A Theoretical Comparison of Graph Neural Network Extensions

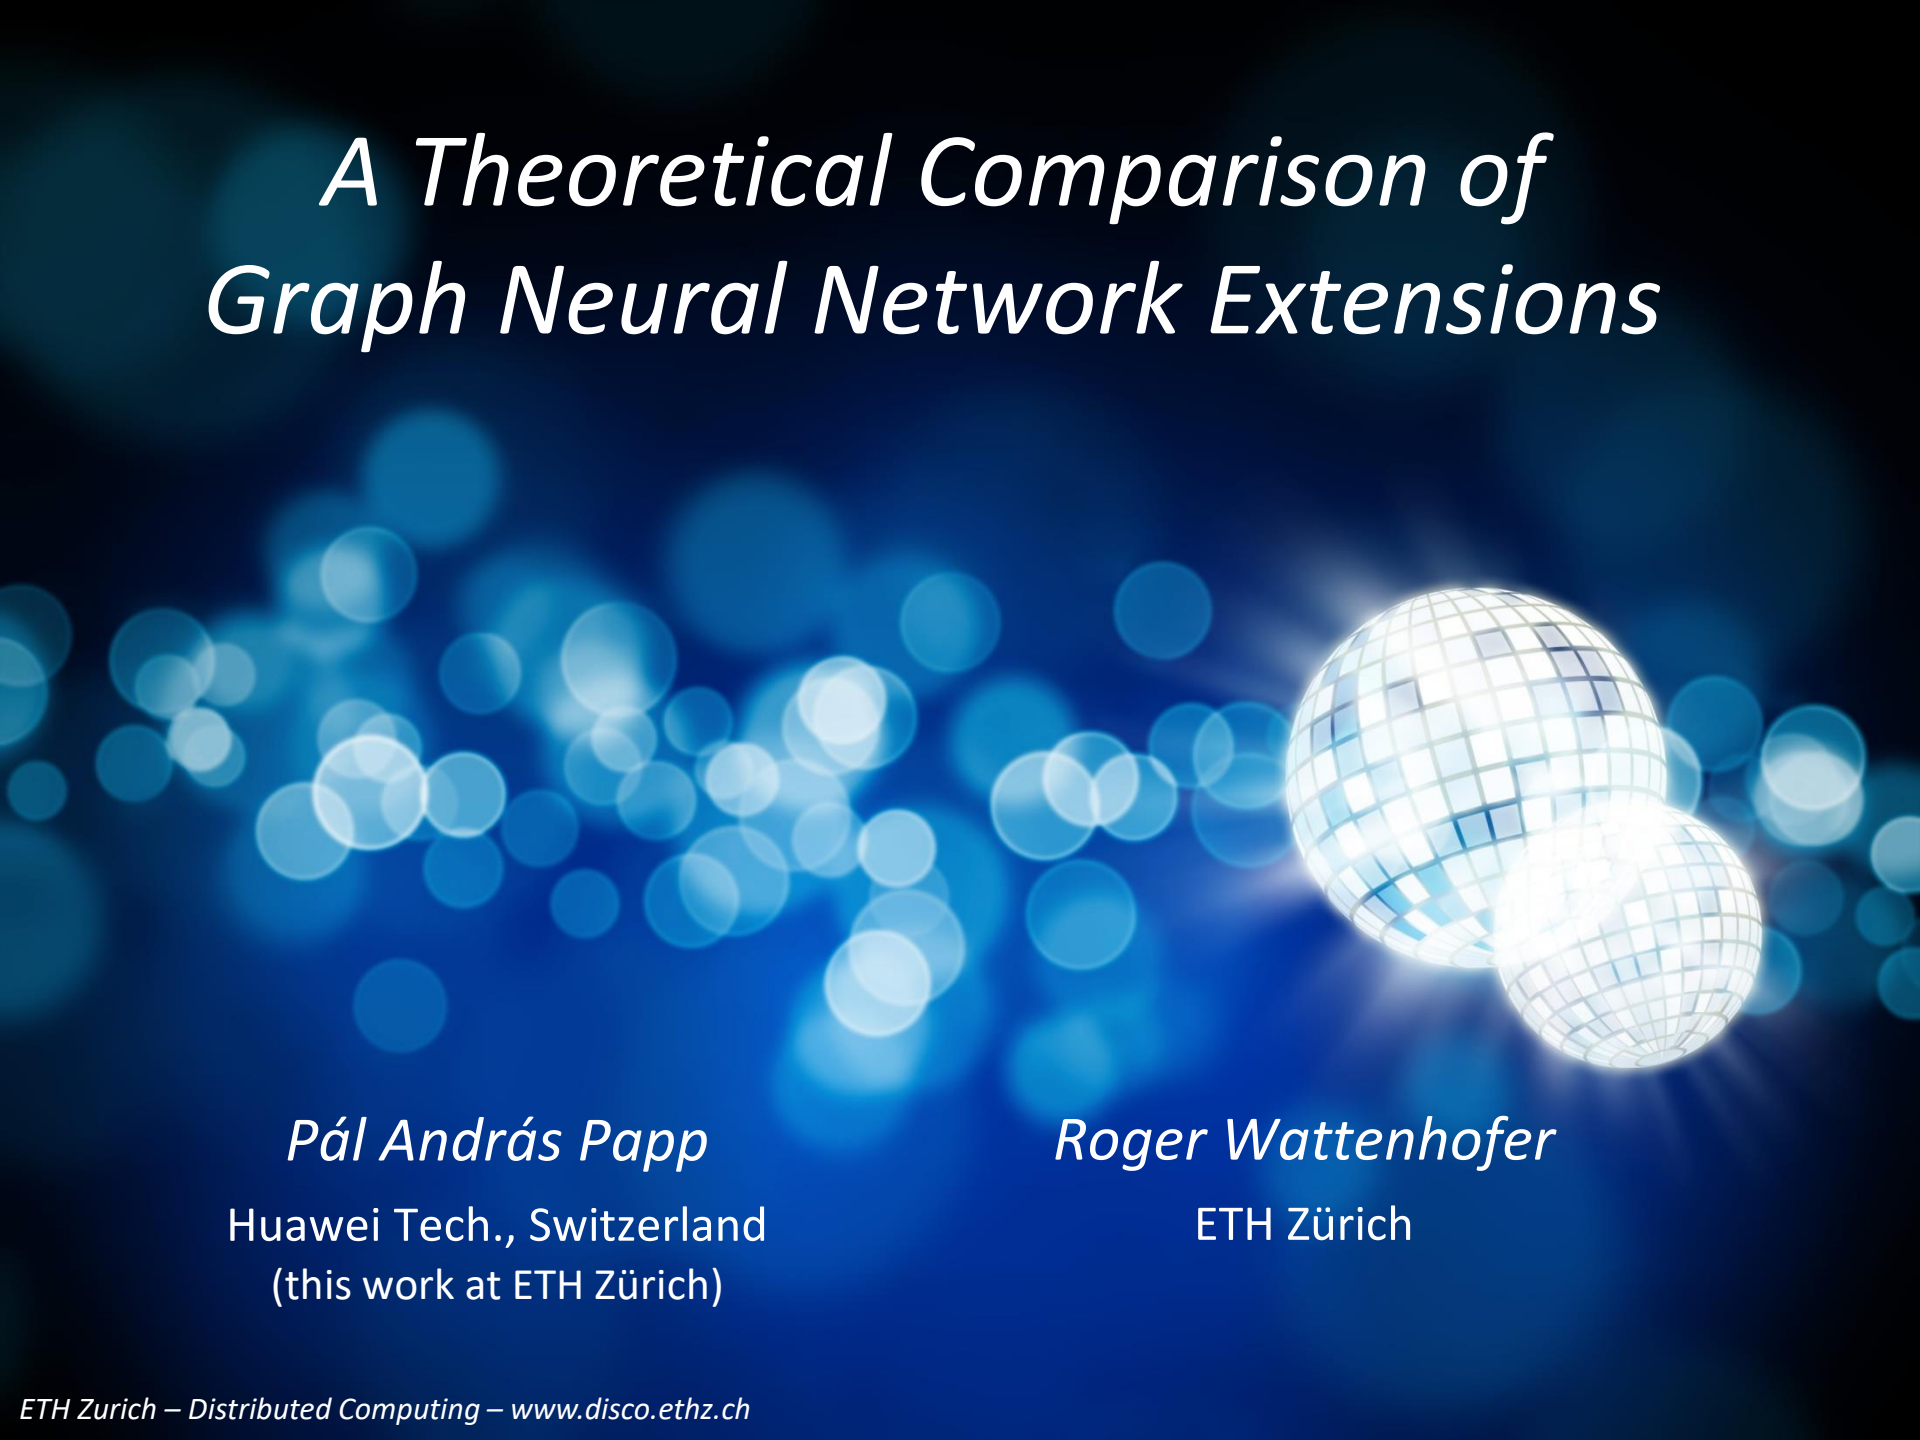

Pál András Papp

Huawei Tech., Switzerland
(this work at ETH Zürich)

Roger Wattenhofer

ETH Zürich

Graph Neural Networks

social networks

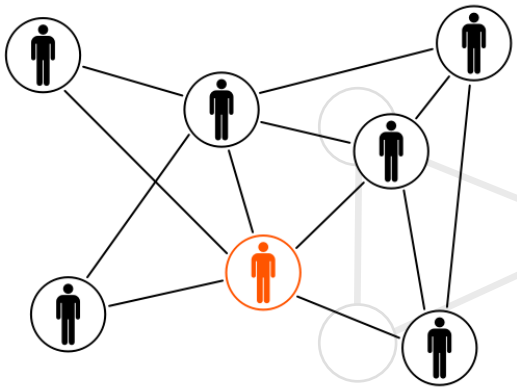

chemo-informatics

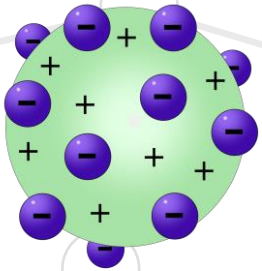

question answering systems

molecule recognition

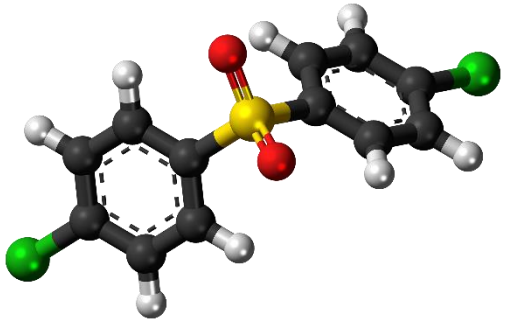

recommender systems

knowledge graphs

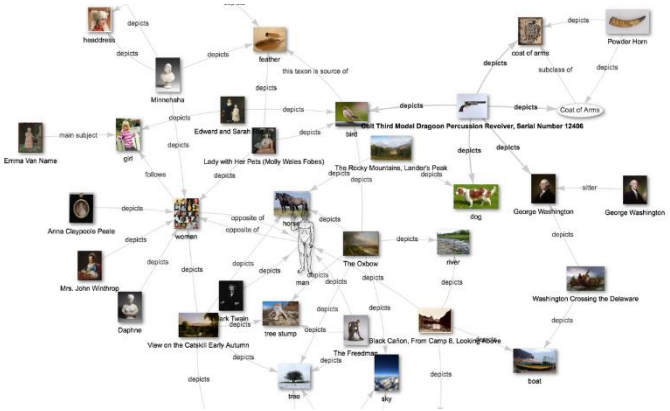

Graph Neural Networks

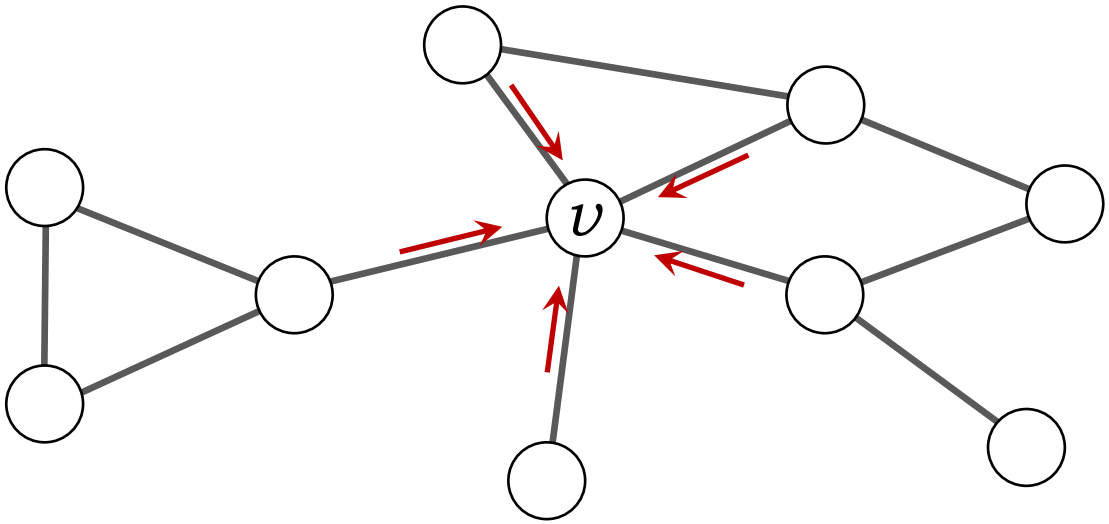

a_v = AGGREGATE ({{ h_u | $u \in N(v)$ }})

$h_v^{(t+1)}$ = UPDATE (h_v, a_v)

Graph Neural Networks

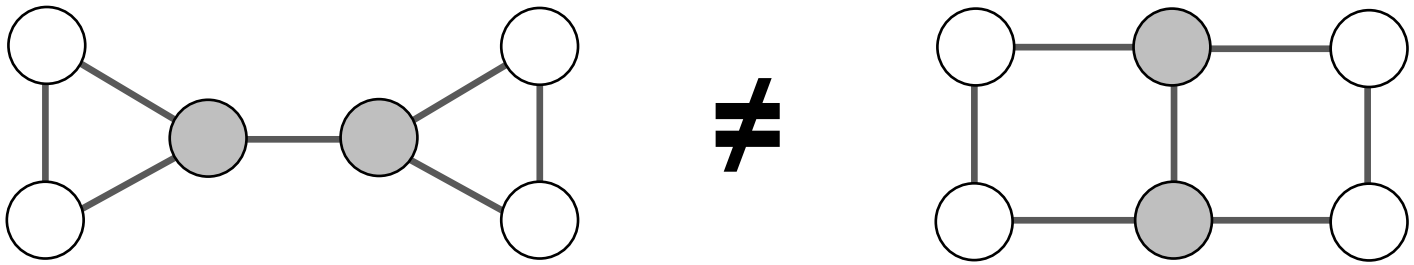

Cannot distinguish some very simple graphs!

k-WL: Higher-Order Methods

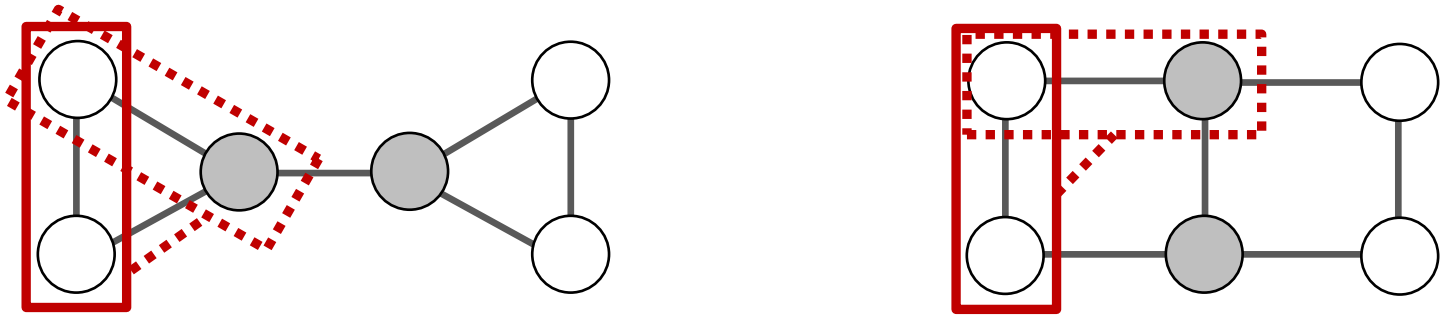

GNN on *k*-tuples of original nodes

S_k: substructures up to size k

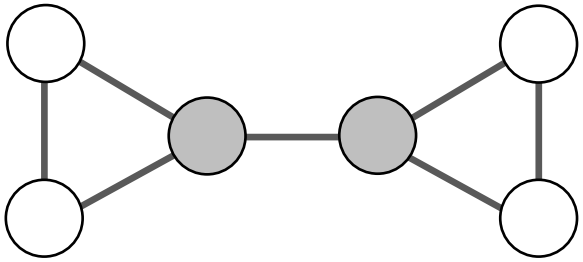

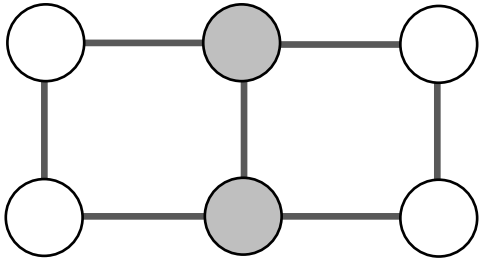

- Preprocess substructures up to size k
- Add as extra node features
- Run standard GNN

S_k: substructures up to size k

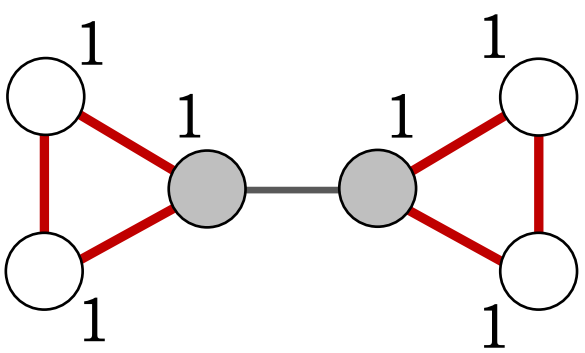

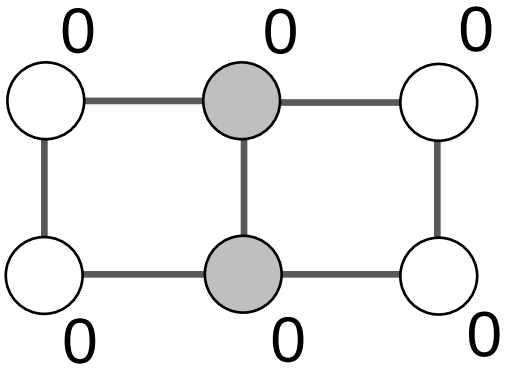

- Preprocess substructures up to size k
- Add as extra node features
- Run standard GNN

N_k: preprocessing up to k hops

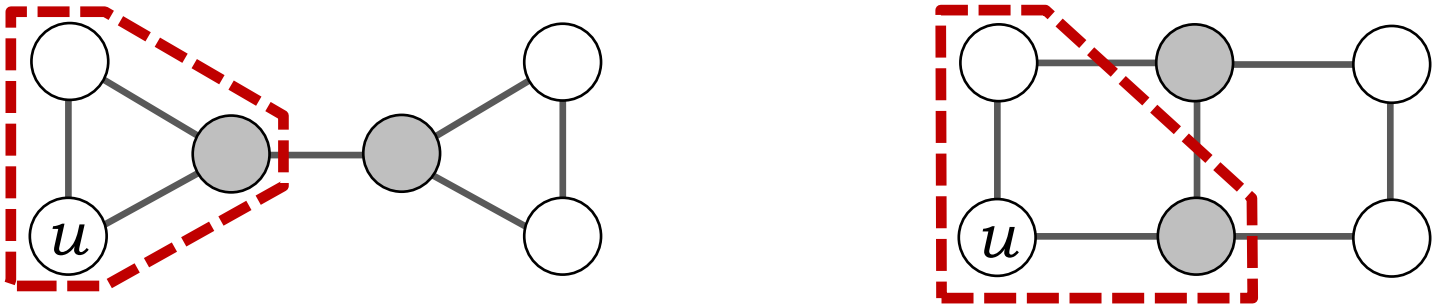

- Preprocess k-hop neighborhood

N_k: preprocessing up to k hops

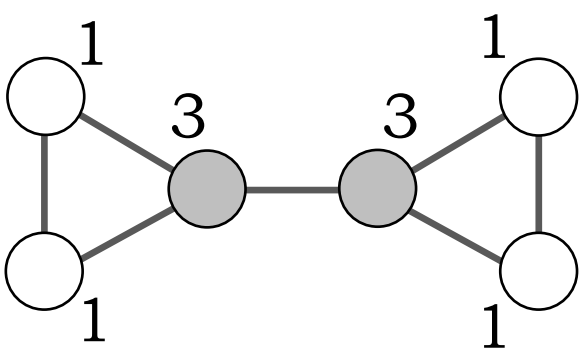

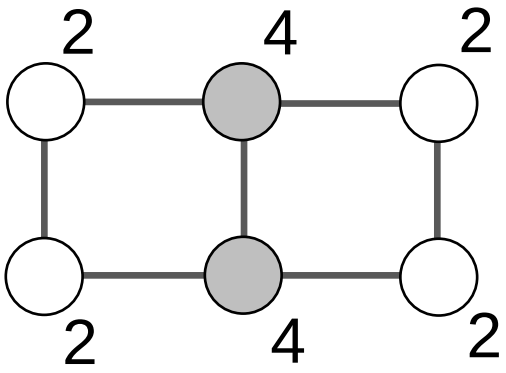

- Preprocess k-hop neighborhood
- Add extra features
- Run standard GNN

Removing some of the nodes

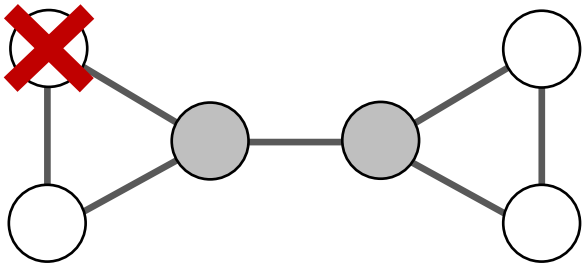

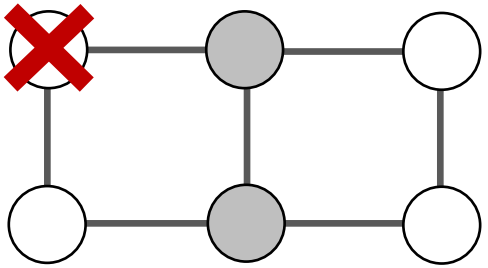

- Multiple runs of a standard GNN
- Removing some nodes in each run
- Aggregate runs in the end

Removing some of the nodes

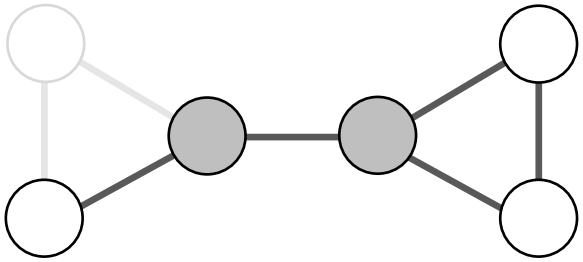

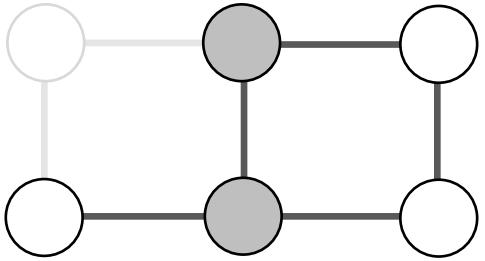

- Multiple runs of a standard GNN
- Removing some nodes in each run
- Aggregate runs in the end

M_k: marking up to k nodes

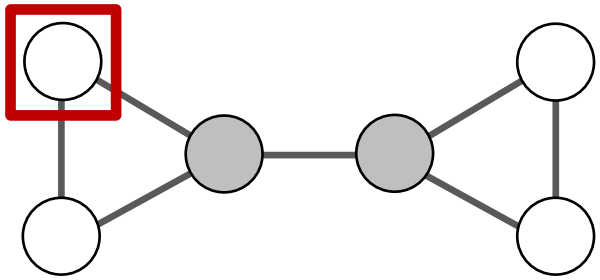

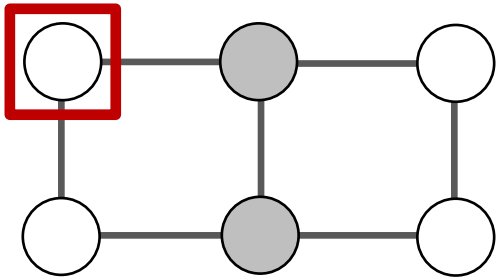

M_k: marking up to k nodes

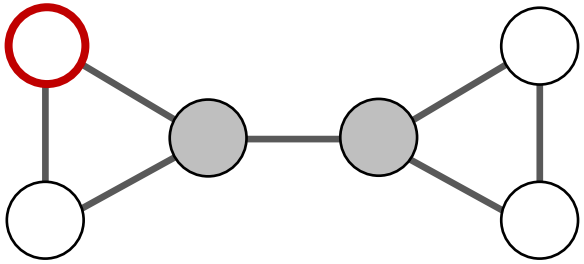

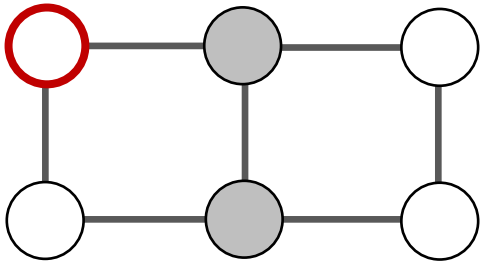

- **Mark** some nodes in each run
- Handle the marked nodes differently

$$a_v = \text{AGGR}_{marked}(\ldots) + \text{AGGR}_{unmarked}(\ldots)$$

M_k: marking up to k nodes

- **Mark** some nodes in each run
- Handle the marked nodes differently

$$a_v = \text{AGGR}_{marked}(\ldots) + \text{AGGR}_{unmarked}(\ldots)$$

Strictly more expressive than node removals!

GNN extensions

→ k-WL: k^{th} order methods

→ S_k: substructures up to size k

→ N_k: k-hop neighborhood

→ M_k: randomly mark k nodes

Comparison of expressiveness

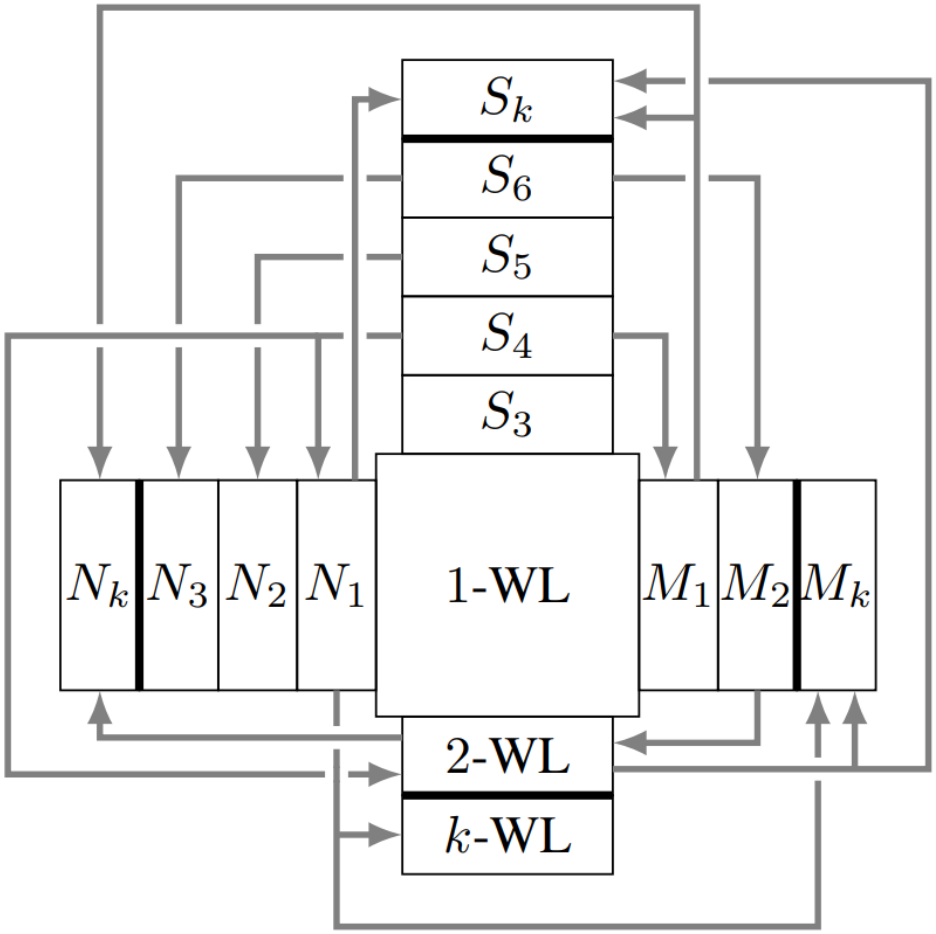

A ⟶ ***B:*** *A can distinguish a pair of graphs that B cannot*

Comparison of expressiveness

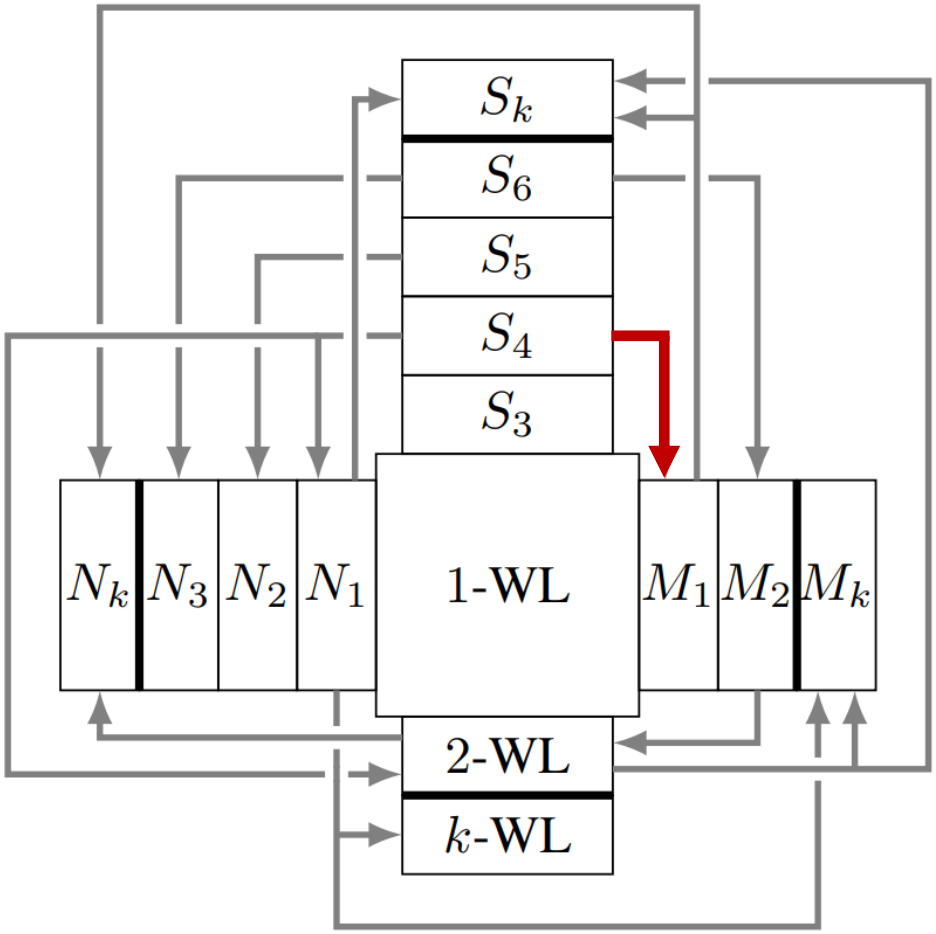

A $\longrightarrow$ ***B:*** *A can distinguish a pair of graphs that B cannot*

Comparison of expressiveness

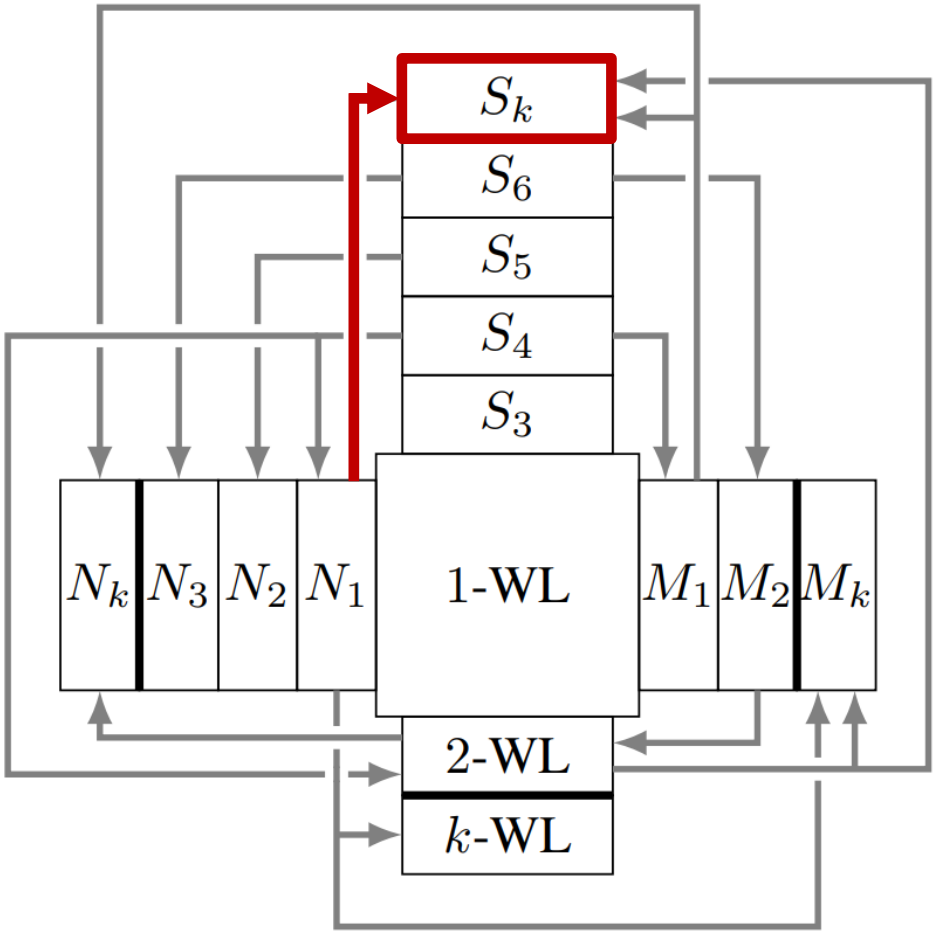

A $\longrightarrow$ ***B:*** *A can distinguish a pair of graphs that B cannot*

Comparison of expressiveness

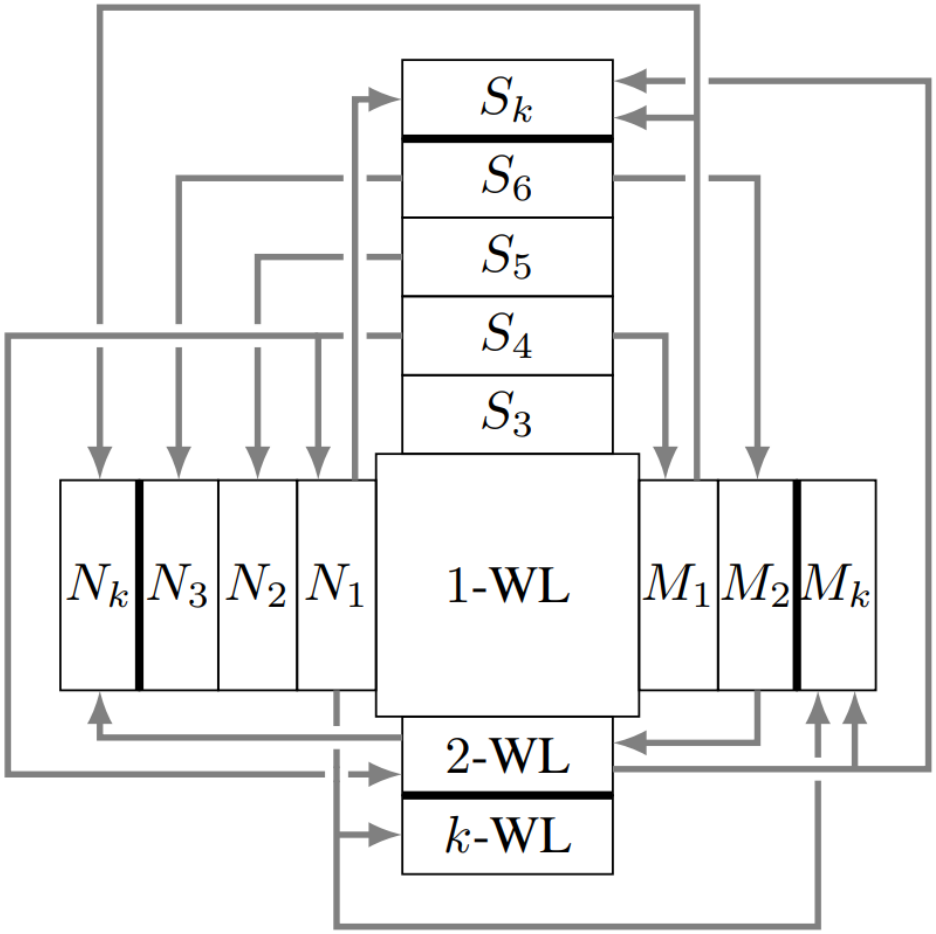

A ⟶ ***B:*** *A can distinguish a pair of graphs that B cannot*

Counting substructures

Can GNNs count **cliques** or **cycles**?

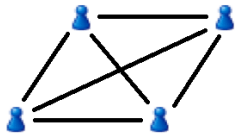

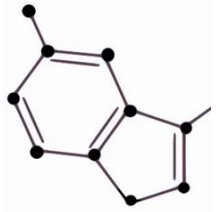

	S_k	N_k	M_k
can count	k-cliques k-cycles	any clique *already with* $k=1$ $(2k+1)$-cycles	$(k+2)$-cliques $(k+1)$-cycles *only if* $d \geq k+1$
cannot count	$(k+1)$-cliques *for small* k *values* $(k+1)$-cycles	$(2k+2)$-cycles	$(k+3)$-cliques *for small* k *values*

A Theoretical Comparison of Graph Neural Network Extensions

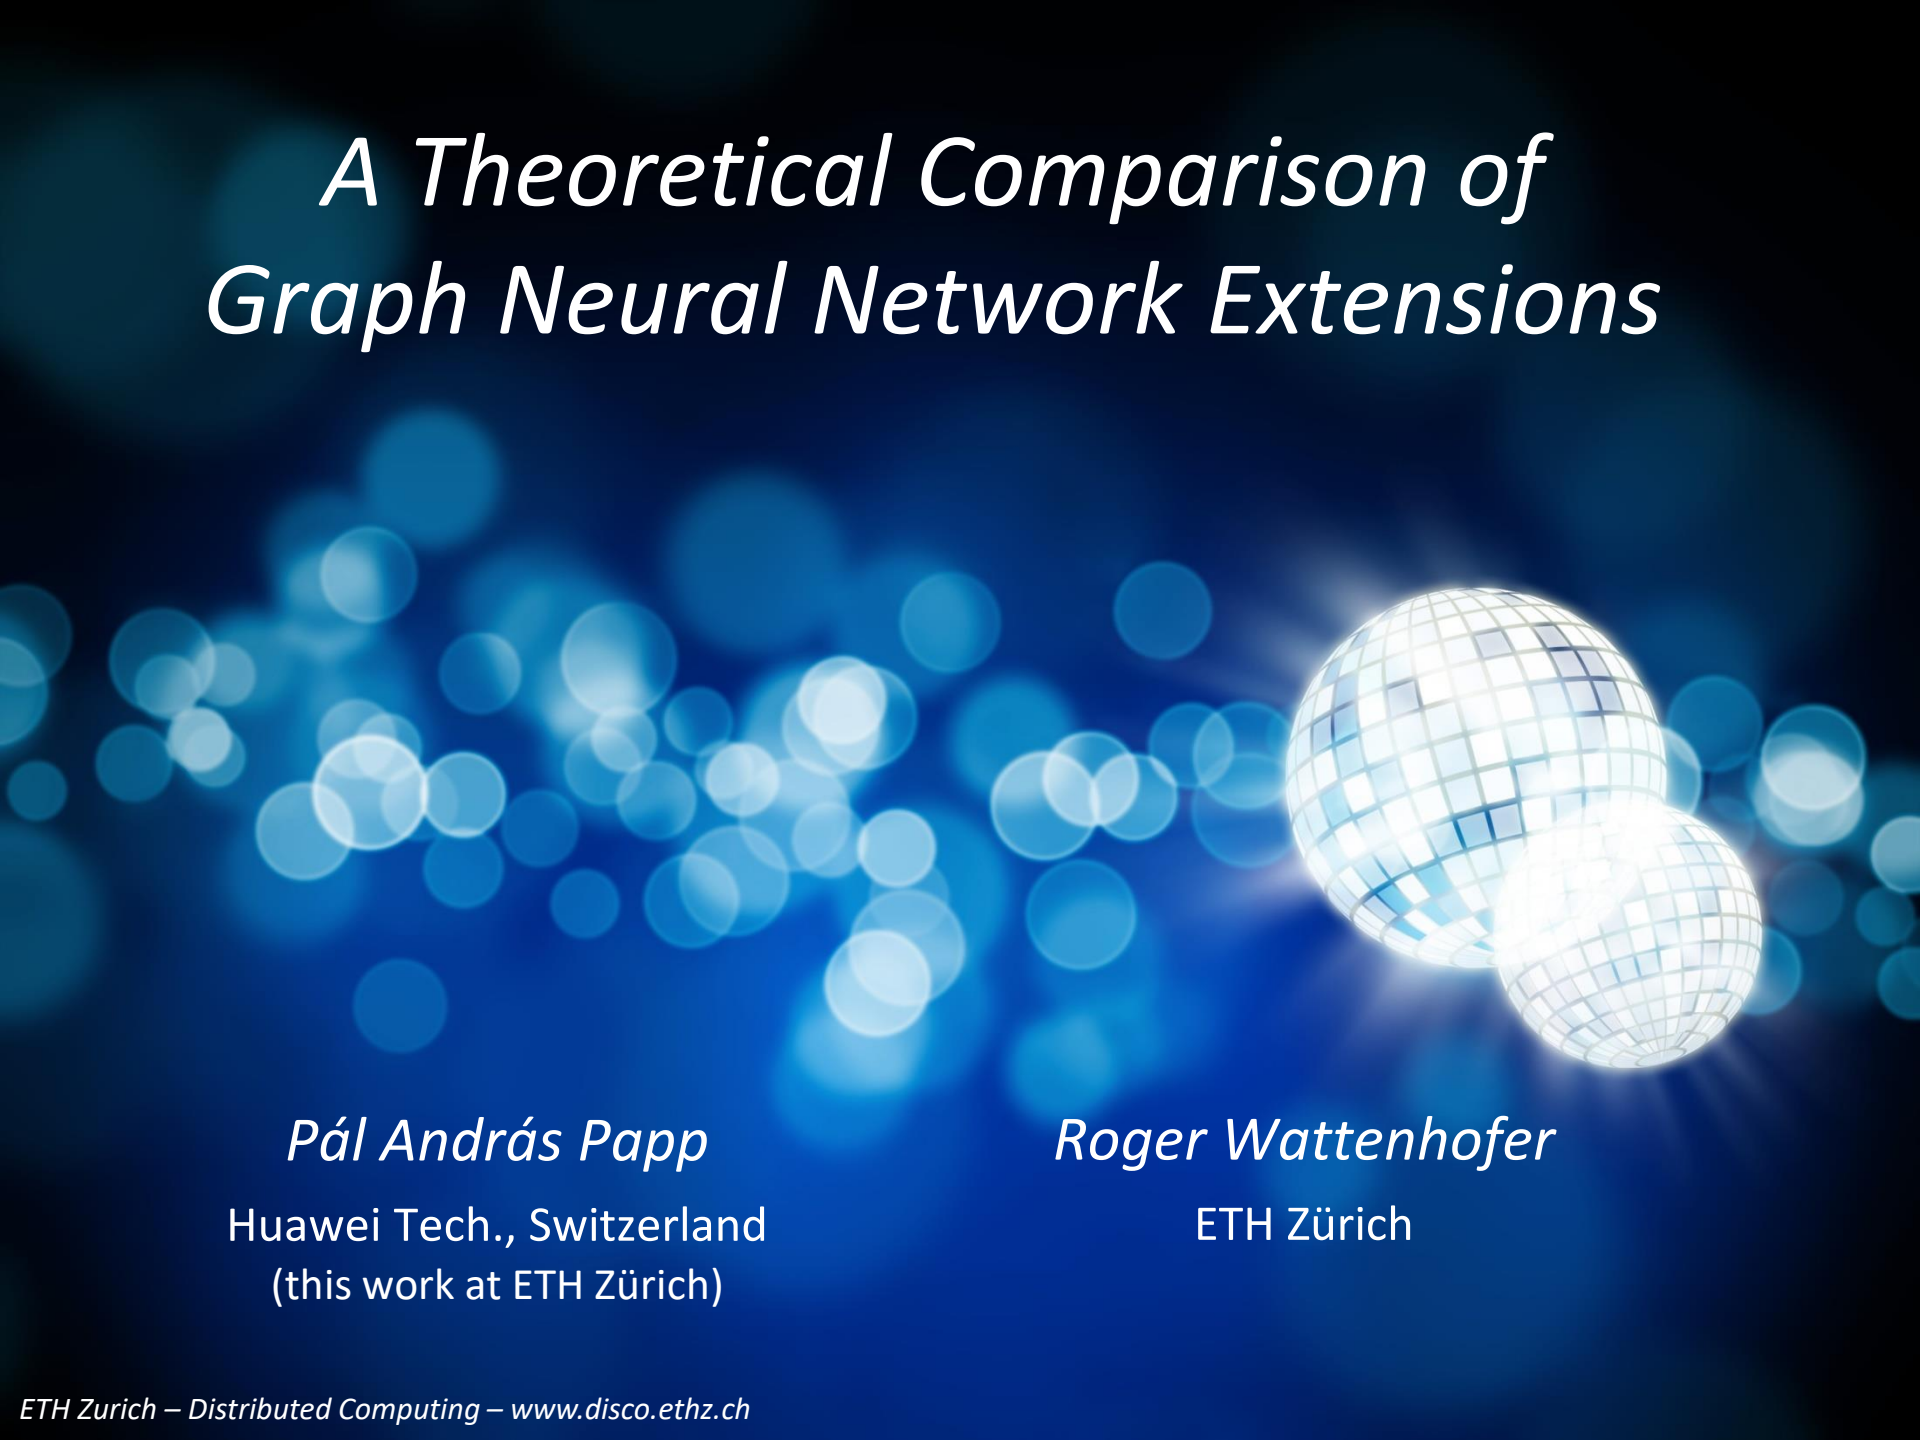

Pál András Papp

Huawei Tech., Switzerland
(this work at ETH Zürich)

Roger Wattenhofer

ETH Zürich