

# GRAIN: Exact Graph Reconstruction from Gradients



Maria Drencheva



Ivo Petrov



Maximilian Baader



Dimitar I. Dimitrov



Martin Vechev

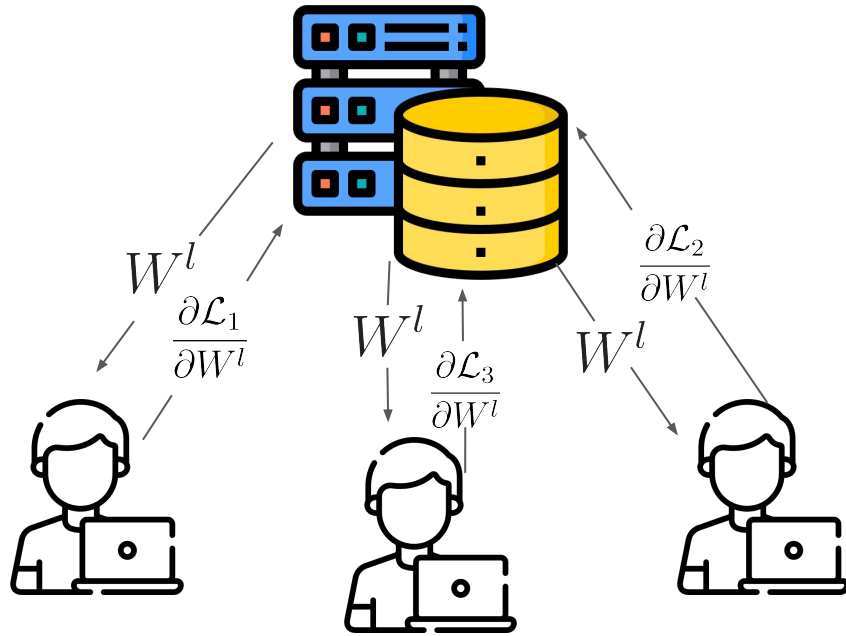
**INSAIT**



**SRILAB**

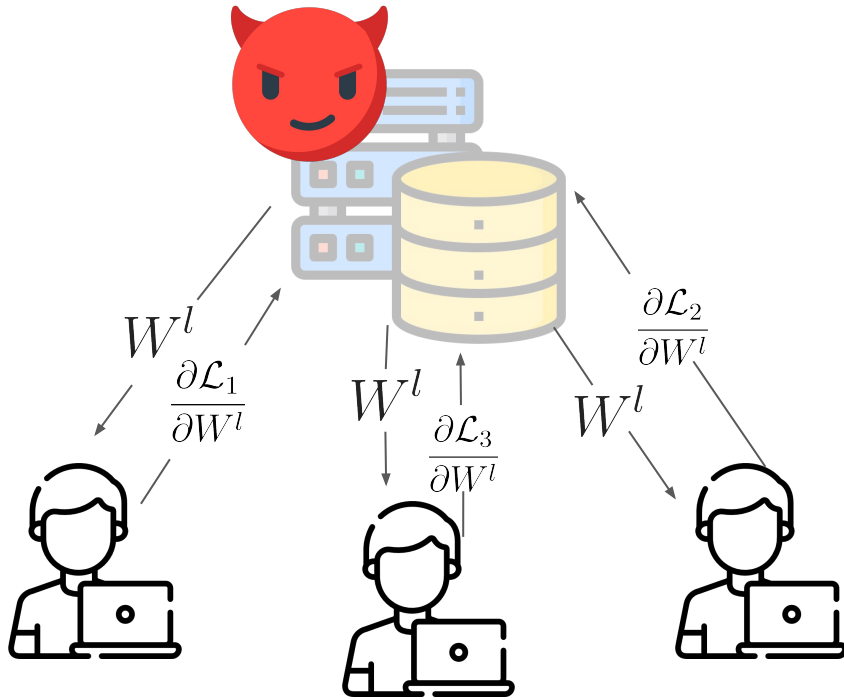
**ETH** zürich

# Federated Learning and Gradient Inversion



**Federated learning** enables training a model across multiple clients without sharing their raw data with a central server; instead, they only share **gradient updates**.

# Federated Learning and Gradient Inversion



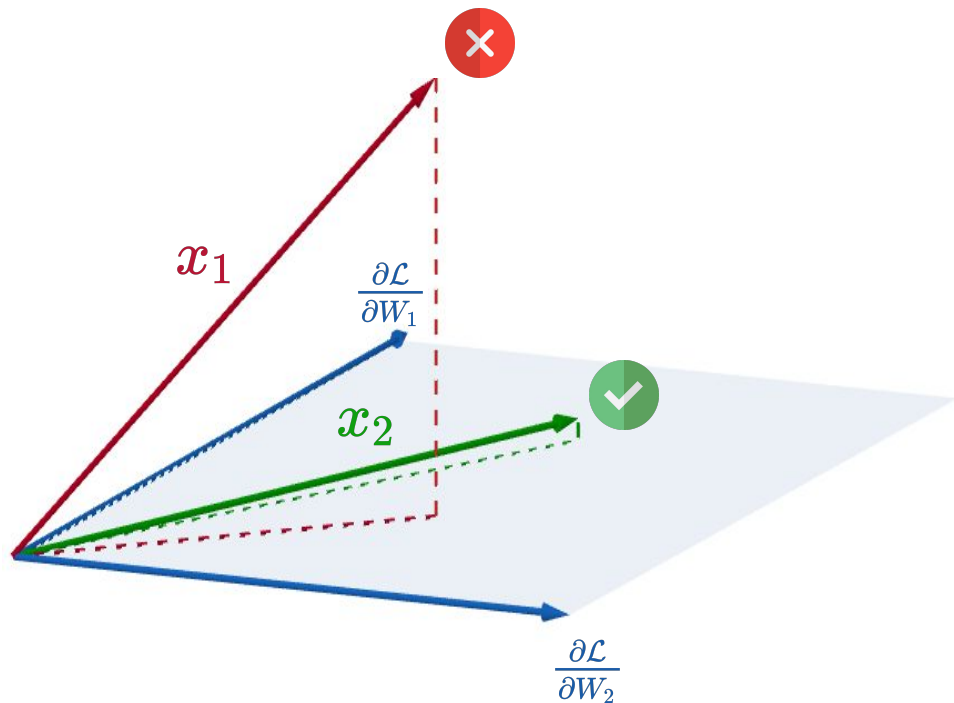
Federated learning enables training a model across multiple clients without sharing their raw data with a central server; instead, they only share **gradient updates**.

However, **Gradient Inversion Attacks** have revealed privacy risks in Federated Learning, as client data can sometimes be reconstructed from the shared gradient updates.

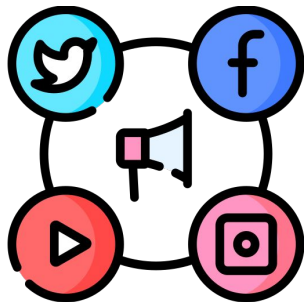
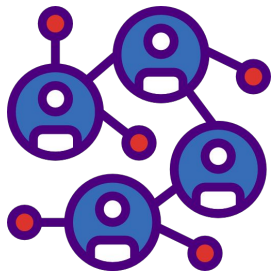
# Spancheck filtering of inputs of linear layers

$$\mathbf{Y} = \mathbf{XW}$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{W}} = \mathbf{X}^T \frac{\partial \mathcal{L}}{\partial \mathbf{Y}}$$

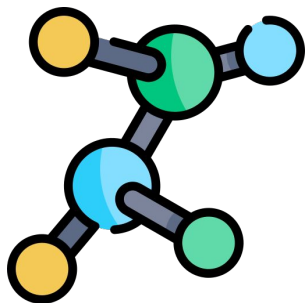
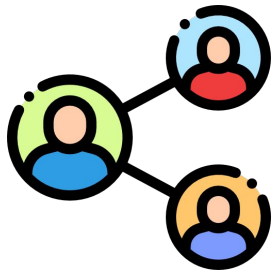


# Graph Neural Networks in Federated Learning

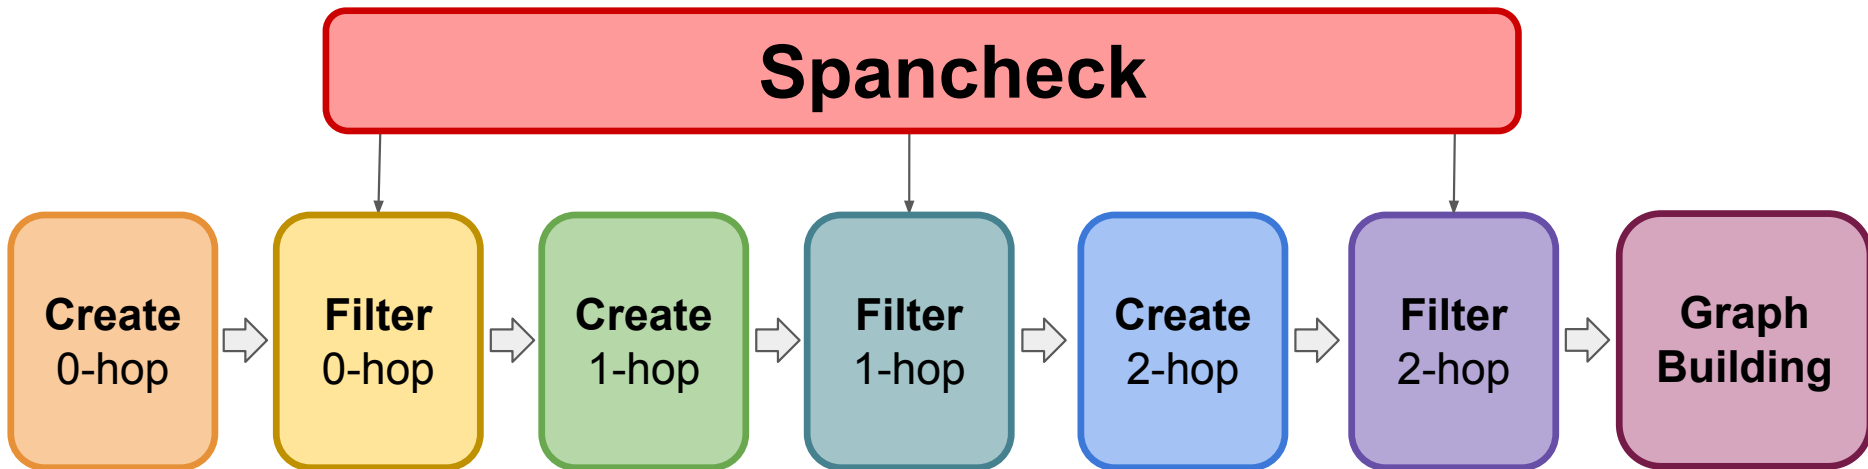


**Graph Neural Networks** allow for models to be trained on graph data, such as molecules or social and citation networks.

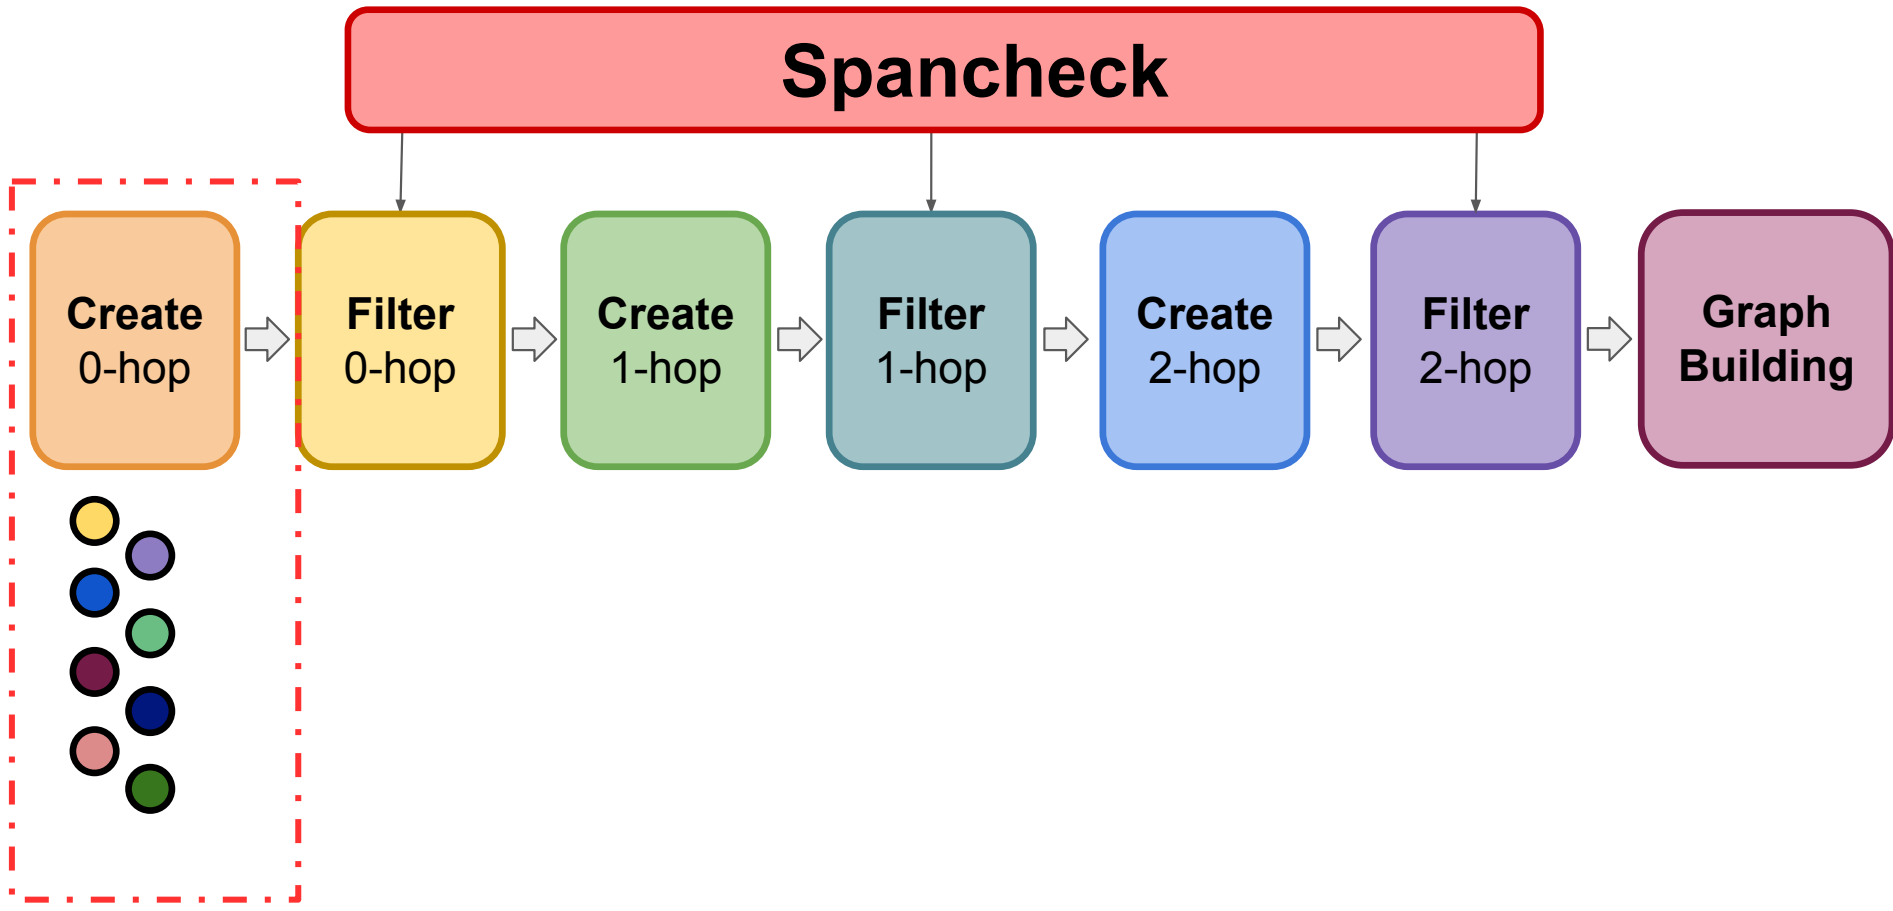
$$X^{l+1} = \text{ReLU}(A^l X^l W^l)$$



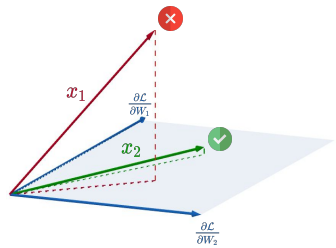
# GRAIN Overview



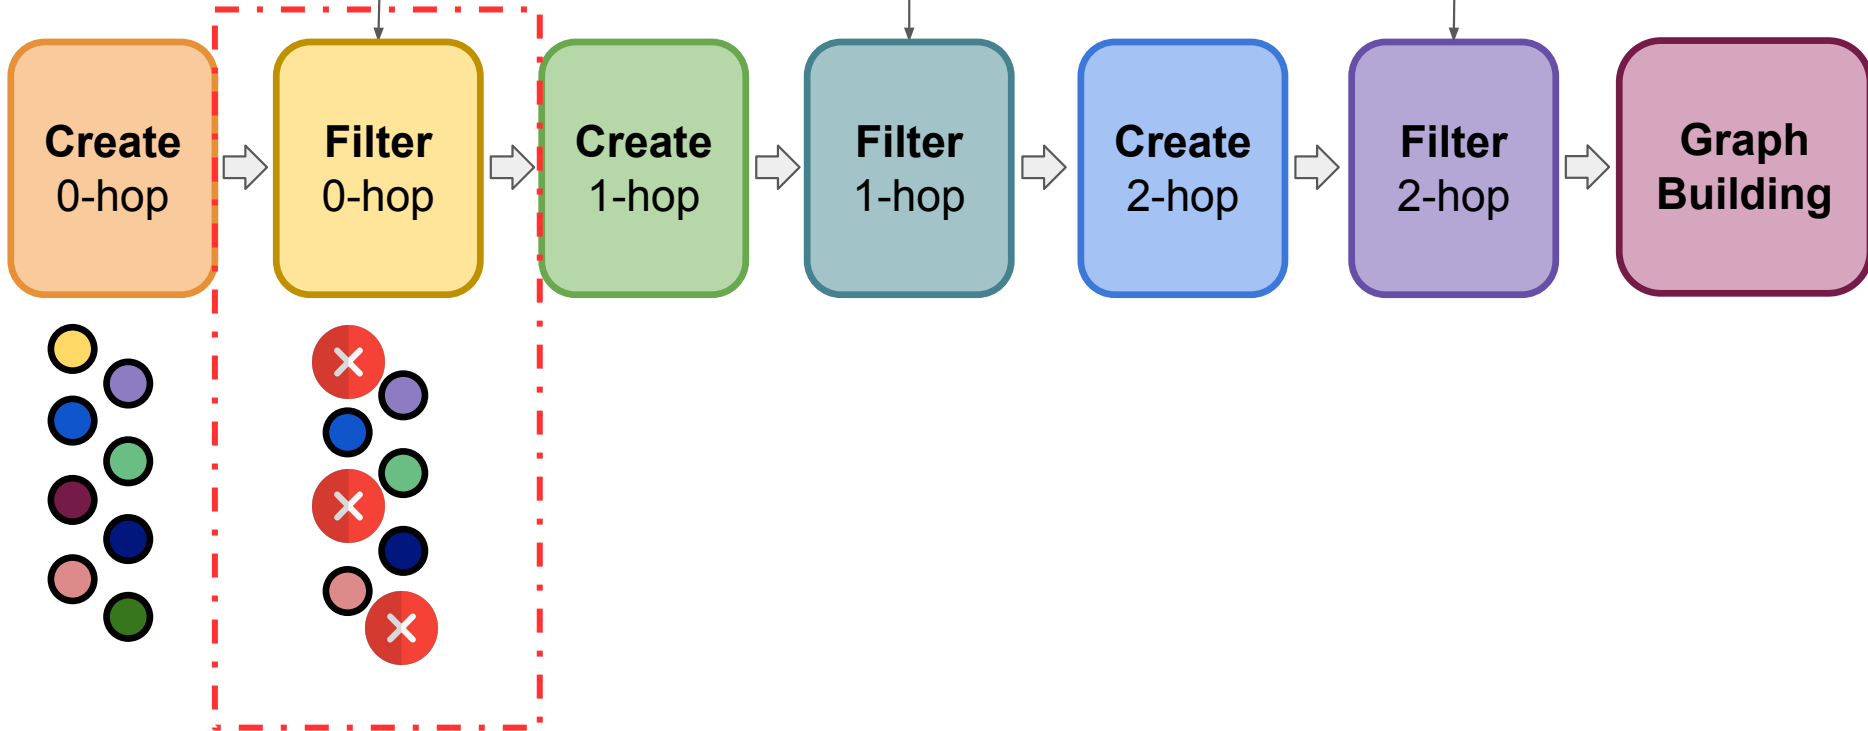
# GRAIN Overview



# GRAIN Overview



## Spancheck



# Single-node reconstruction

Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer

Atom type



# Single-node reconstruction

Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer

Atom type



# Single-node reconstruction









Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer

Atom type



















# Single-node reconstruction

Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer

Atom type	Number of bonds
	
	
	
	
	

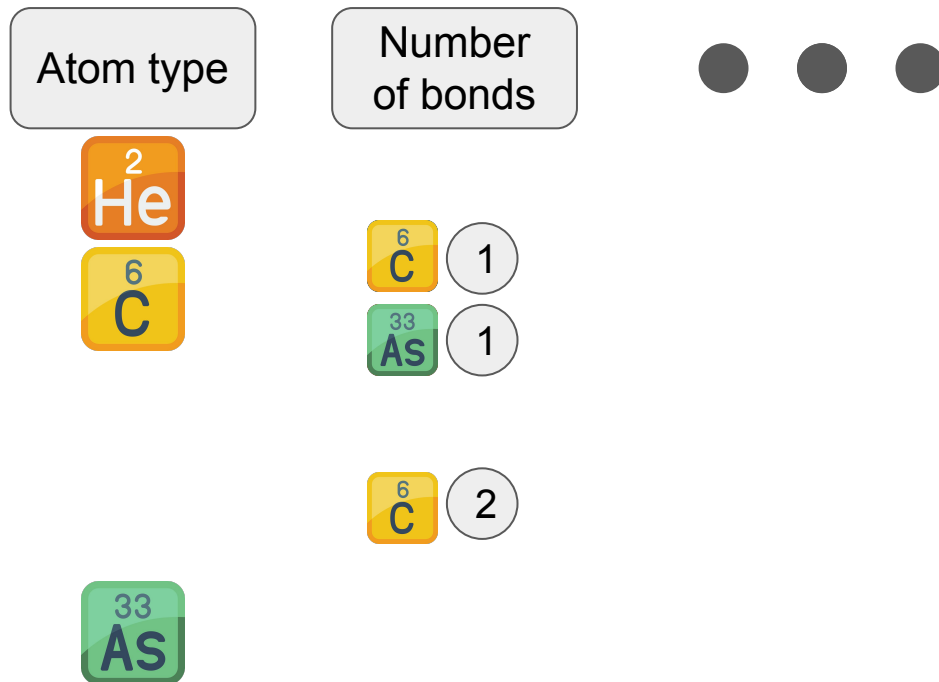
# Single-node reconstruction

Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer

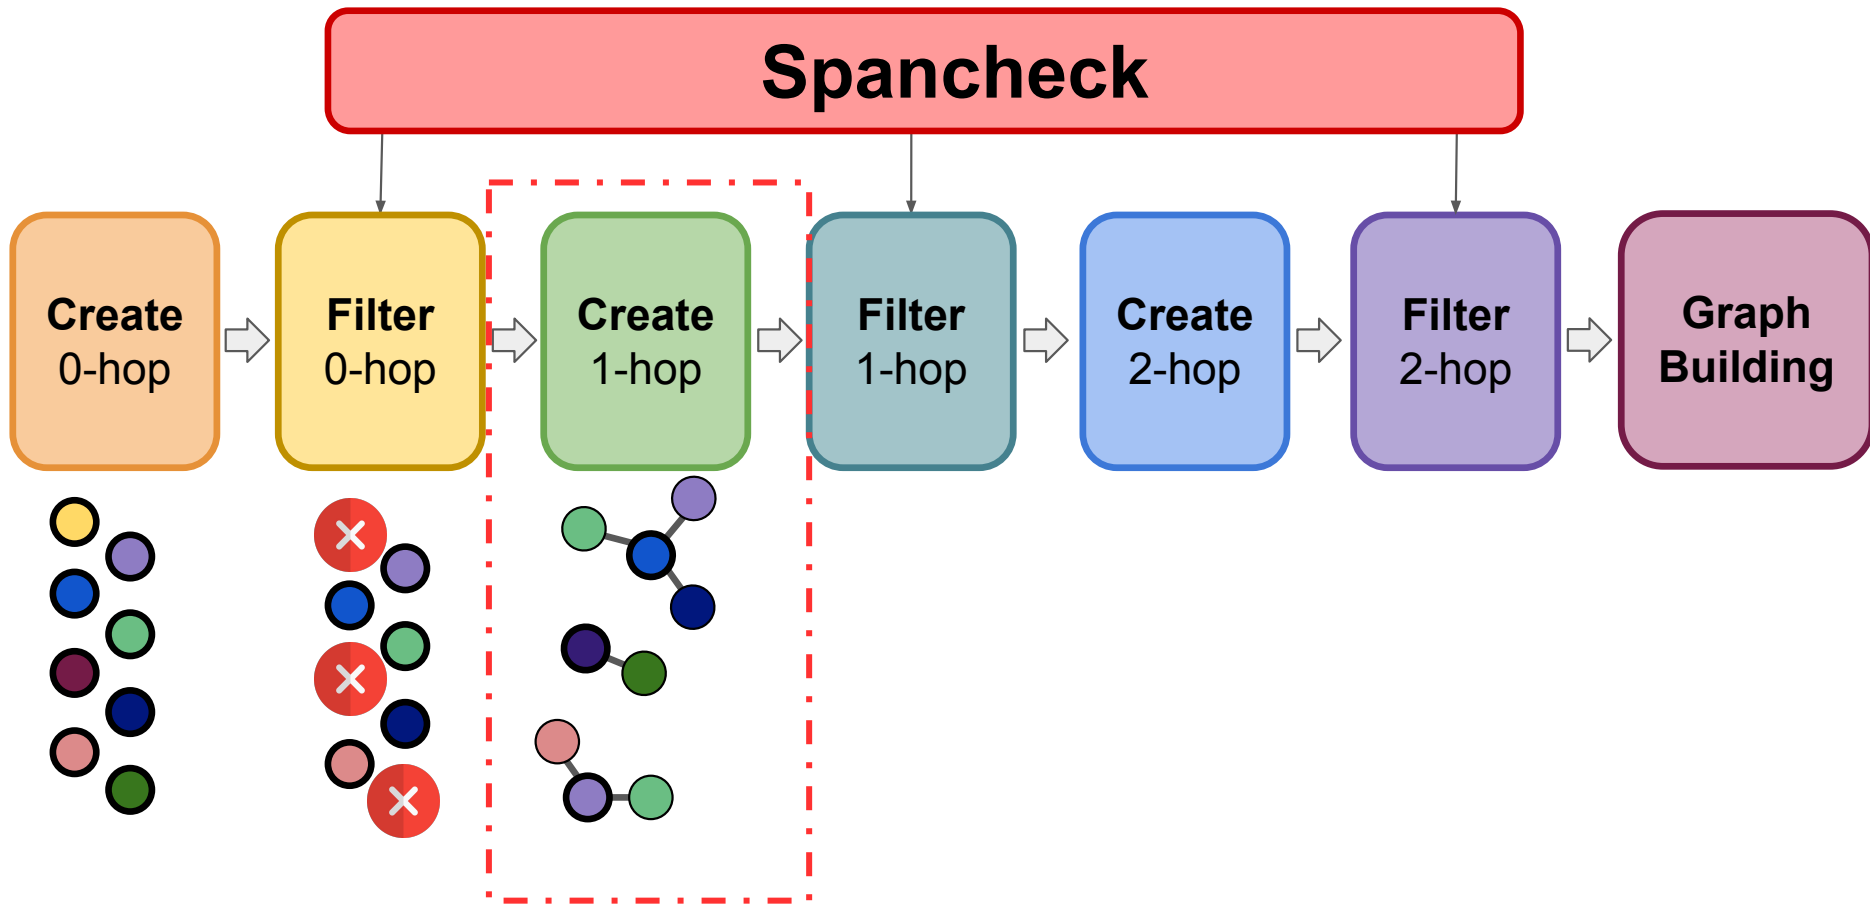
Atom type	Number of bonds		
		1	
		1	
		1	
		2	
		2	
		2	
			

# Single-node reconstruction

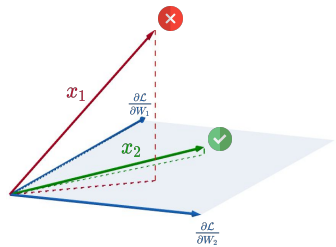
Spancheck - applied on the gradient of the loss w.r.t the weight matrix  $W^0$  of the first GNN layer



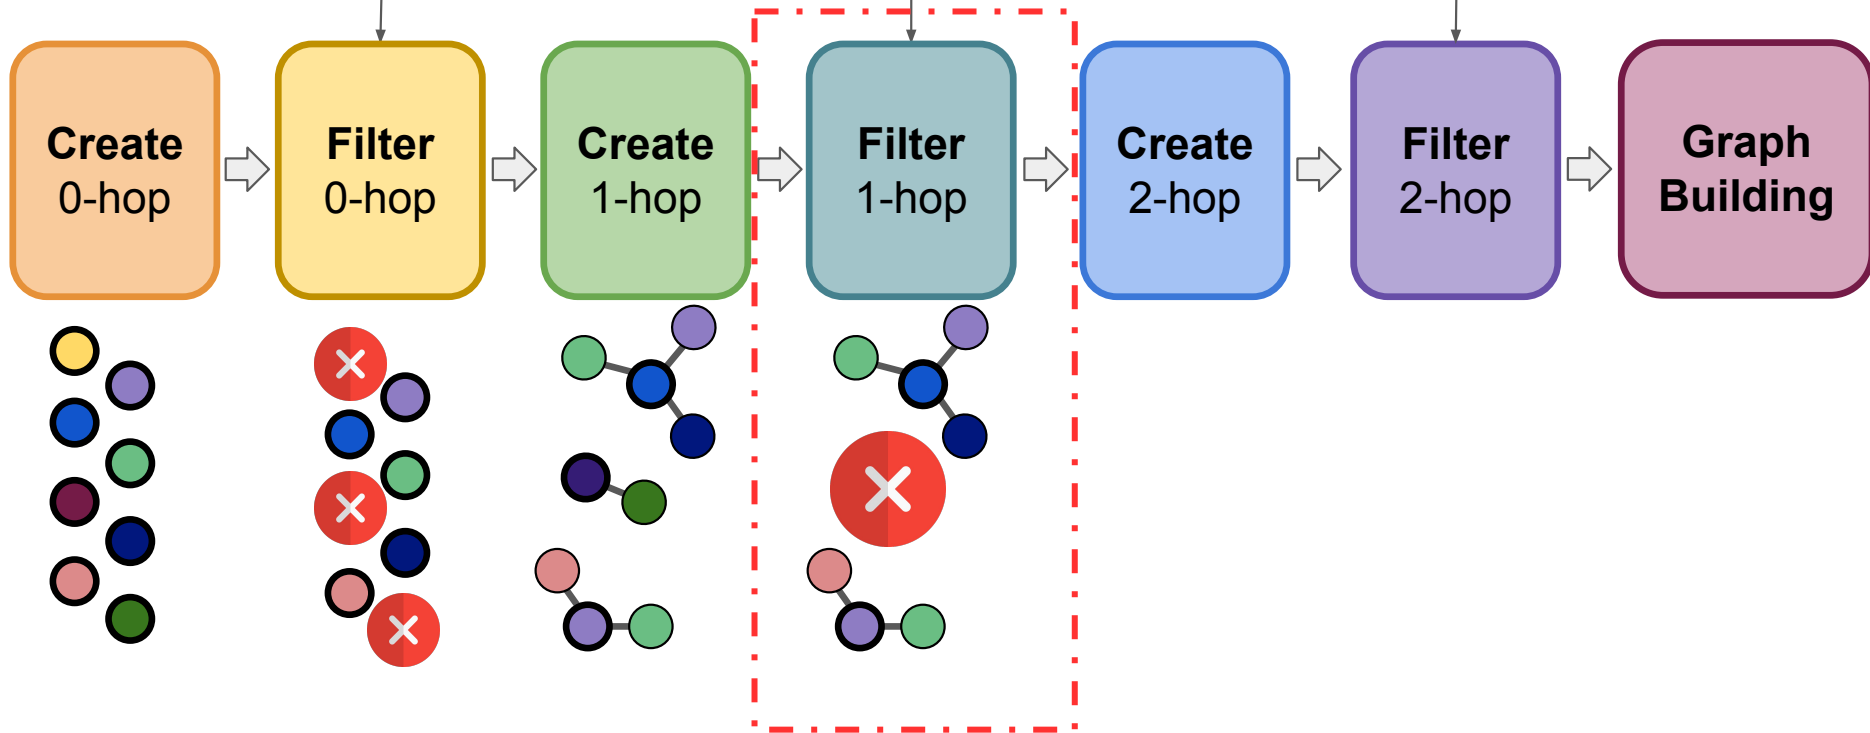
# GRAIN Overview



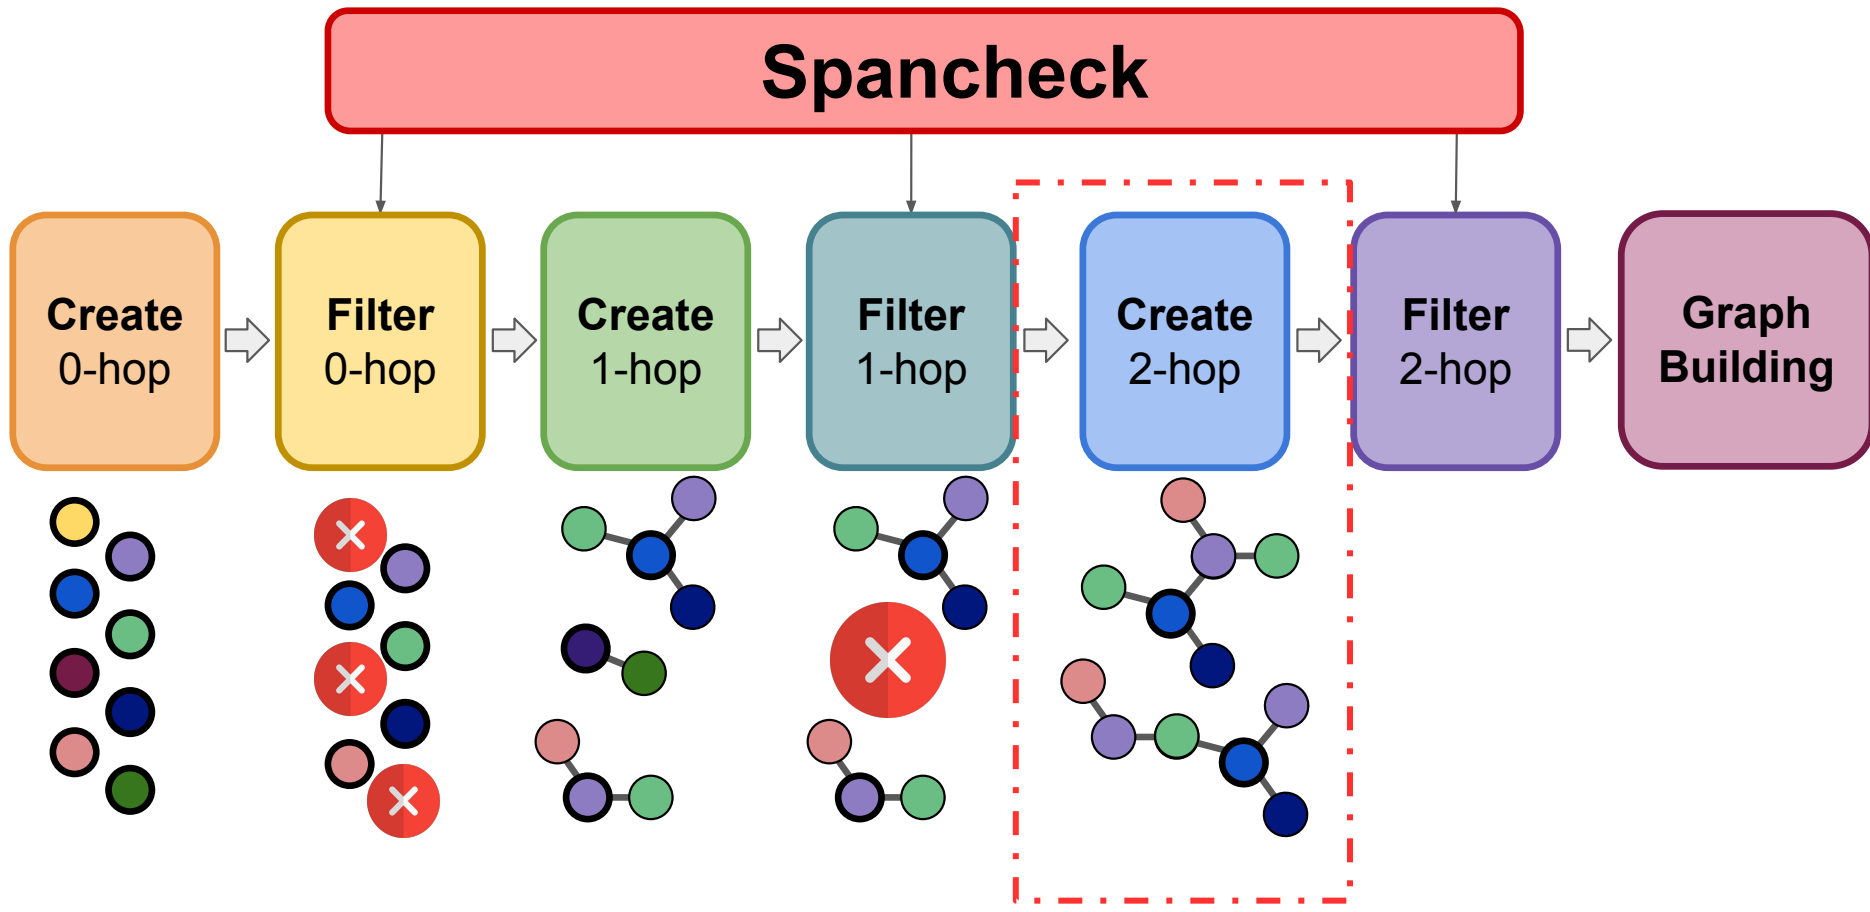
# GRAIN Overview



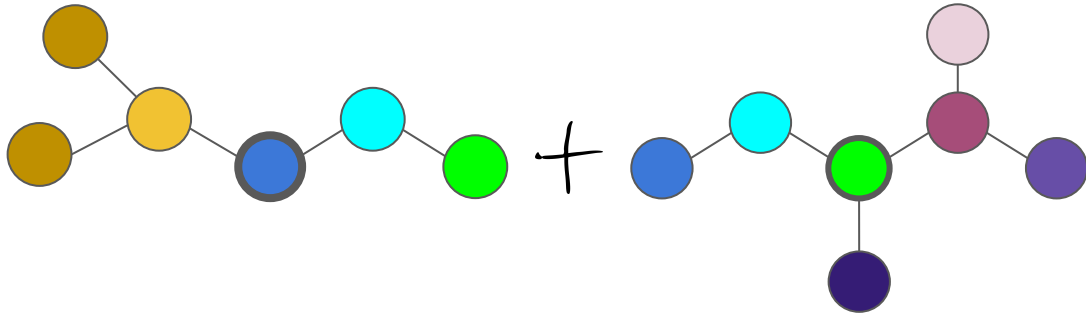
## Spancheck



# GRAIN Overview

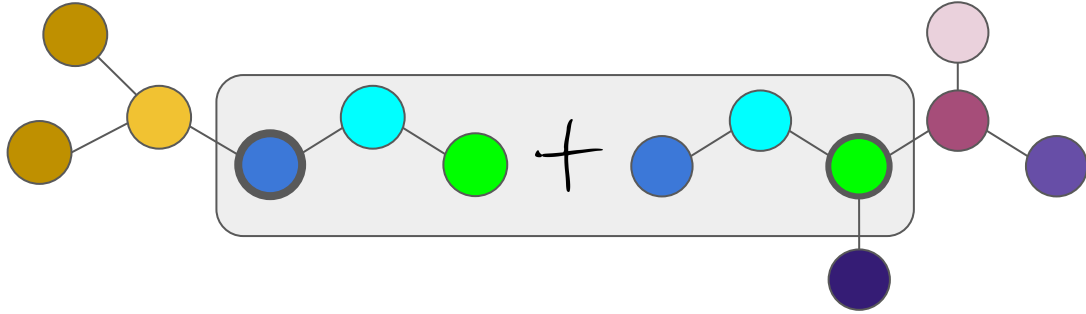


# Building Algorithm - Graph gluing



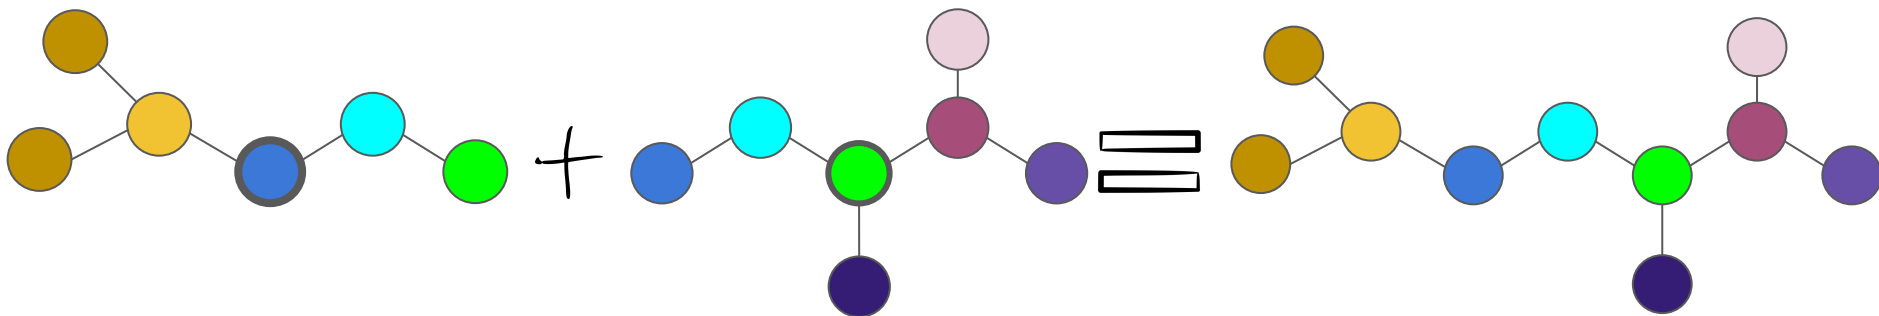
# Building Algorithm - Graph gluing

Gluing building blocks into larger graphs



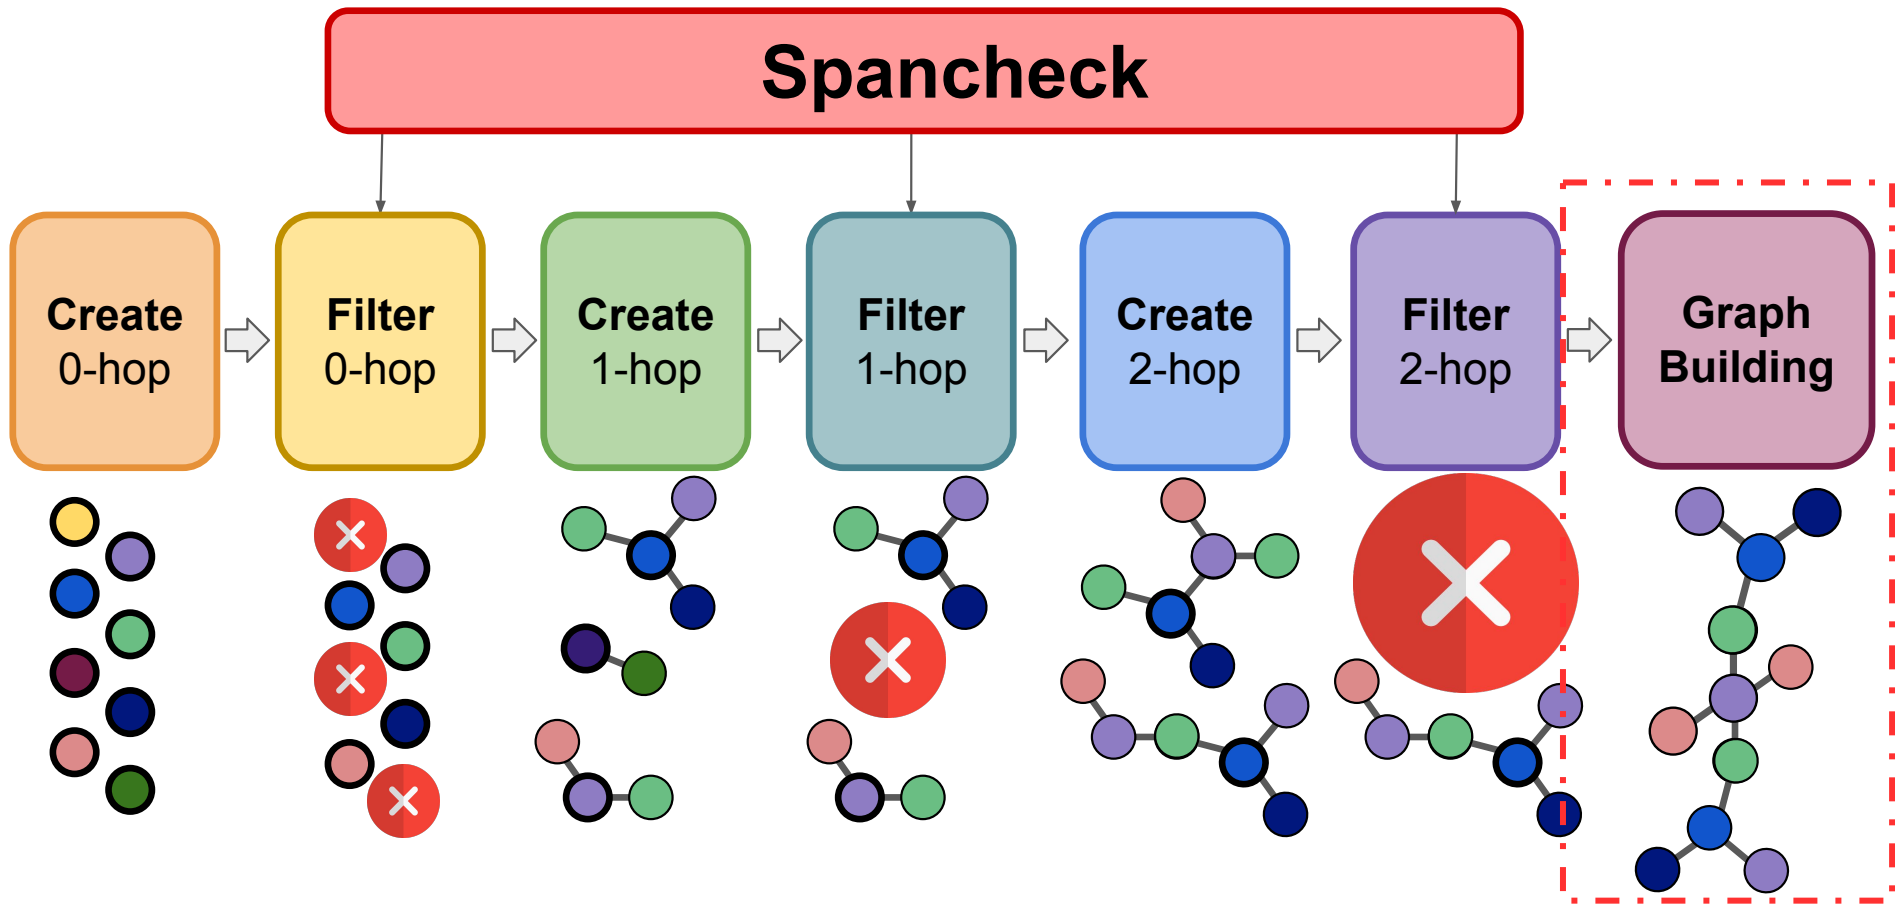
# Building Algorithm - Graph gluing

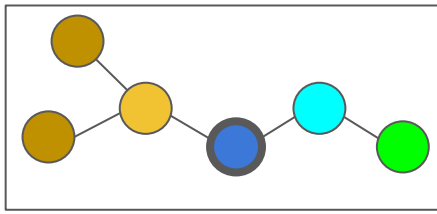
Gluing building blocks into larger graphs

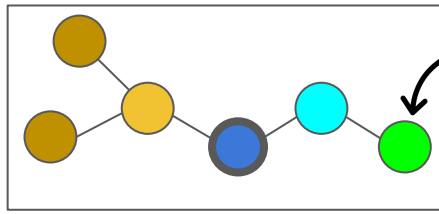




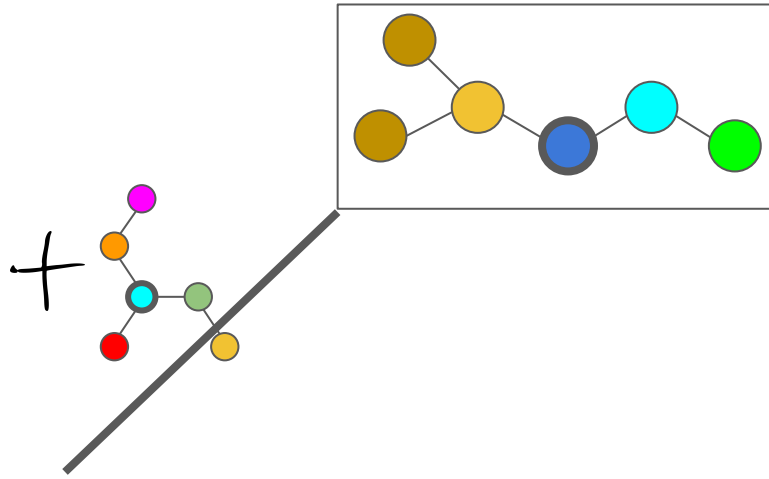
# GRAIN Overview

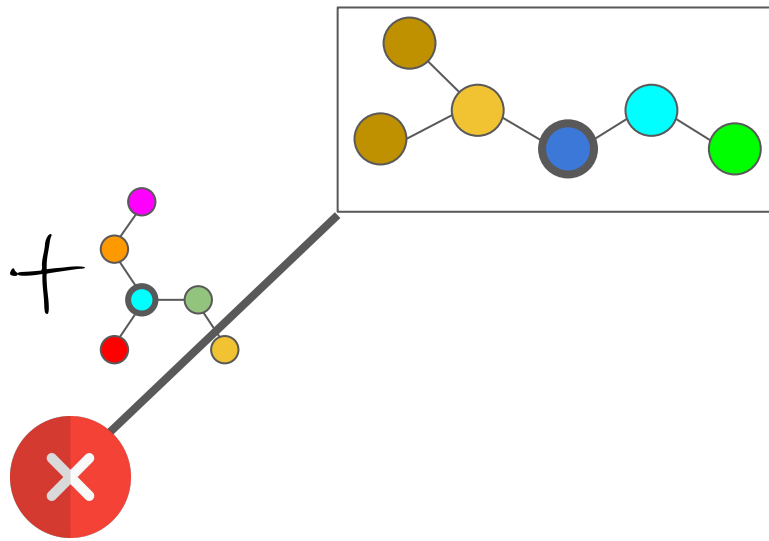


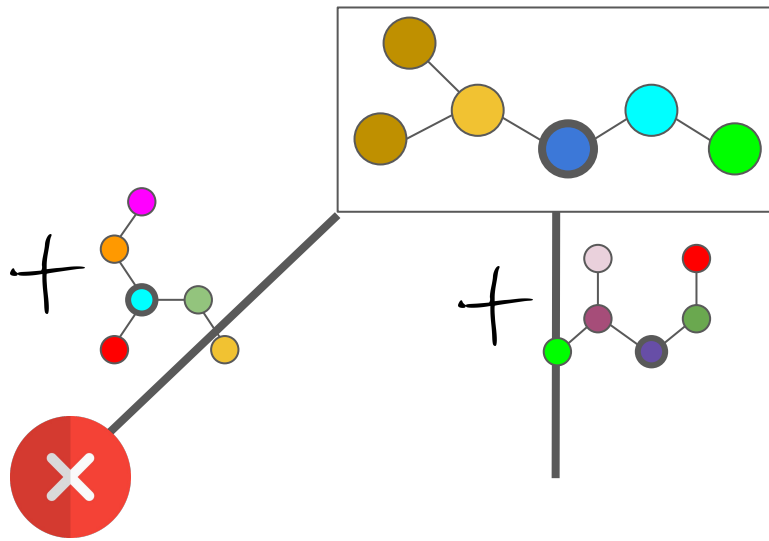


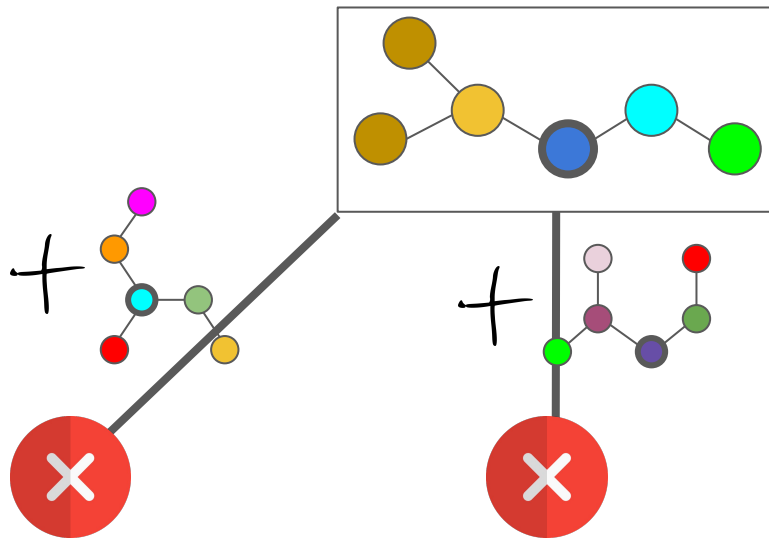


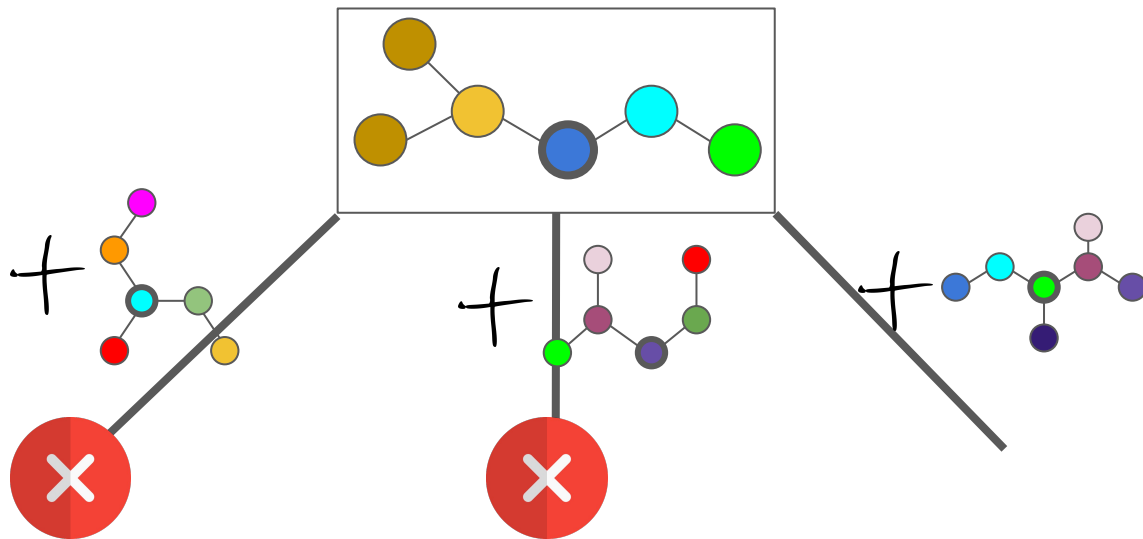
Choosing this node to  
extend the graph

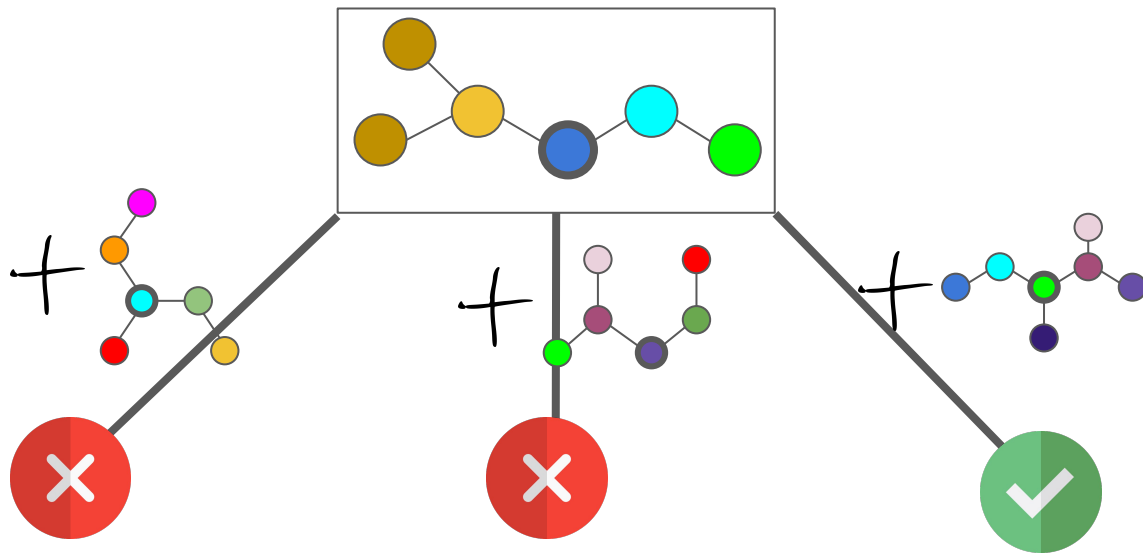


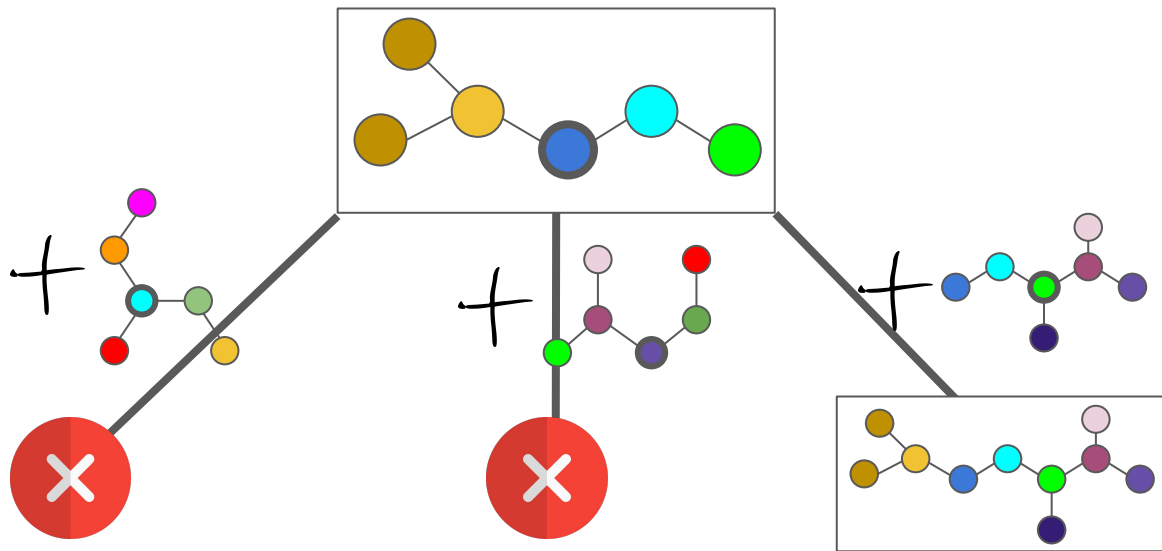


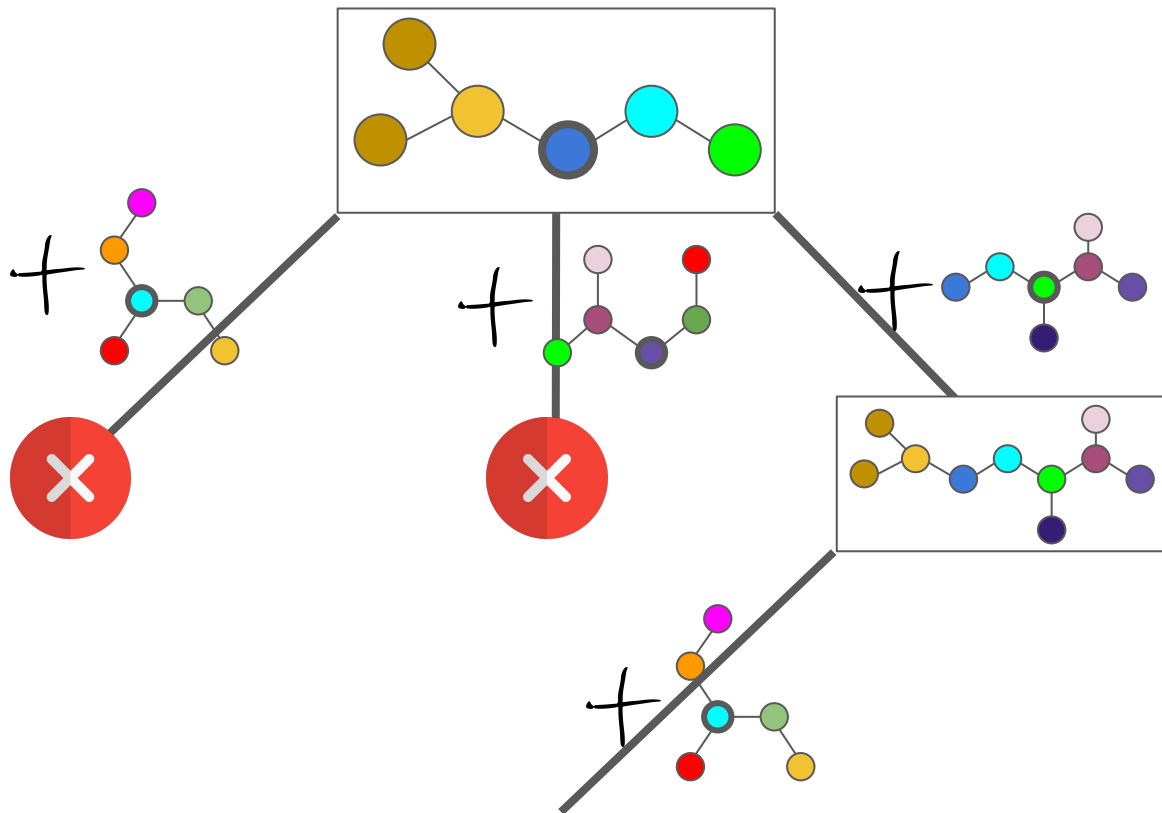


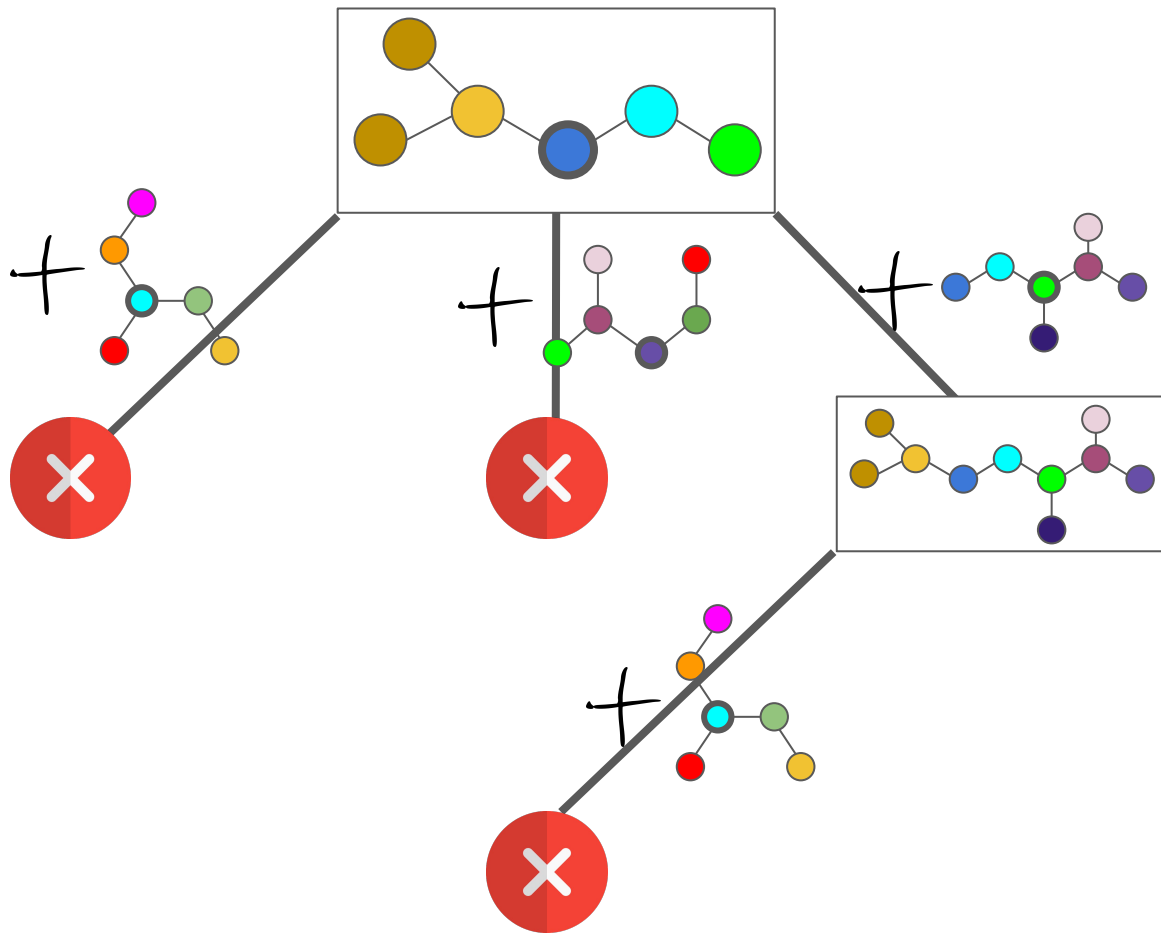


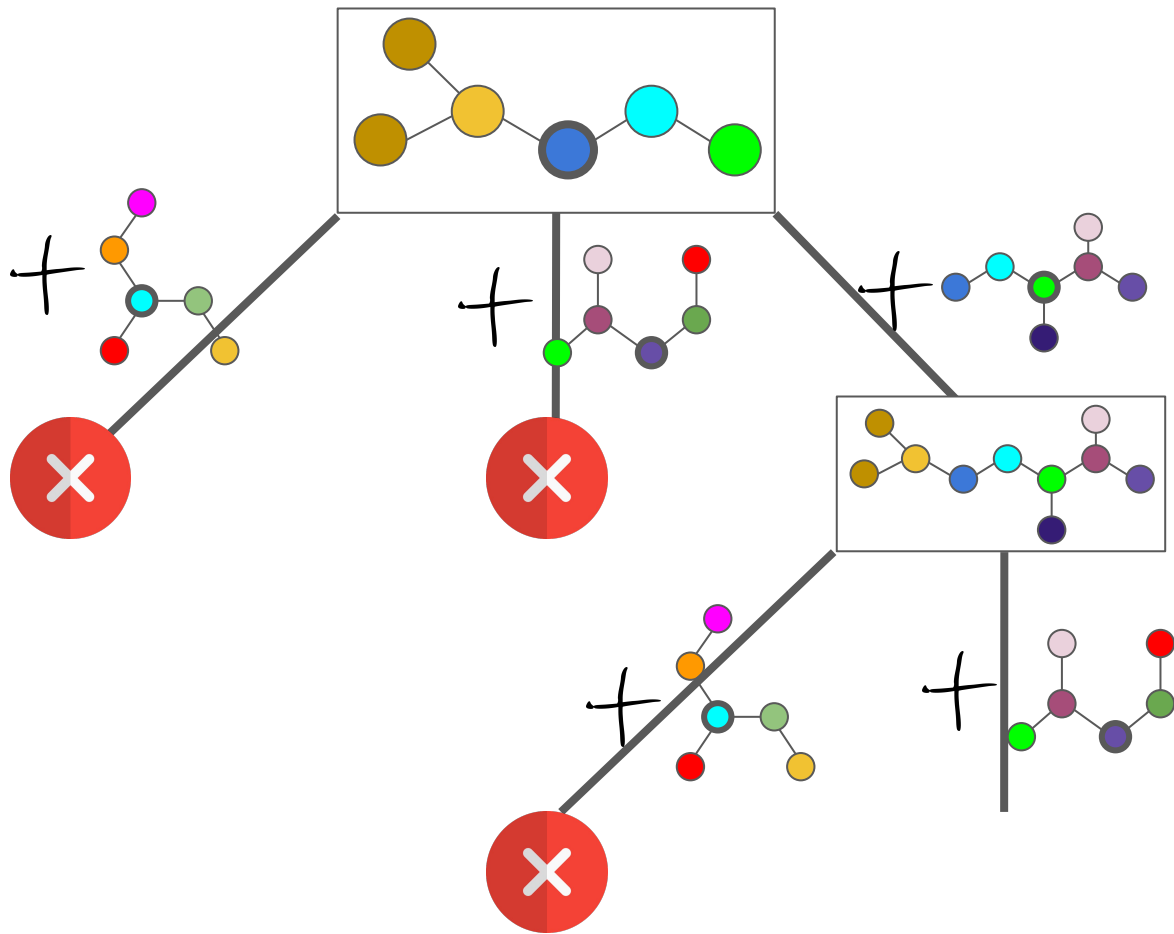


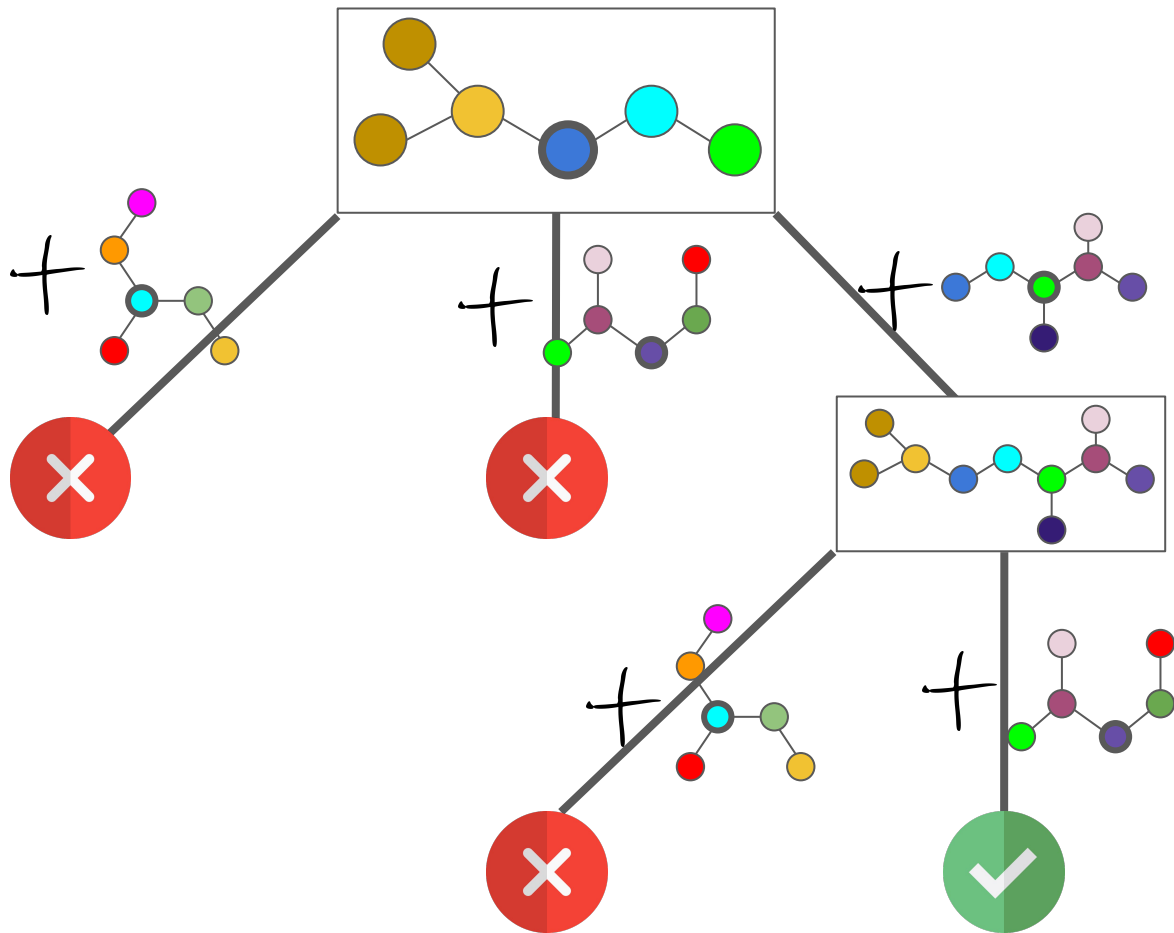


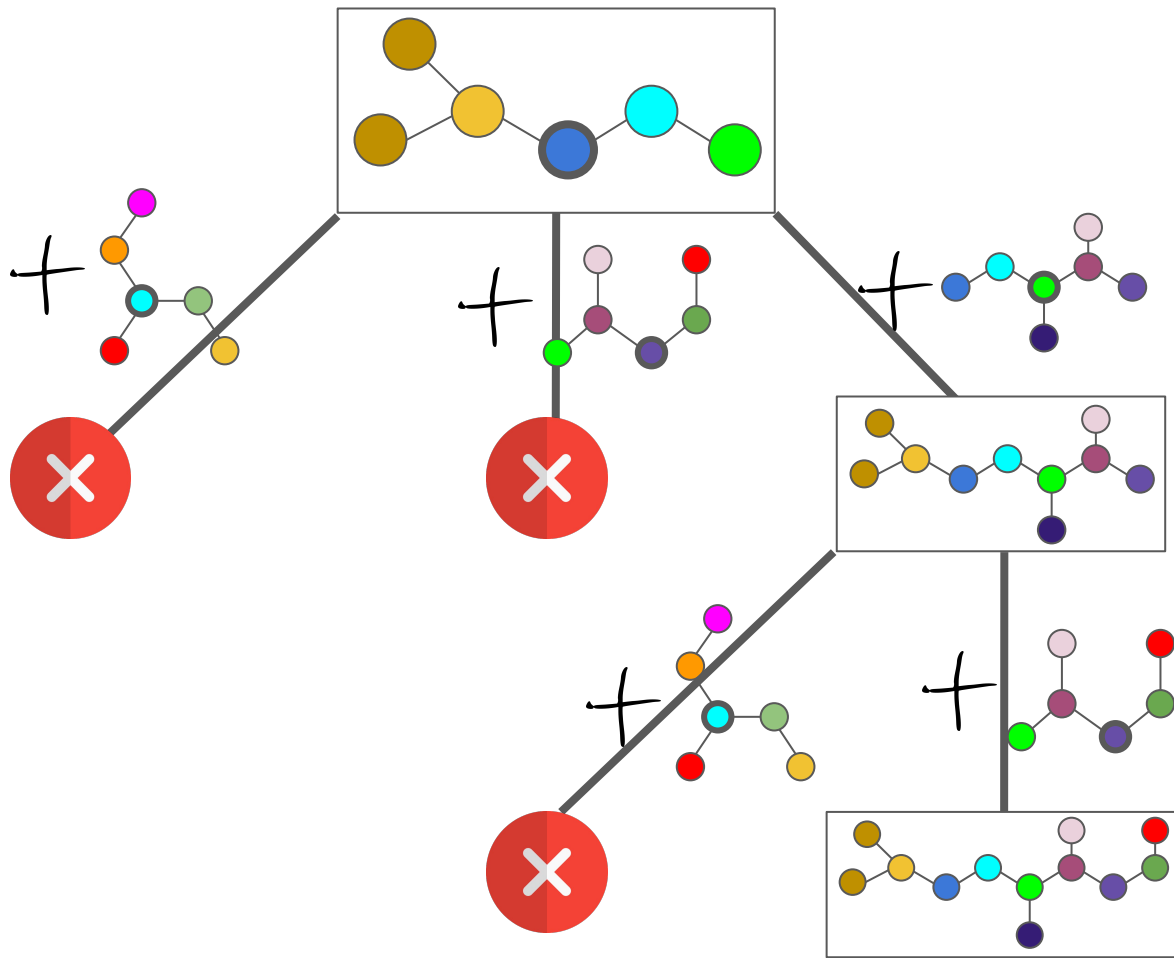












# Evaluation - GCNs vs GATs

Baselines: DLG [1] and TabLeak [2] and the same with the adjacency matrix given (+A)

		GCN					GAT				
		GSM-0	GSM-1	GSM-2	FULL	Time [h]	GSM-0	GSM-1	GSM-2	FULL	Time [h]
Tox21	GRAIN	<b>86.9</b> <sup>+4.2</sup> <sub>-5.7</sub>	<b>83.9</b> <sup>+5.2</sup> <sub>-6.9</sub>	<b>82.6</b> <sup>+5.7</sup> <sub>-7.4</sub>	<b>68.0 ± 1.7</b>	14.3	92.9 <sup>+3.8</sup> <sub>-5.8</sub>	<b>90.7</b> <sup>+5.0</sup> <sub>-7.1</sub>	<b>89.9</b> <sup>+5.8</sup> <sub>-7.2</sub>	<b>75.0 ± 1.8</b>	10.8
	DLG	31.8 <sup>+4.5</sup> <sub>-4.3</sub>	20.3 <sup>+5.5</sup> <sub>-4.8</sub>	22.8 <sup>+6.6</sup> <sub>-5.6</sub>	1.0 ± 0.2	3.3	96.0 ± 0.32	9.3 <sup>+4.4</sup> <sub>-4.9</sub>	6.5 <sup>+3.9</sup> <sub>-4.1</sub>	2.0 ± 0.3	<b>4.2</b>
	DLG +A	54.7 <sup>+3.9</sup> <sub>-4.2</sub>	60.1 <sup>+4.6</sup> <sub>-5.2</sub>	76.7 <sup>+3.6</sup> <sub>-4.8</sub>	1.0 ± 0.2	<b>3.1</b>	<b>96.5 ± 0.34</b>	69.7 <sup>+4.1</sup> <sub>-4.2</sub>	81.3 <sup>+3.4</sup> <sub>-3.6</sub>	2.0 ± 0.3	4.5
	TabLeak	25.1 <sup>+5.1</sup> <sub>-4.3</sub>	12.4 <sup>+5.5</sup> <sub>-4.3</sub>	10.8 <sup>+5.6</sup> <sub>-3.9</sub>	1.0 ± 0.2	13.1	73.7 <sup>+2.6</sup> <sub>-2.0</sub>	7.2 <sup>+5.2</sup> <sub>-4.9</sub>	10.0 ± 4.8	1.0 ± 0.2	6.0
	TabLeak +A	55.6 <sup>+3.9</sup> <sub>-3.9</sub>	57.7 <sup>+4.1</sup> <sub>-4.6</sub>	73.8 <sup>+2.8</sup> <sub>-3.5</sub>	1.0 ± 0.2	12.3	75.1 <sup>+2.5</sup> <sub>-1.9</sub>	74.9 <sup>+2.1</sup> <sub>-1.9</sub>	84.2 <sup>+1.5</sup> <sub>-1.3</sub>	1.0 ± 0.2	6.0

[1] Zhu et. al. in “Deep Leakage from Gradients”

[2] Vero et. al. in “TabLeak: Tabular Data Leakage in Federated Learning”

# Evaluation – GRAIN on different domains

		GAT				
		GSM-0	GSM-1	GSM-2	FULL	Min/Rec
Citation Network	GRAIN	<b>79.3<sup>+4.7</sup><sub>-6.3</sub></b>	<b>69.1<sup>+6.1</sup><sub>-6.4</sub></b>	<b>69.6<sup>+6.2</sup><sub>-6.0</sub></b>	<b>61.0 ± 1.6</b>	<b>0.8</b>
	DLG	67.7 <sup>+3.9</sup> <sub>-3.7</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 ± 0.0	31.0
	DLG +A	67.7 <sup>+4.0</sup> <sub>-3.7</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 ± 0.0	27.7
	TabLeak	67.7 <sup>+3.9</sup> <sub>-3.8</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 ± 0.0	153.0
	TabLeak +A	67.7 <sup>+4.0</sup> <sub>-3.7</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 ± 0.0	148.7
Social Network	GRAIN	<b>97.2<sup>+1.6</sup><sub>-1.9</sub></b>	<b>93.5<sup>+3.4</sup><sub>-4.2</sub></b>	<b>96.3<sup>+1.9</sup><sub>-2.3</sub></b>	<b>79.0 ± 1.8</b>	<b>0.2</b>
	DLG	44.7 <sup>+2.3</sup> <sub>-2.3</sub>	2.2 <sup>+3.1</sup> <sub>-2.2</sub>	0.0 <sup>+0.0</sup> <sub>-0.0</sub>	0.0 ± 0.0	26.3
	DLG +A	57.4 <sup>+3.7</sup> <sub>-3.9</sub>	69.5 <sup>+3.6</sup> <sub>-4.0</sub>	88.6 <sup>+2.0</sup> <sub>-2.1</sub>	0.0 ± 0.0	21.6
	TabLeak	50.8 <sup>+12.4</sup> <sub>-8.9</sub>	13.9 <sup>+13.5</sup> <sub>-12.3</sub>	7.9 <sup>+11.9</sup> <sub>-7.9</sub>	0.0 ± 0.0	204.5
	TabLeak +A	52.6 <sup>+3.3</sup> <sub>-3.3</sub>	68.1 <sup>+4.1</sup> <sub>-3.9</sub>	82.7 <sup>+4.0</sup> <sub>-4.9</sub>	0.0 ± 0.0	254.5
Chemical dataset	GRAIN	92.9 <sup>+3.8</sup> <sub>-5.8</sub>	<b>90.7<sup>+5.0</sup><sub>-7.1</sub></b>	<b>89.9<sup>+5.8</sup><sub>-7.2</sub></b>	<b>75.0 ± 1.8</b>	10.8
	DLG	96.0 ± 0.32	9.3 <sup>+4.4</sup> <sub>-4.9</sub>	6.5 <sup>+3.9</sup> <sub>-4.1</sub>	2.0 ± 0.3	<b>4.2</b>
	DLG +A	<b>96.5 ± 0.34</b>	69.7 <sup>+4.1</sup> <sub>-4.2</sub>	81.3 <sup>+3.4</sup> <sub>-3.6</sub>	2.0 ± 0.3	4.5
	TabLeak	73.7 <sup>+2.6</sup> <sub>-2.0</sub>	7.2 <sup>+5.2</sup> <sub>-4.9</sub>	10.0 ± 4.8	1.0 ± 0.2	6.0
	TabLeak +A	75.1 <sup>+2.5</sup> <sub>-1.9</sub>	74.9 <sup>+2.1</sup> <sub>-1.9</sub>	84.2 <sup>+1.5</sup> <sub>-1.3</sub>	1.0 ± 0.2	6.0

# Evaluation - Model Width and Depth

		GSM-0	GSM-1	GSM-2	FULL
$L = 2,$ $d' = 300$ (default)	GRAIN	<b>86.9<sup>+4.2</sup><sub>-5.7</sub></b>	<b>83.9<sup>+5.2</sup><sub>-6.9</sub></b>	<b>82.6<sup>+5.7</sup><sub>-7.4</sub></b>	<b>68.0 <math>\pm</math> 1.7</b>
	DLG	31.8 <sup>+4.5</sup> <sub>-4.3</sub>	20.3 <sup>+5.5</sup> <sub>-4.8</sub>	22.8 <sup>+6.6</sup> <sub>-5.6</sub>	1.0 $\pm$ 0.2
	DLG + A	54.7 <sup>+3.9</sup> <sub>-4.2</sub>	60.1 <sup>+4.6</sup> <sub>-5.2</sub>	76.7 <sup>+3.6</sup> <sub>-4.8</sub>	1.0 $\pm$ 0.2
	TabLeak	25.1 <sup>+5.1</sup> <sub>-4.3</sub>	12.4 <sup>+5.5</sup> <sub>-4.3</sub>	10.8 <sup>+5.6</sup> <sub>-3.9</sub>	1.0 $\pm$ 0.2
	TabLeak + A	55.6 <sup>+3.9</sup> <sub>-3.9</sub>	57.7 <sup>+4.1</sup> <sub>-4.6</sub>	73.8 <sup>+2.8</sup> <sub>-3.5</sub>	1.0 $\pm$ 0.2
$L = 3,$ $d' = 300$	GRAIN	<b>82.5<sup>+5.7</sup><sub>-7.7</sub></b>	<b>80.7<sup>+6.3</sup><sub>-7.7</sub></b>	<b>80.4<sup>+6.2</sup><sub>-7.8</sub></b>	<b>63.0 <math>\pm</math> 1.6</b>
	DLG	20.3 <sup>+4.3</sup> <sub>-3.4</sub>	7.8 <sup>+5.1</sup> <sub>-3.3</sub>	8.2 <sup>+5.3</sup> <sub>-3.4</sub>	1.0 $\pm$ 0.2
	DLG + A	43.0 <sup>+3.7</sup> <sub>-3.6</sub>	48.0 <sup>+4.3</sup> <sub>-4.5</sub>	66.0 <sup>+3.7</sup> <sub>-4.6</sub>	1.0 $\pm$ 0.2
	TabLeak	16.5 <sup>+3.8</sup> <sub>-2.9</sub>	8.8 <sup>+4.4</sup> <sub>-3.1</sub>	8.0 <sup>+4.3</sup> <sub>-3.0</sub>	1.0 $\pm$ 0.2
	TabLeak + A	47.5 <sup>+4.0</sup> <sub>-4.2</sub>	48.1 <sup>+4.8</sup> <sub>-5.0</sub>	62.9 <sup>+4.3</sup> <sub>-4.4</sub>	1.0 $\pm$ 0.2
$L = 2,$ $d' = 200$	GRAIN	<b>84.6<sup>+4.6</sup><sub>-6.4</sub></b>	<b>81.4<sup>+5.8</sup><sub>-6.9</sub></b>	<b>80.5<sup>+5.9</sup><sub>-7.2</sub></b>	<b>62.0 <math>\pm</math> 1.6</b>
	DLG	30.8 <sup>+4.5</sup> <sub>-4.1</sub>	18.9 <sup>+5.8</sup> <sub>-4.9</sub>	22.2 <sup>+6.7</sup> <sub>-5.4</sub>	1.0 $\pm$ 0.2
	DLG + A	50.3 <sup>+4.2</sup> <sub>-4.2</sub>	53.4 <sup>+5.3</sup> <sub>-5.9</sub>	68.7 <sup>+4.9</sup> <sub>-6.1</sub>	3.0 $\pm$ 0.4
	TabLeak	22.1 <sup>+4.8</sup> <sub>-3.7</sub>	10.3 <sup>+5.3</sup> <sub>-3.6</sub>	8.9 <sup>+5.5</sup> <sub>-3.6</sub>	1.0 $\pm$ 0.2
	TabLeak + A	55.0 <sup>+4.8</sup> <sub>-5.0</sub>	62.1 <sup>+4.9</sup> <sub>-5.9</sub>	76.7 <sup>+3.6</sup> <sub>-4.7</sub>	1.0 $\pm$ 0.2

## Evaluation - Miscellaneous Settings

	GSM-0	GSM-1	GSM-2	FULL
Default	$86.9^{+4.2}_{-5.7}$	$83.9^{+5.2}_{-6.9}$	$82.6^{+5.7}_{-7.4}$	<b><math>68.0 \pm 1.7</math></b>
$\sigma = \text{GELU}$	$82.0^{+5.3}_{-6.7}$	$79.1^{+6.0}_{-7.4}$	$78.4^{+6.2}_{-8.0}$	$61.0 \pm 1.6$
Pre-trained	$73.5^{+6.4}_{-7.4}$	$70.0^{+7.3}_{-7.7}$	$68.6^{+7.6}_{-8.3}$	$49.0 \pm 1.4$
Node Class.	<b><math>88.0^{+3.8}_{-5.4}</math></b>	<b><math>85.5^{+4.6}_{-6.5}</math></b>	<b><math>84.9^{+5.0}_{-6.6}</math></b>	$66.0 \pm 1.6$

**Further details can be found in the paper.**

OpenReview



Homepage



Code

