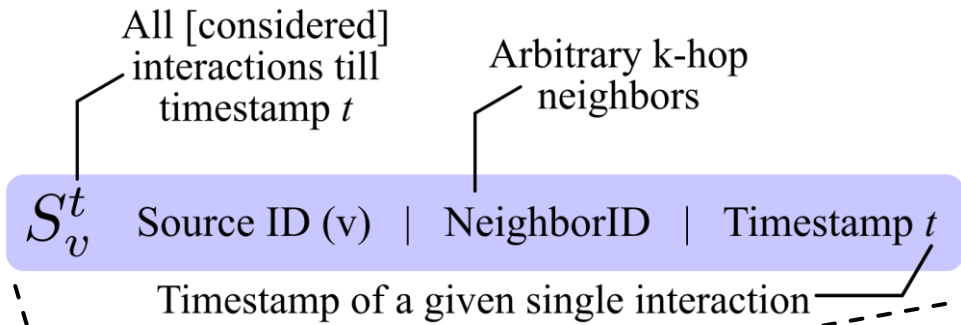
Encoding HO Structures



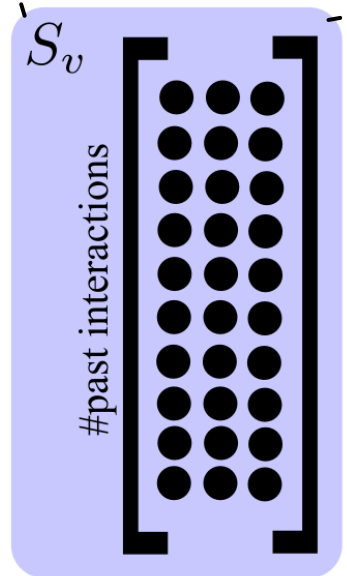
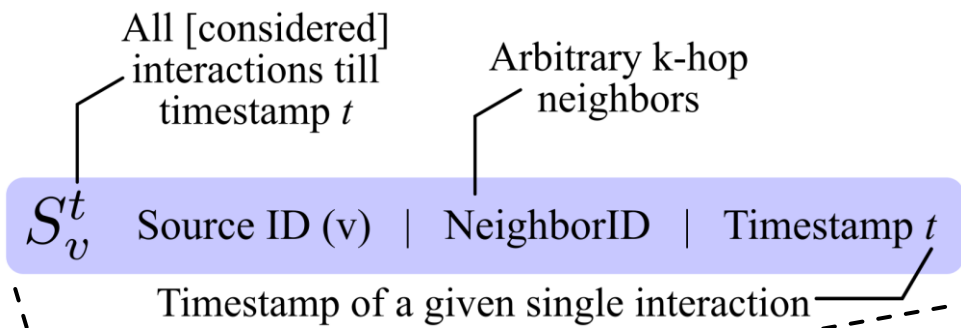
Encoding HO Structures



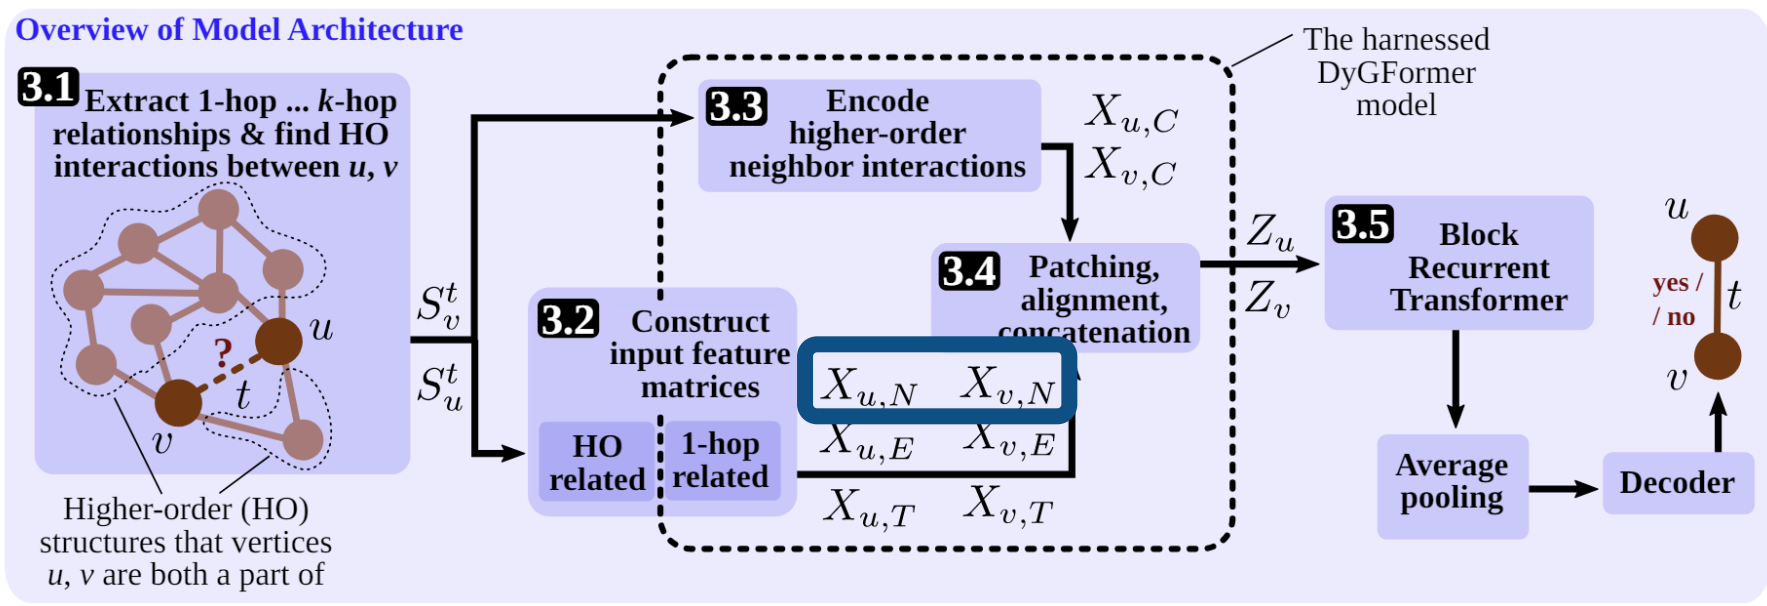
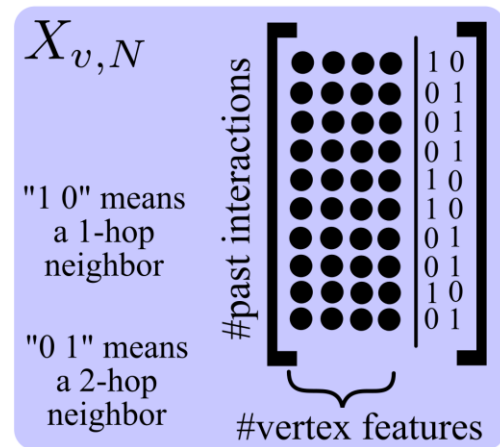
Encoding HO Structures



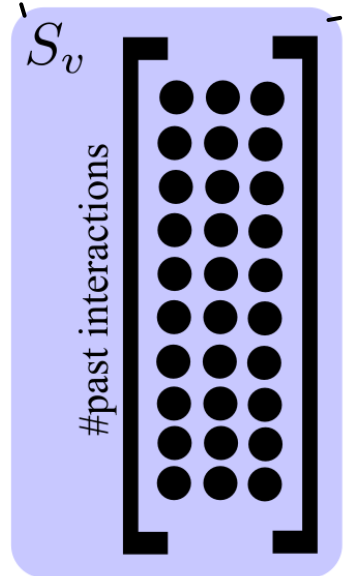
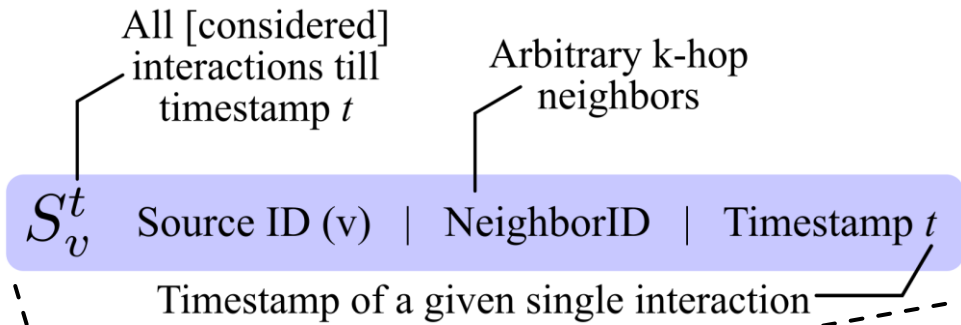
Encoding HO Structures



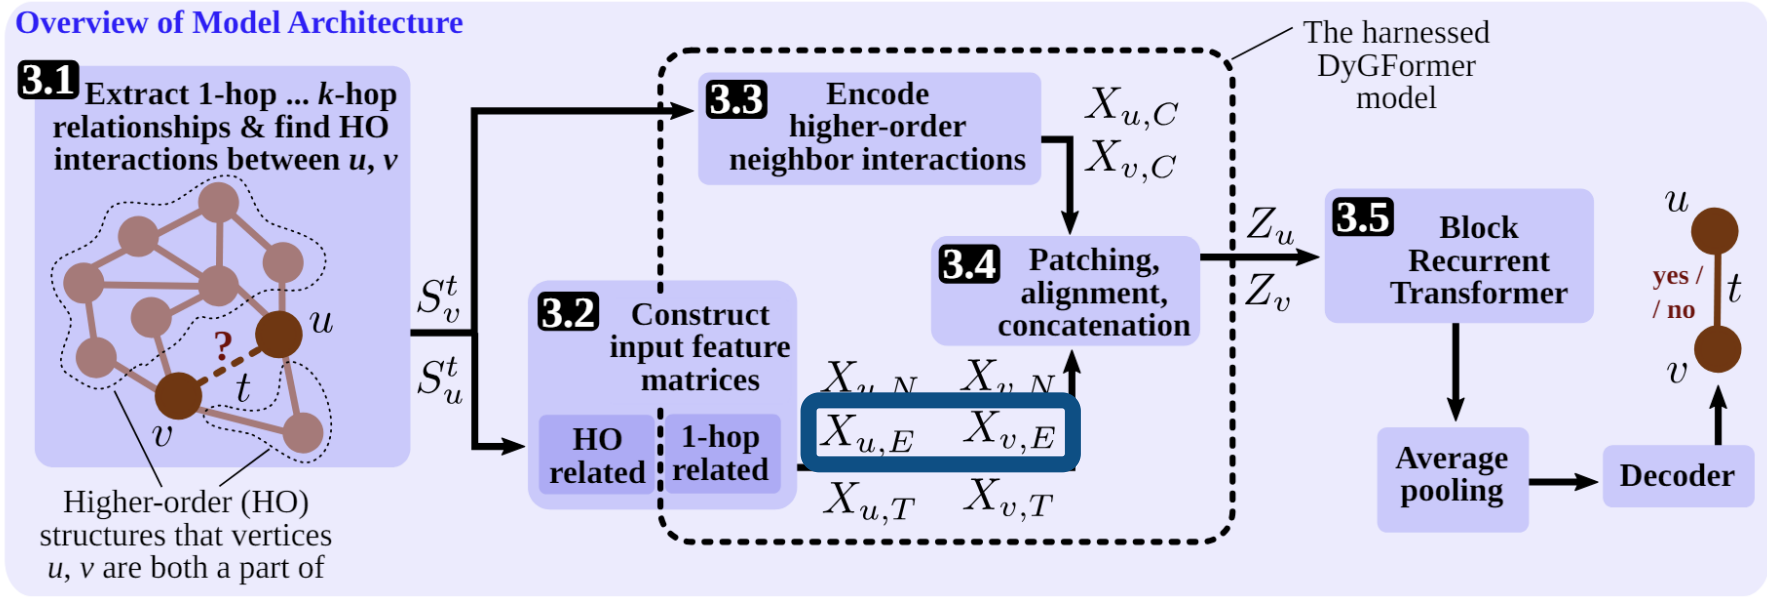
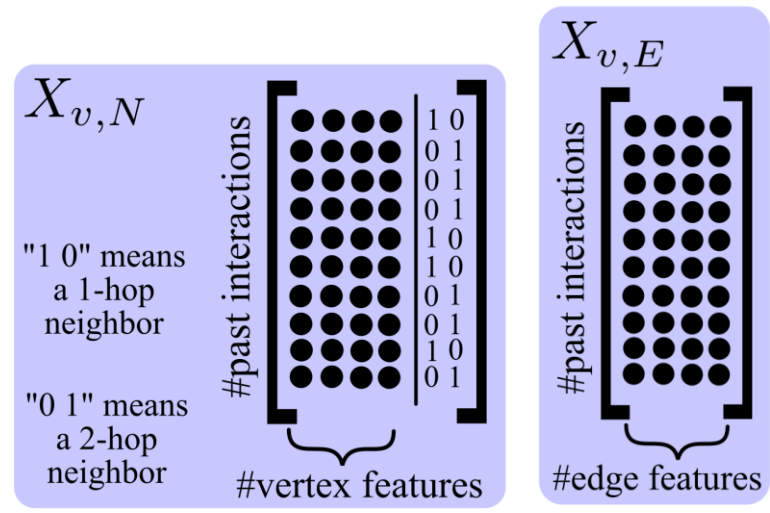
Historical neighborhood of vertex v



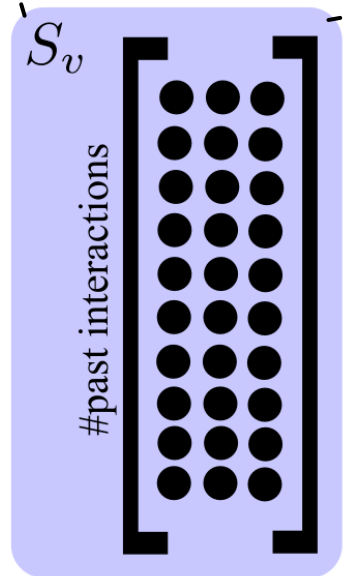
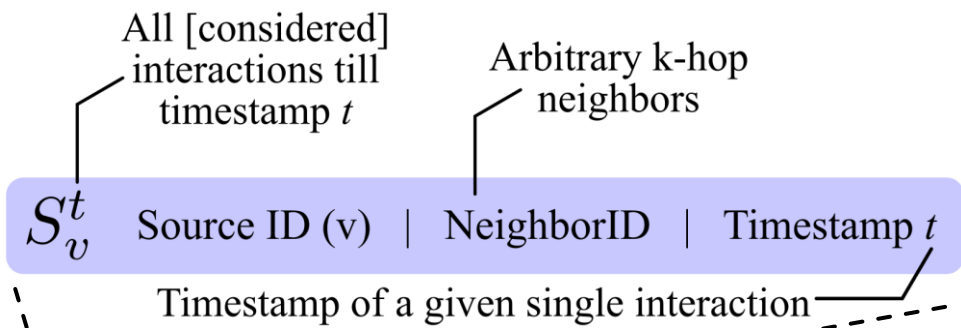
Encoding HO Structures



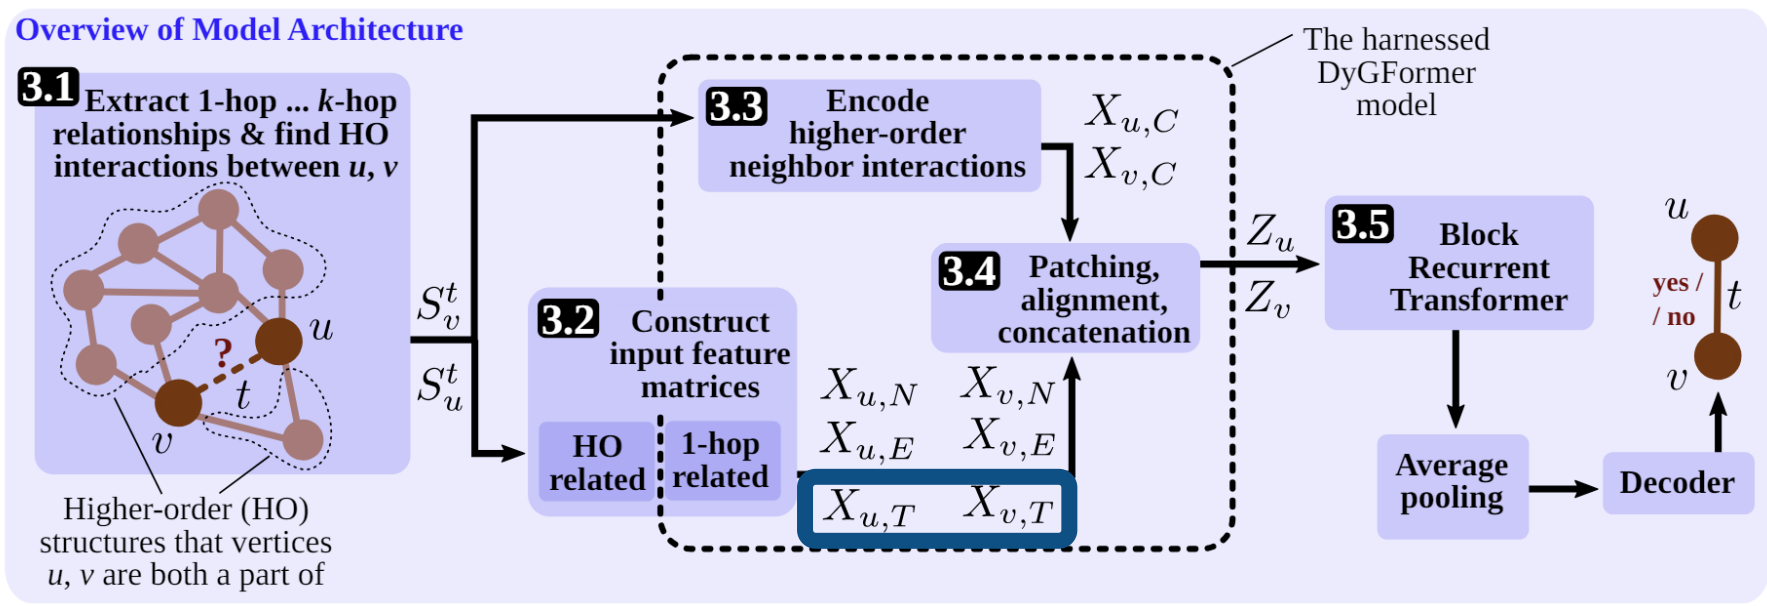
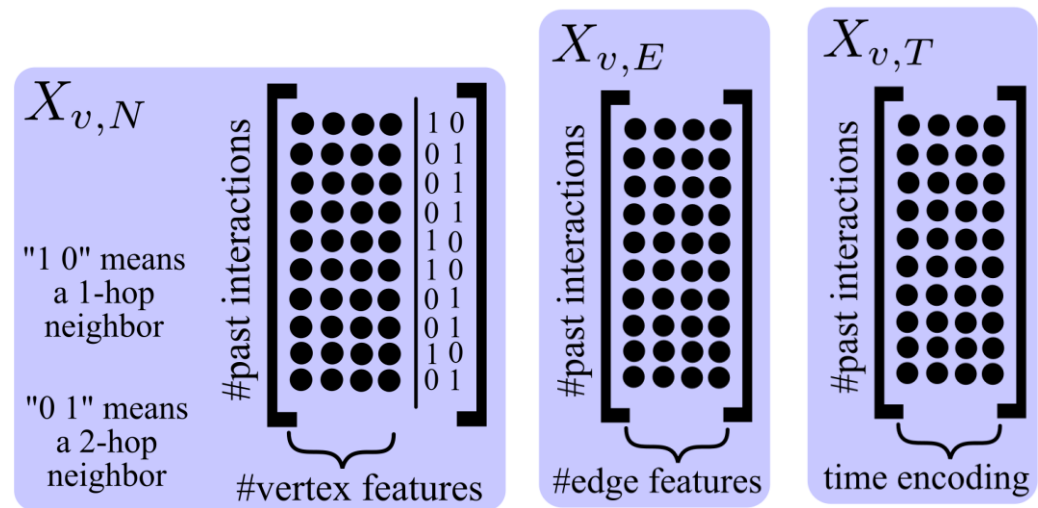
Historical neighborhood of vertex v



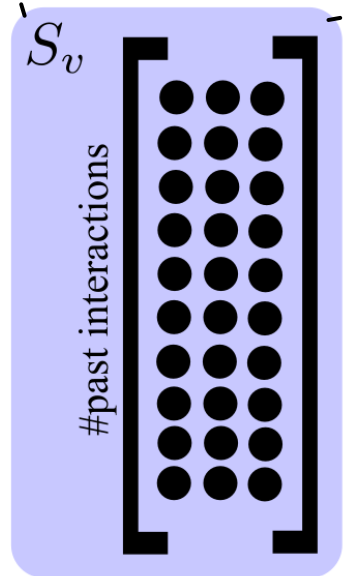
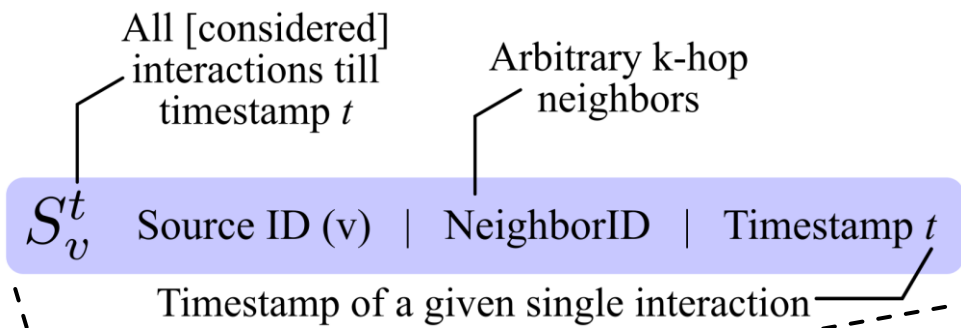
Encoding HO Structures



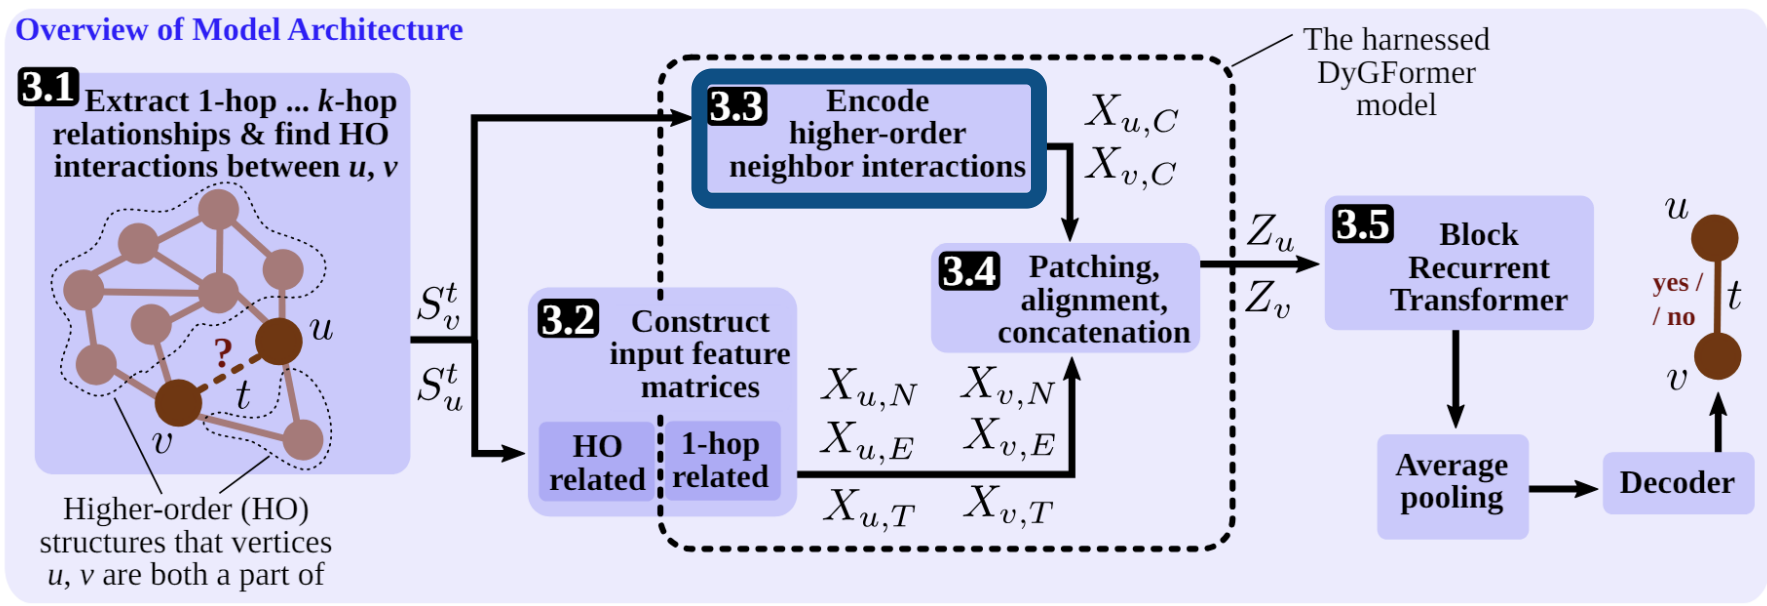
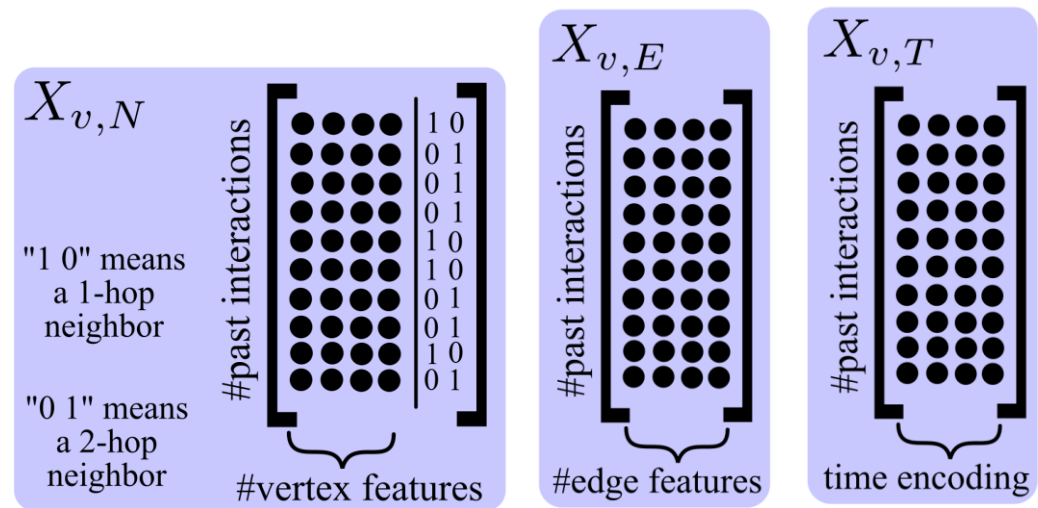
Historical neighborhood of vertex v



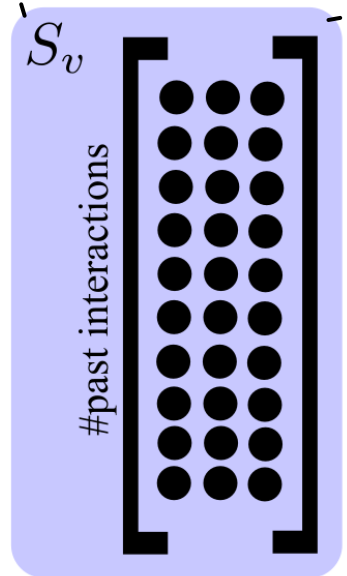
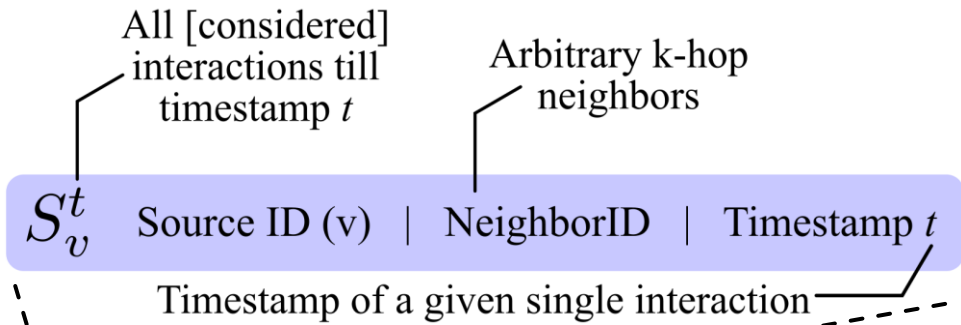
Encoding HO Structures



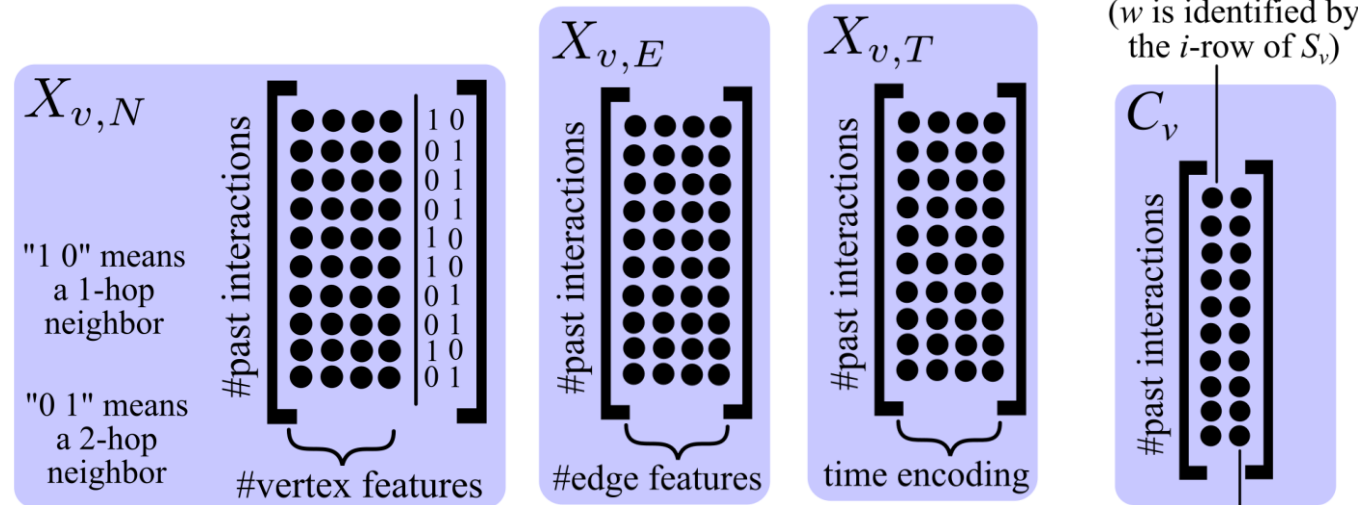
Historical neighborhood of vertex v



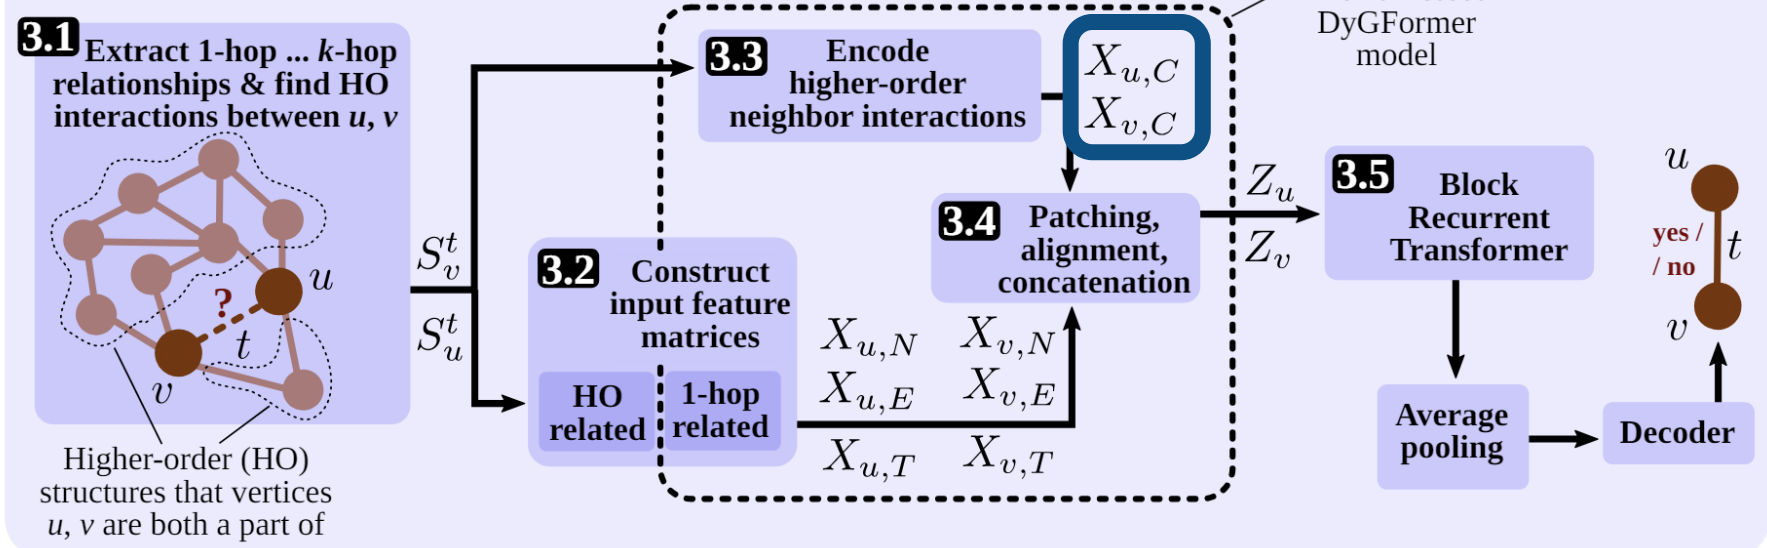
Encoding HO Structures



Historical neighborhood of vertex v

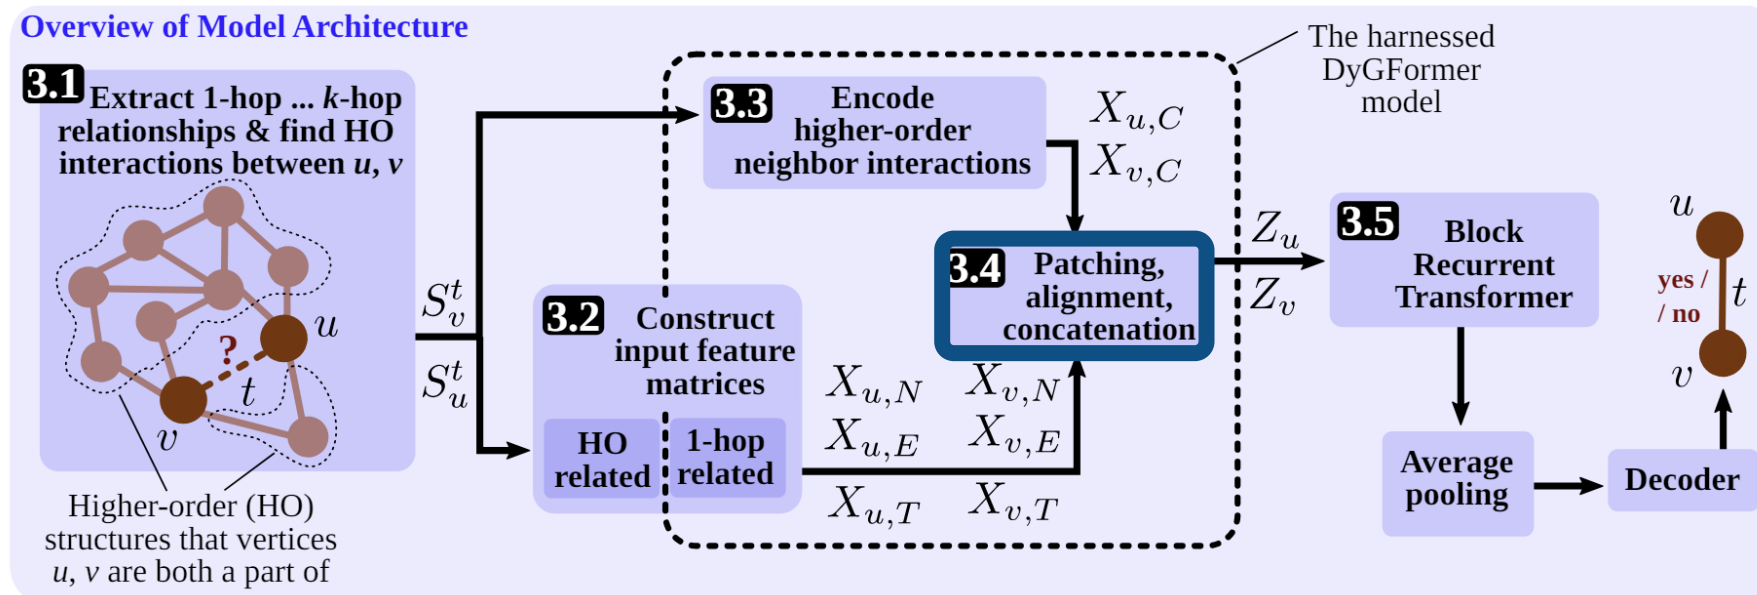


Overview of Model Architecture

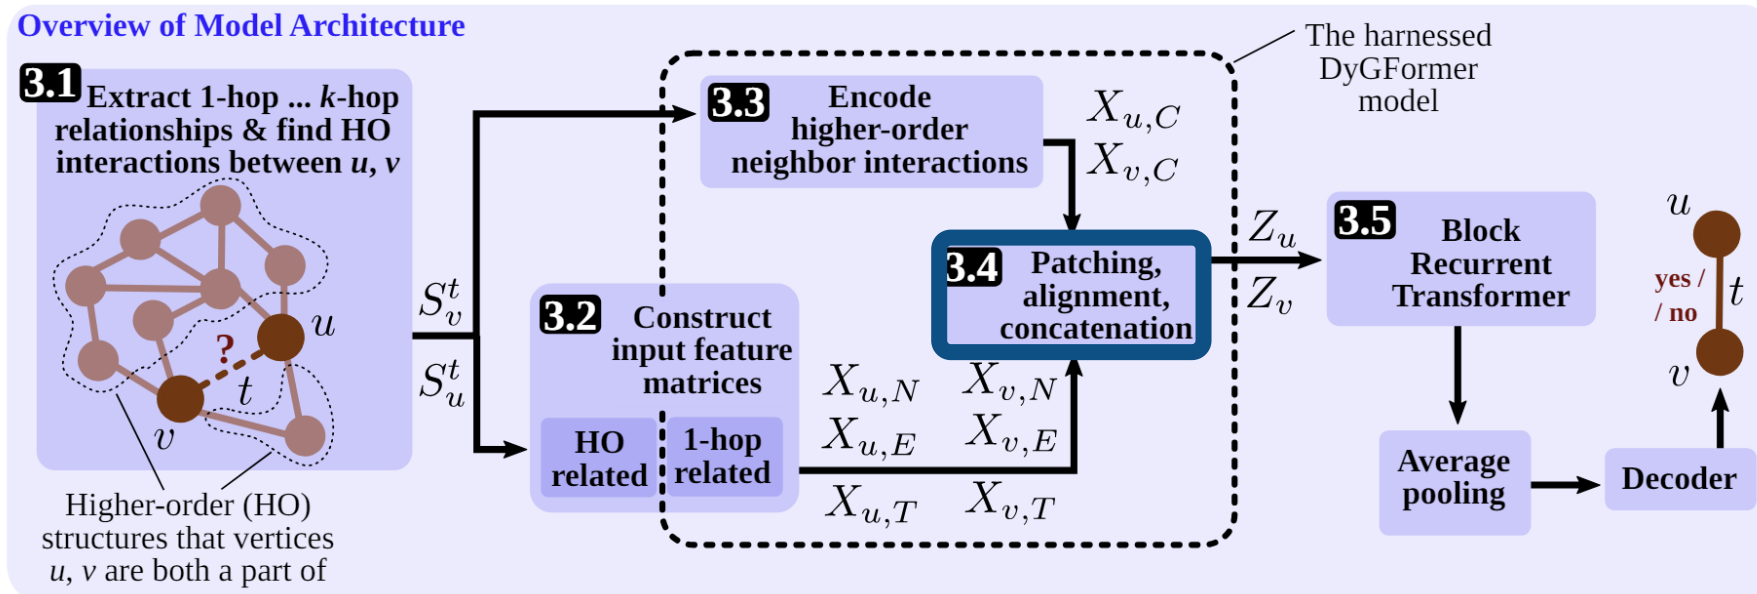
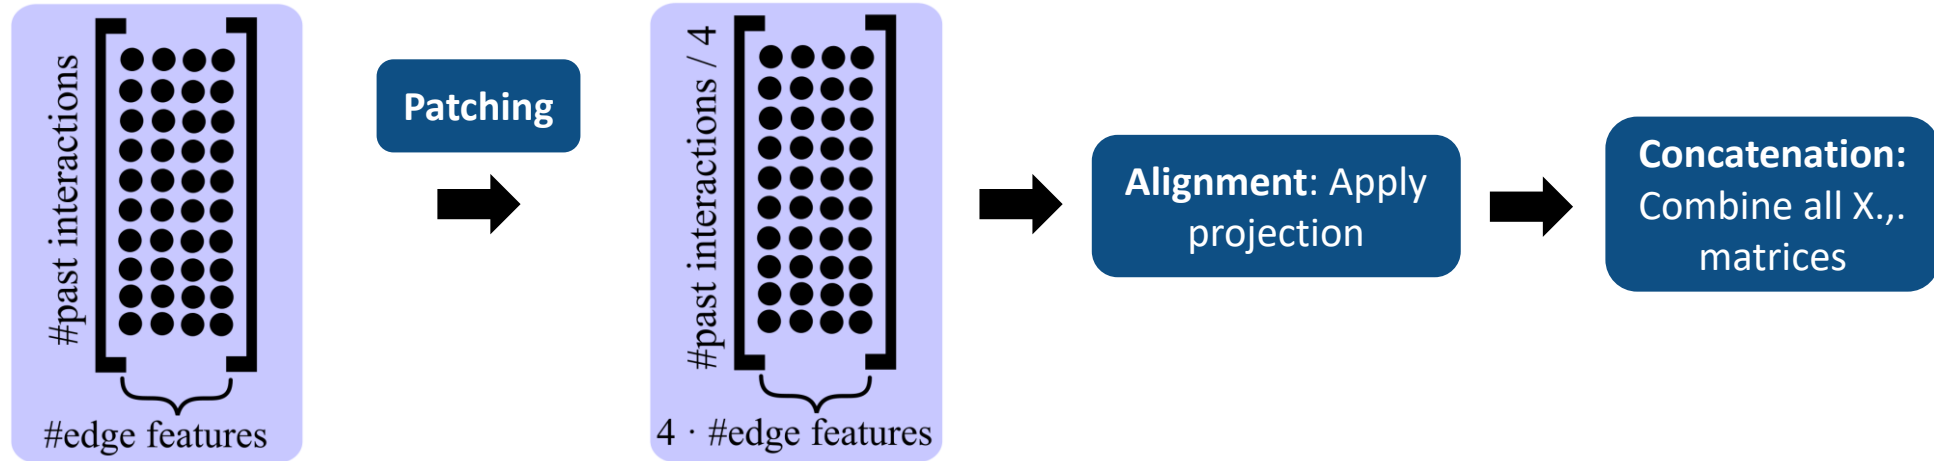


$$X_{v,C} = MLP(C_v)$$

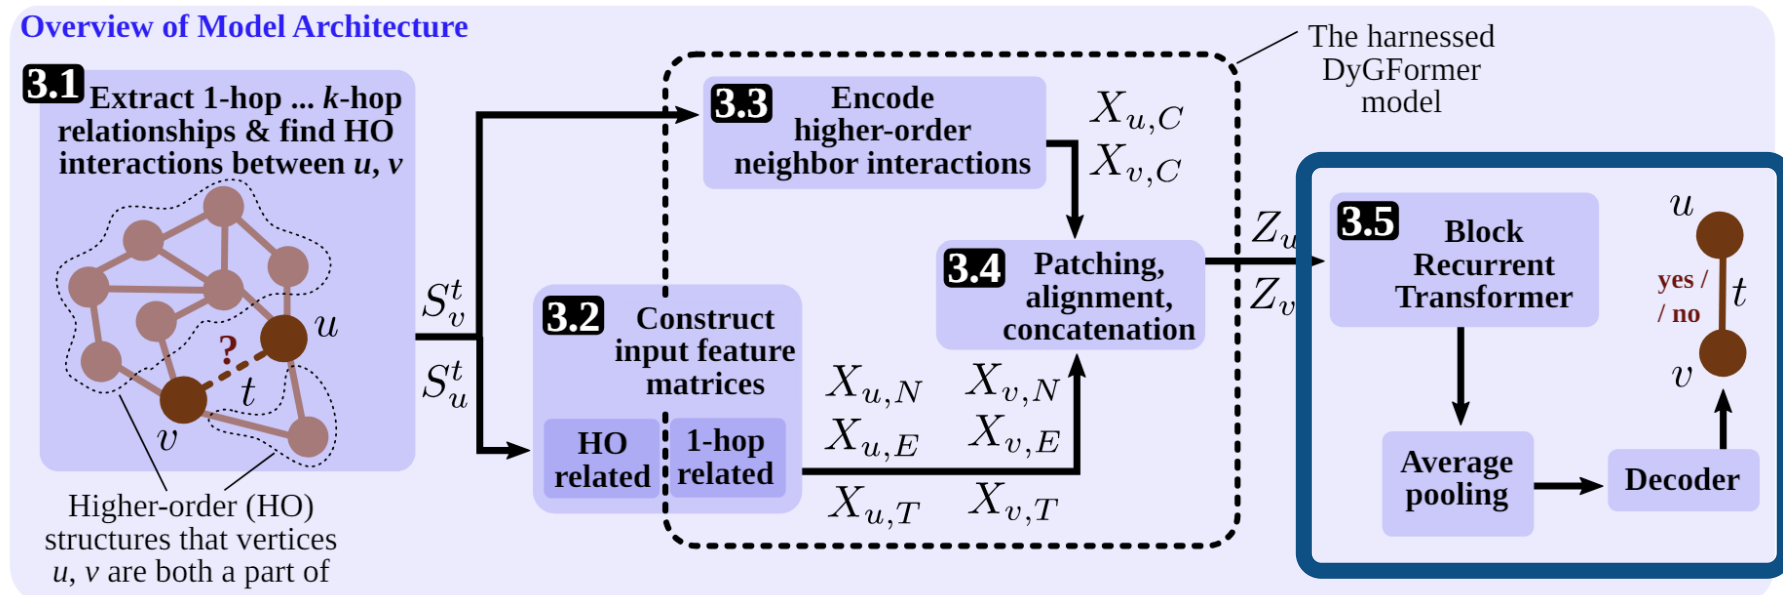
Encoding HO Structures



Encoding HO Structures

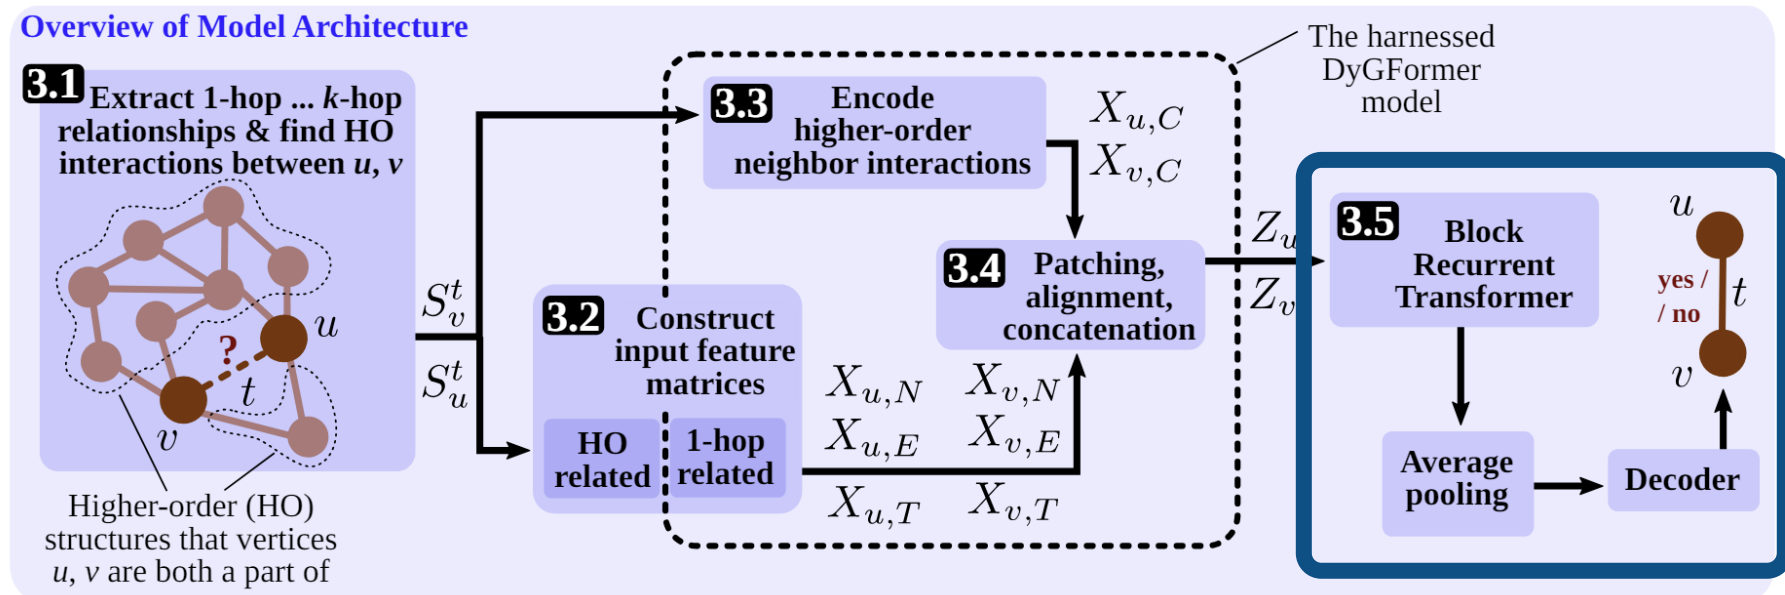


Making Predictions



Making Predictions

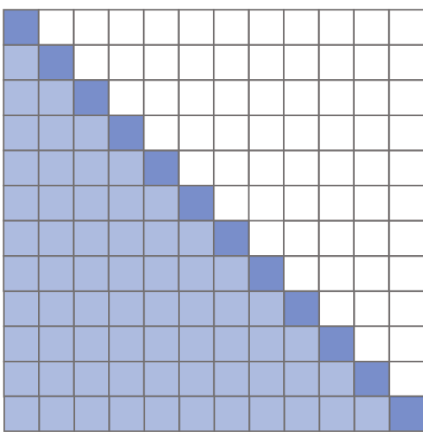
We use the module as a black box, any scheme in the domain of [efficient] Transformers could be applied



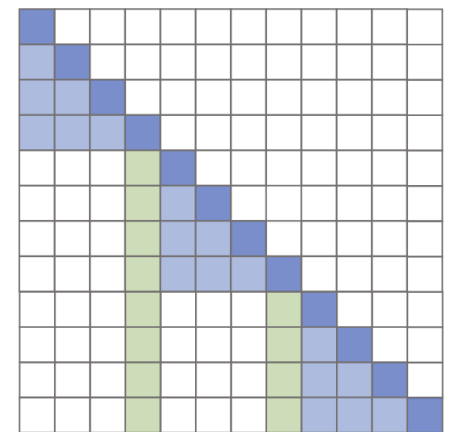
[1] Tay et al. Efficient Transformers: A Survey. 2020

Making Predictions

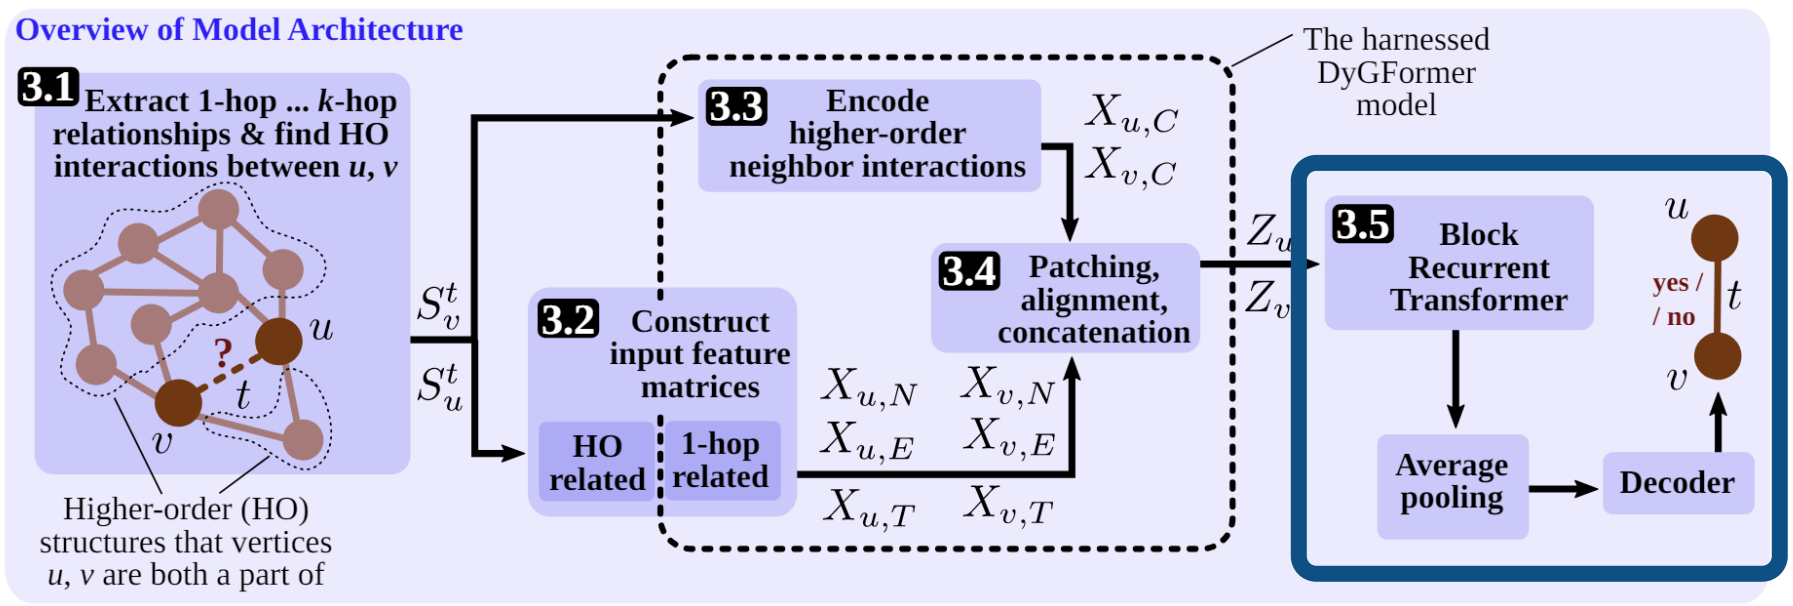
Attention matrix [1]



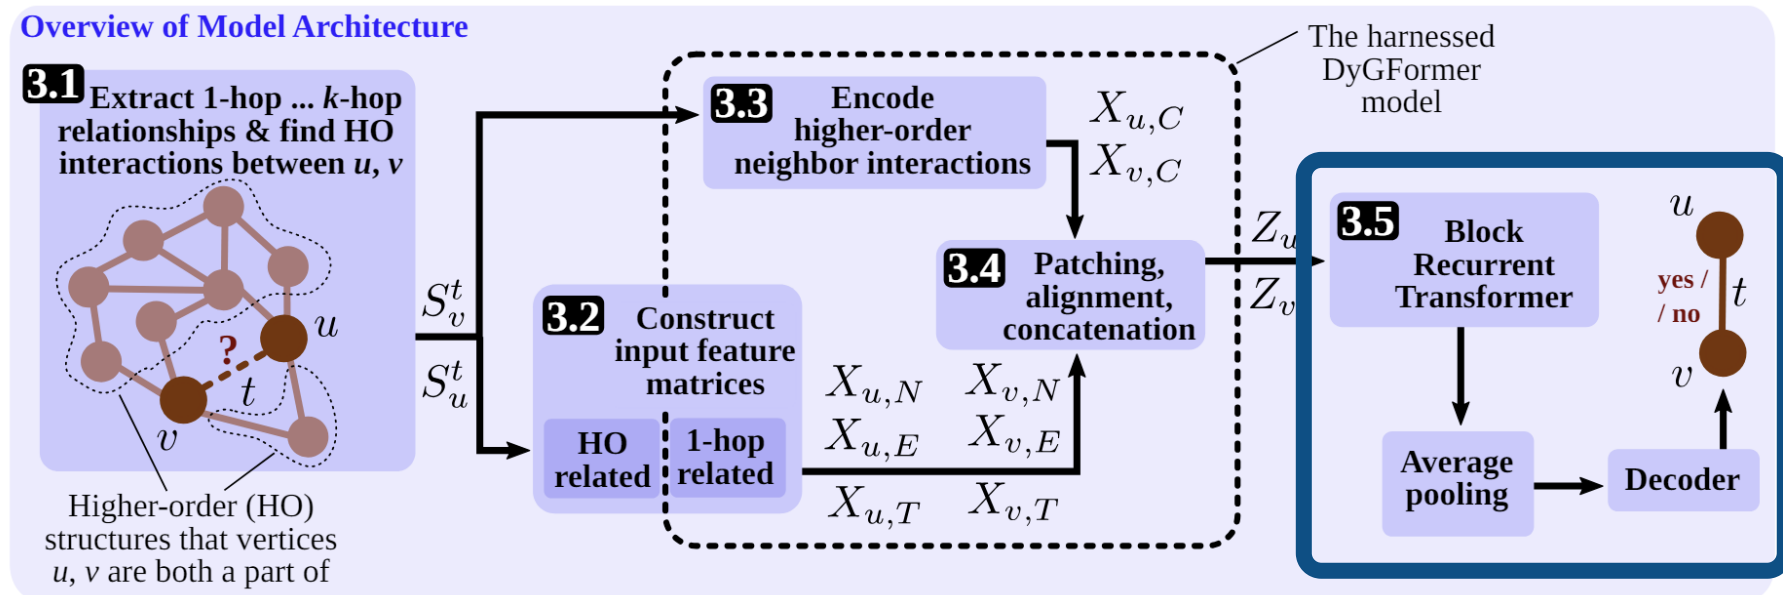
Sparse attention matrix [1]



We use the module as a black box, any scheme in the domain of [efficient] Transformers could be applied

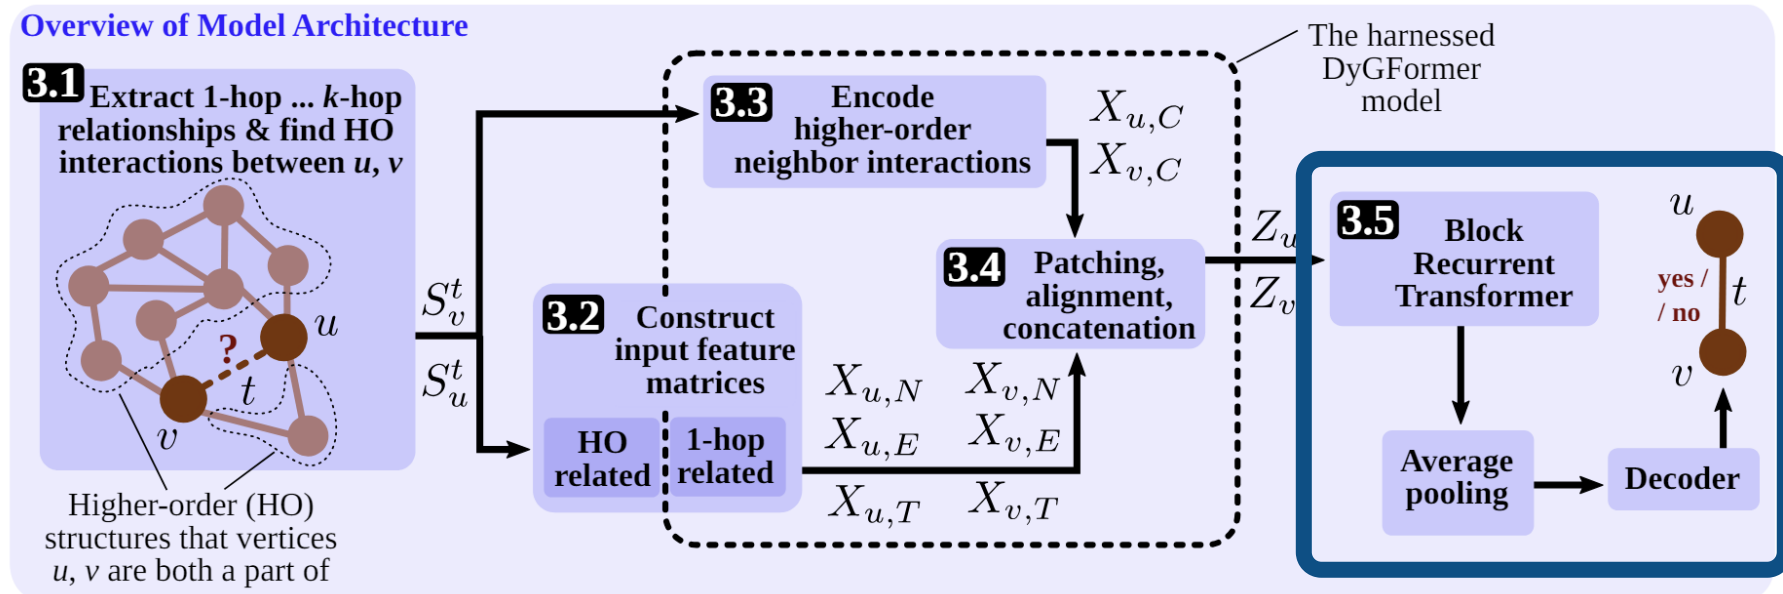


Making Predictions



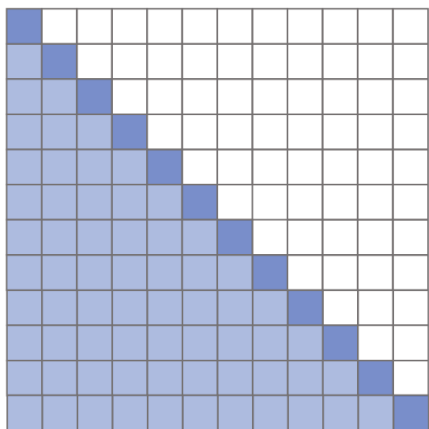
Making Predictions

Scheme harnessed:
Block-Recurrent Transformer [1]



Making Predictions

Standard attention

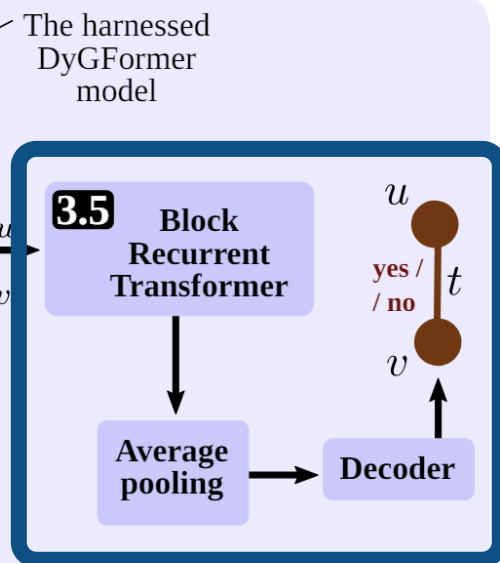
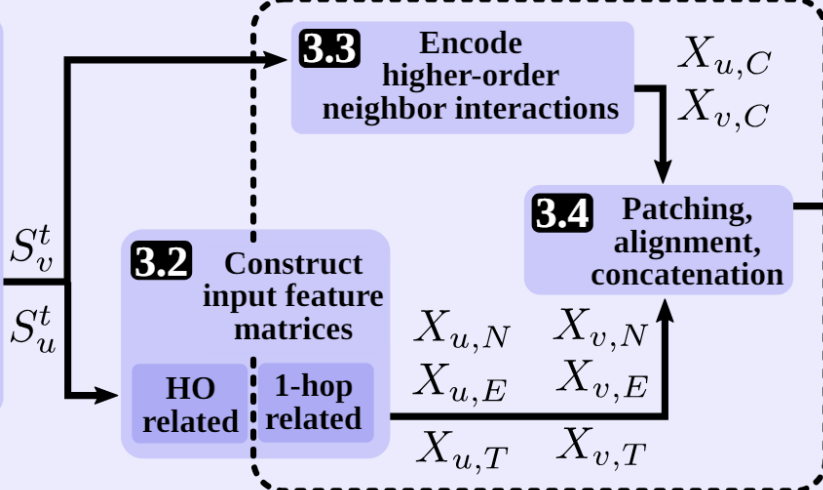
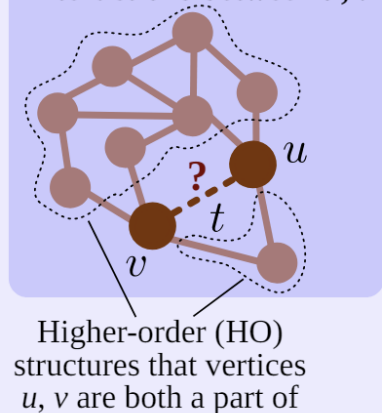


Large storage requirements when #tokens is growing

Scheme harnessed:
Block-Recurrent Transformer [1]

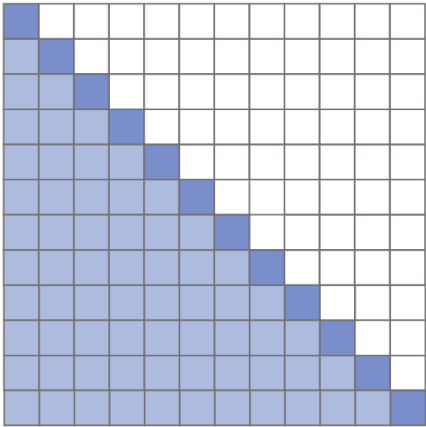
Overview of Model Architecture

3.1 Extract 1-hop ... k -hop relationships & find HO interactions between u, v



Making Predictions

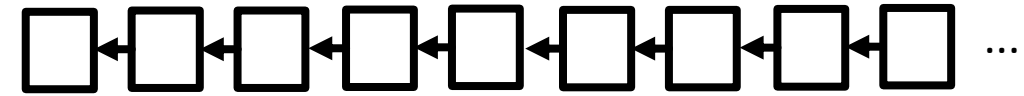
Standard attention



Large storage requirements when #tokens is growing

Scheme harnessed:
Block-Recurrent Transformer [1]

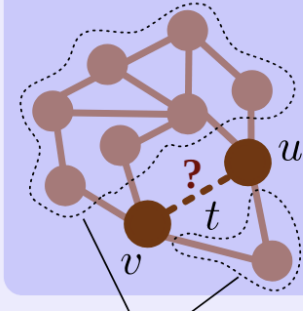
Standard RNN



Low accuracy when #tokens is growing

Overview of Model Architecture

3.1 Extract 1-hop ... k -hop relationships & find HO interactions between u, v



Higher-order (HO) structures that vertices u, v are both a part of

S_v^t
 S_u^t

3.2 Construct input feature matrices

HO related	1-hop related
$X_{u,N}$	$X_{v,N}$
$X_{u,E}$	$X_{v,E}$
$X_{u,T}$	$X_{v,T}$

3.3 Encode higher-order neighbor interactions

$X_{u,C}$
 $X_{v,C}$

3.4 Patching, alignment, concatenation

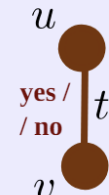
Z_u
 Z_v

The harnessed DyGFormer model

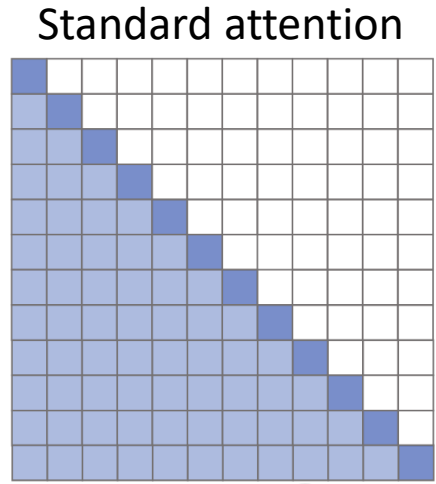
3.5 Block Recurrent Transformer

Average pooling

Decoder

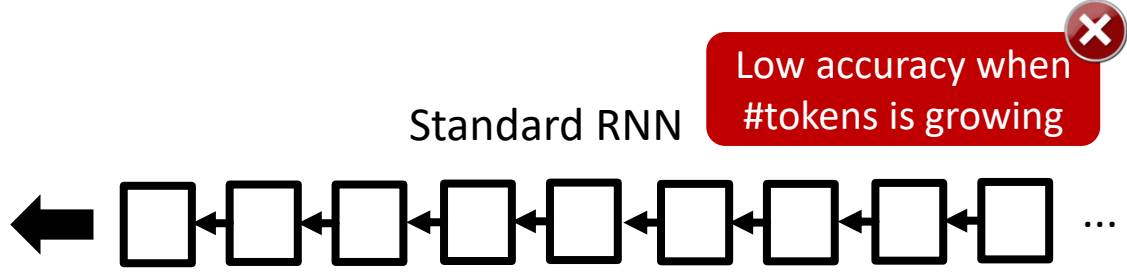
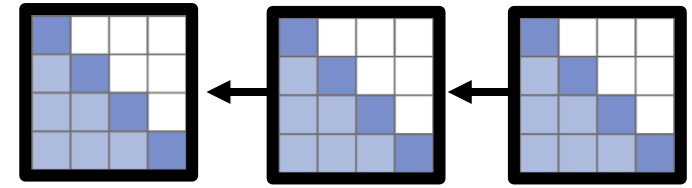


Making Predictions



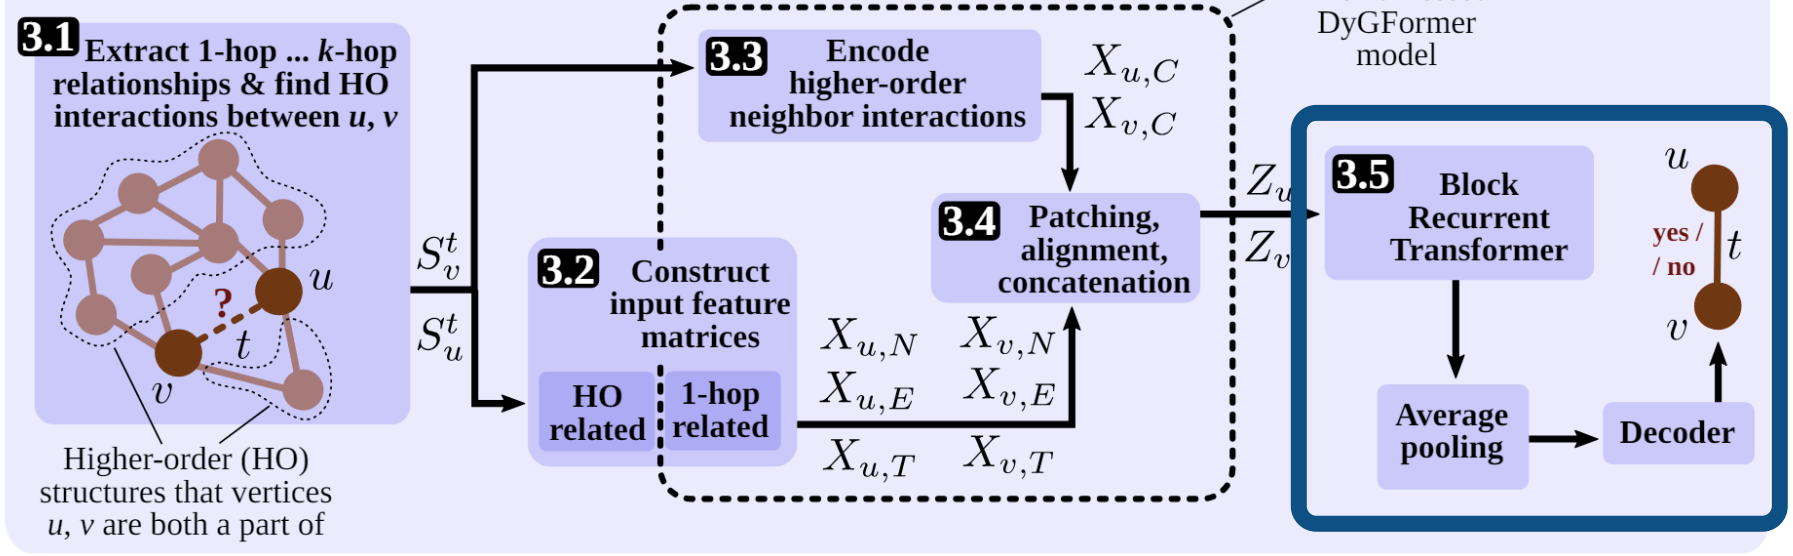
Large storage requirements when #tokens is growing

Scheme harnessed:
Block-Recurrent Transformer [1]



Low accuracy when #tokens is growing

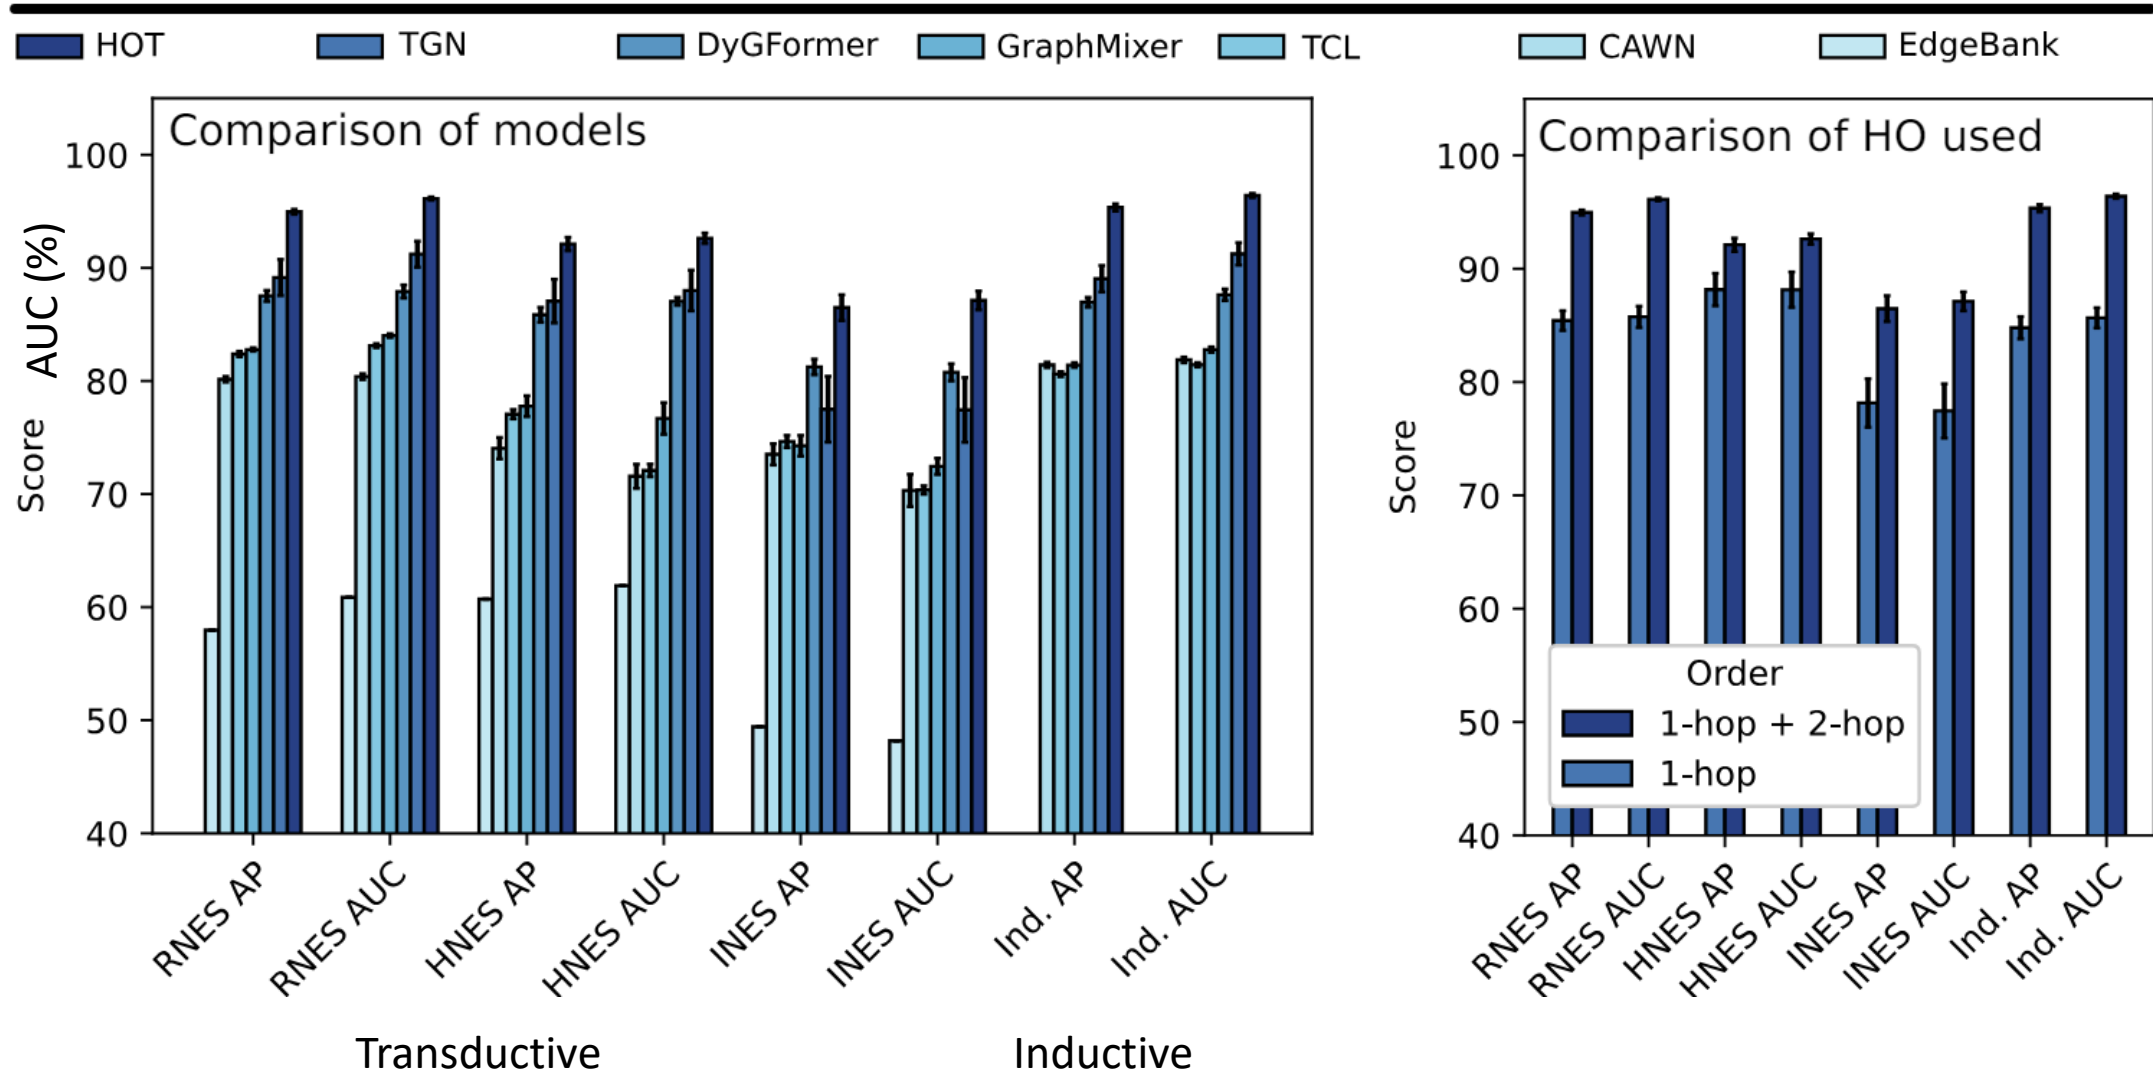
Overview of Model Architecture



HOT Evaluation

HOT successfully leverages 2-hop interactions to make its predictions more accurate

The MOOC graph dataset

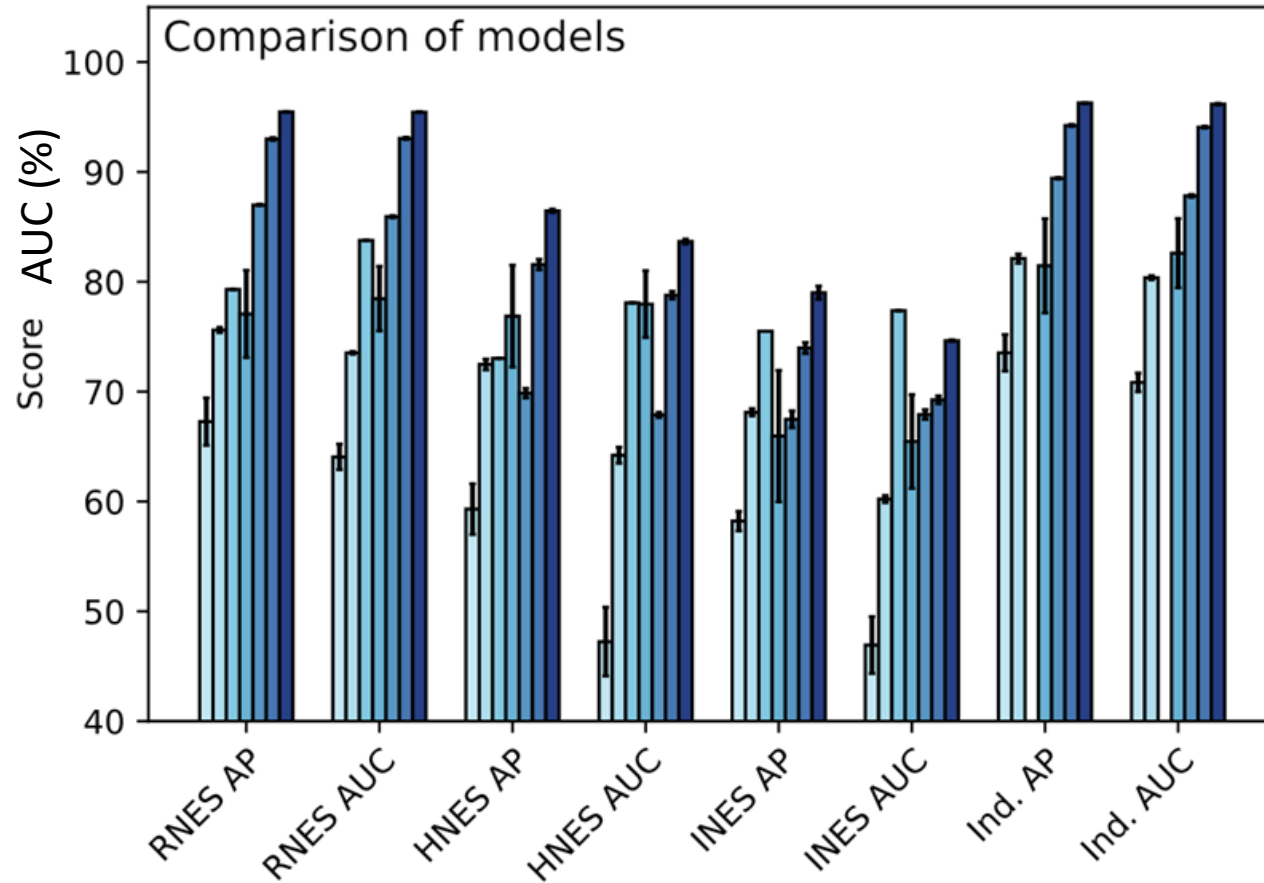


HOT Evaluation

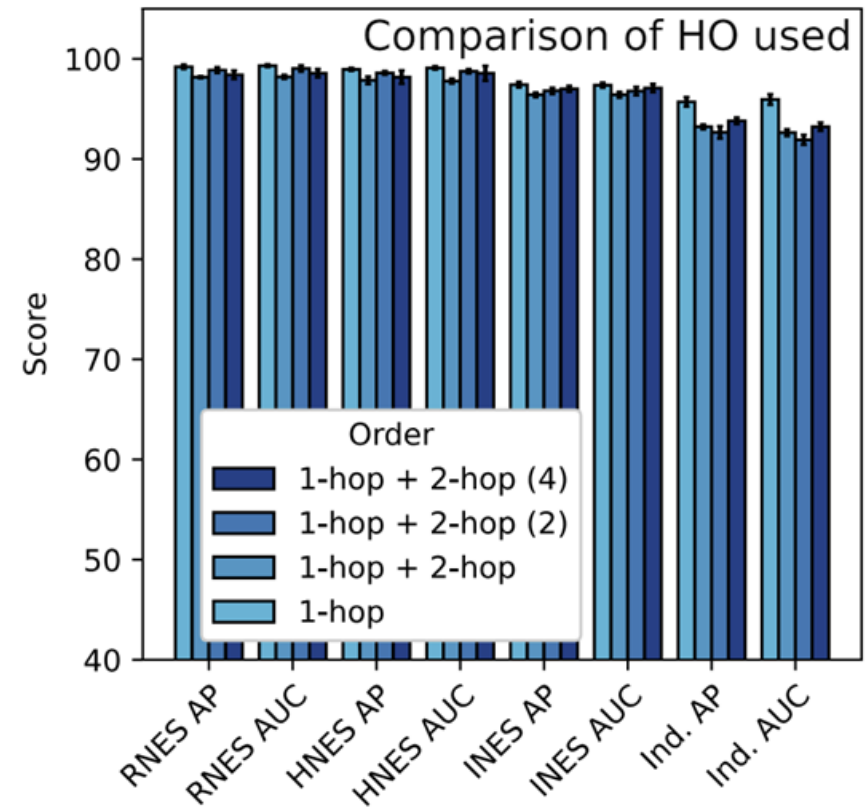
Here, the additional HO graph structural information does not seem to be adding significant amount of value to the model in this case.

The CanParl graph dataset

■ HOT
 ■ DyGFormer
 ■ CAWN
 ■ TGN
 ■ EdgeBank
 ■ GraphMixer
 ■ TCL

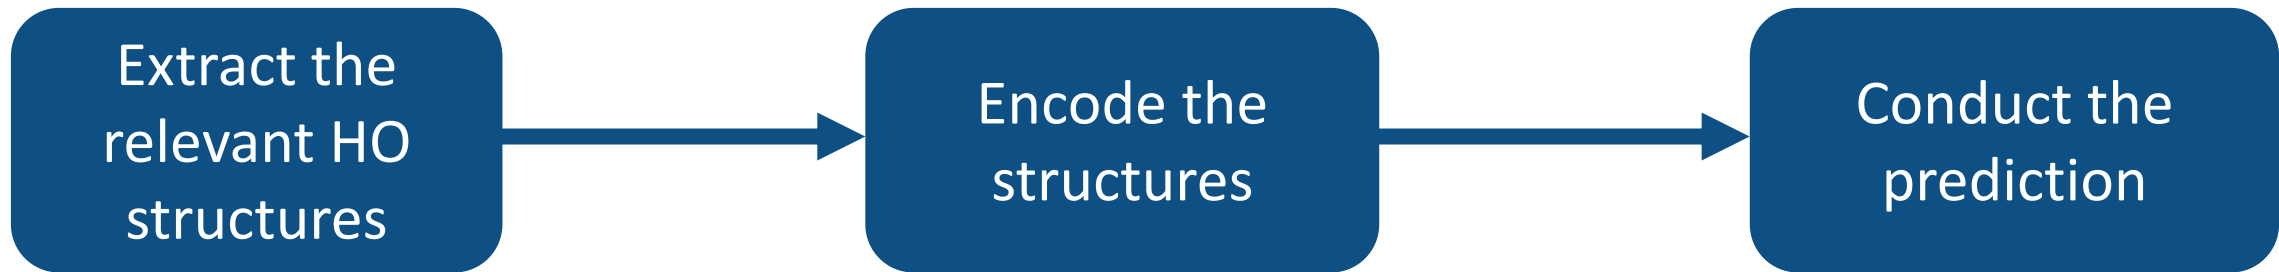


Transductive



Inductive

HO-Enhanced Pipeline for Dynamic Link Prediction



HO-Enhanced Pipeline for Dynamic Link Prediction

@ LOG'23

HOT: Higher-Order Dynamic Graph Representation Learning with Efficient Transformers

Maciej Besta^{1*} Afonso Claudino Catarino^{1*} Lukas Gianinazzi¹ Nils Blach¹
Piotr Nyczyk² Hubert Niewiadomski² Torsten Hoefler¹

¹Department of Computer Science, ETH Zurich; ²Cledar

HO-Enhanced Pipeline for Dynamic Link Prediction

