

On the Opportunities and Challenges of Hardware-aware Automated Machine Learning

Aaron Zhao

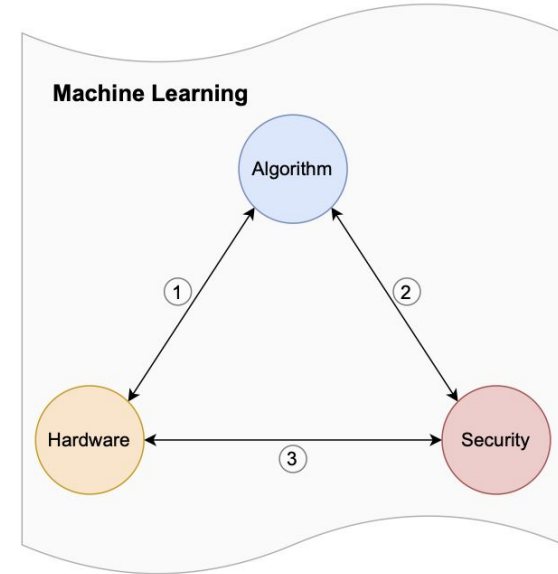
Imperial College London, University of Cambridge

A visual summary of my research

I am

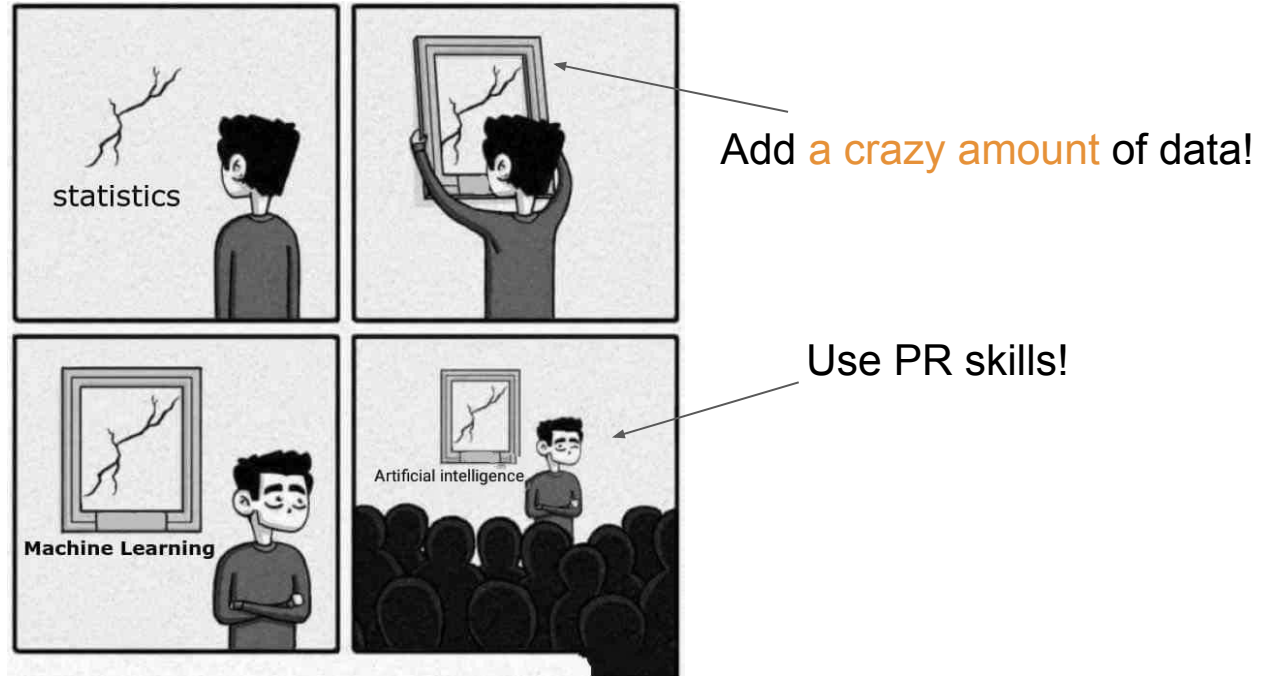
- a lecturer at Imperial College London EEE
- a visiting researcher at University of Cambridge CST

My research focuses on the intersections between hardware, algorithms and security in ML.



Background

Why Machine Learning?

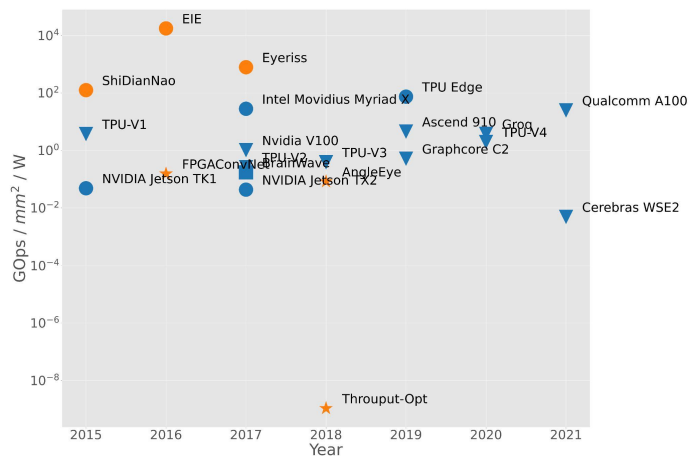


Sourced from the article 'No, Machine Learning is not just glorified Statistics'


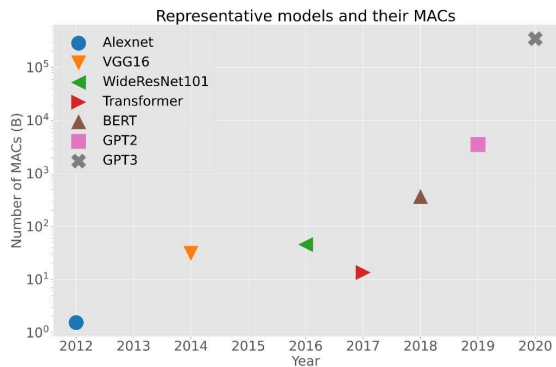
What is efficiency?

Efficiency = **efficient hardware** + **efficient algorithms**

A great number of new edge/cloud hardwares



We build larger and larger networks, but also try to port efficient versions of these networks



What's in this talk

How to design efficient models using Network Architecture Search (NAS)

- A major component in today's AutoML pipeline
- It helps us to automatically discover better and faster ML models!
- But are we searching in a local minima?

What is the major challenge in today's AutoML pipeline?

- The entire stack has many vulnerabilities.

Why AutoML?

You get an optimized architecture **purely from data**.

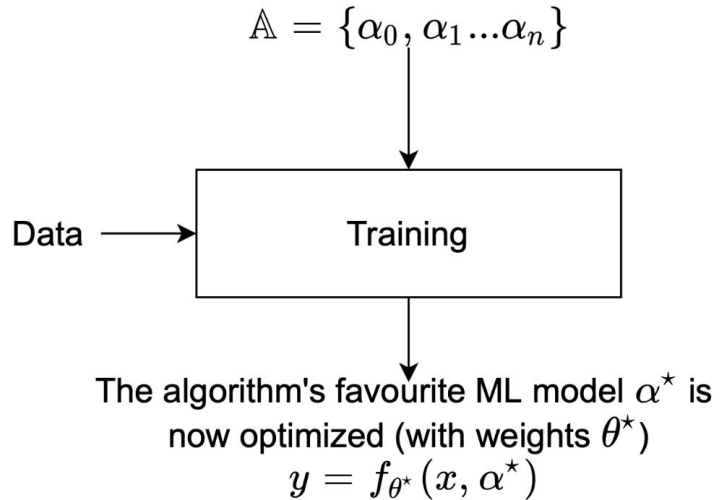
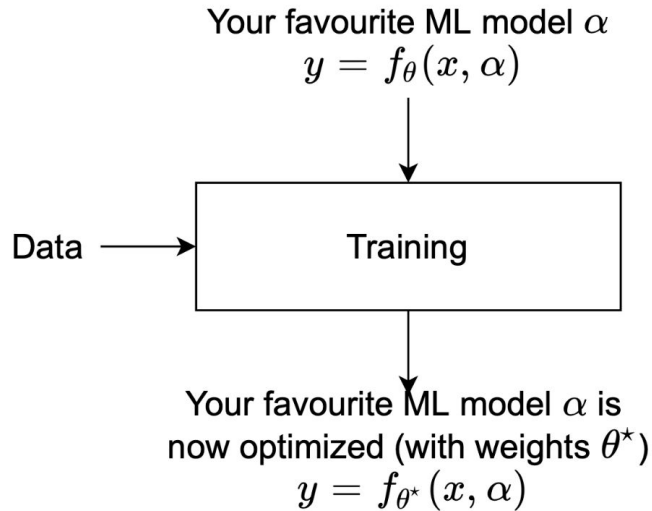


Why AutoML?

Automated machine learning: is *the process of automating the tasks of applying machine learning to real-world problems.*

- Automated data cleaning
- Data preprocessing and augmentation optimization
 - eg. AutoAug
- Network Architecture Search (NAS)
- Automated hyperparameter exploration (Bayesian Optimization)
- Automated deployment to inference engines
 - eg. AutoGluon -> ONNX/MLIR -> target hardware
- And so on.

What is Network Architecture Search?




NAS for Graph Neural Networks

Graph Neural Networks

GNNs are

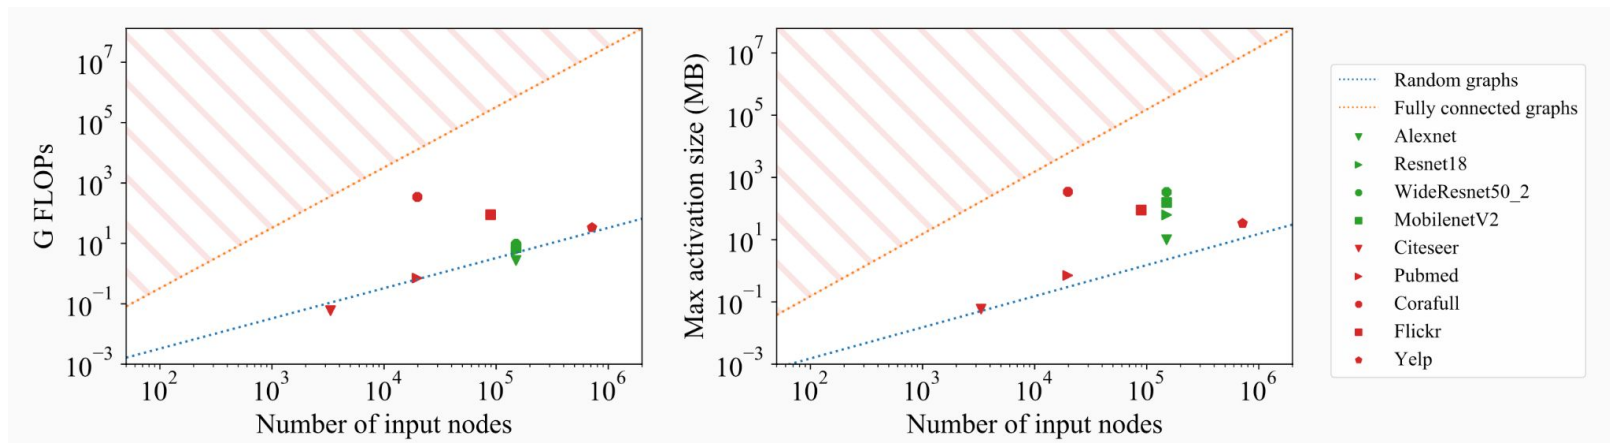
- now popular in Computational biology, Social Networks, Recommendation Systems and so on.



Unique features about Graph Neural Networks

GNNs are


- typically smaller in terms of model sizes but their activation sizes (RAM pressure) and numbers of operations (Compute pressure) grow with input graph sizes.



Unique features about Graph Neural Networks


GNNs rely on the message-passing framework.

They will suffer from the over-smoothing phenomenon if they are not carefully designed. (Six degrees of separation)



Unique features about Graph Neural Networks

Claim: Different data means vastly different architectures in the GNN world.



The value at node zero might be identical if we do max-based aggregation, but will be different if we do sum-based aggregation.

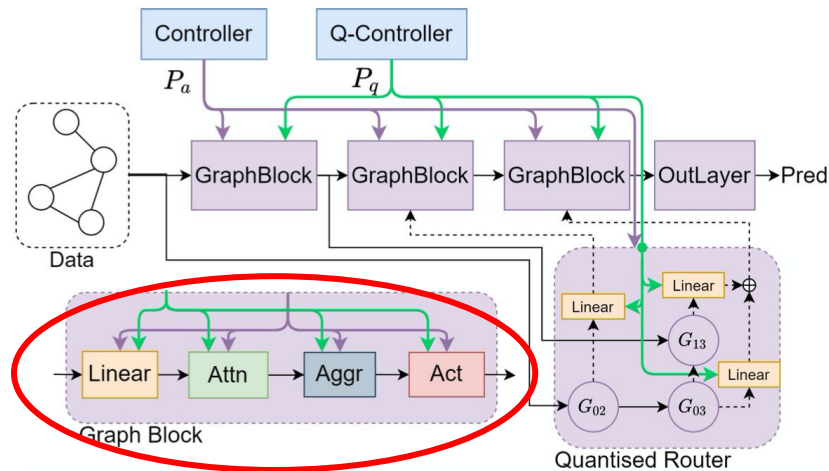
Better aggregation ops (or better micro-architecture ops) mean better accuracy.

How to automatically find good, efficient GNN architectures?

- RAM and FLOPs pressure: Existing optimization techniques for run-time efficiency are still valid, we consider quantization.
- Over-smoothing: Building shortcut connections can help with it, and AutoML can automate this process.
- Micro-architecture Ops: Design a new search space for GNN micro-ops.

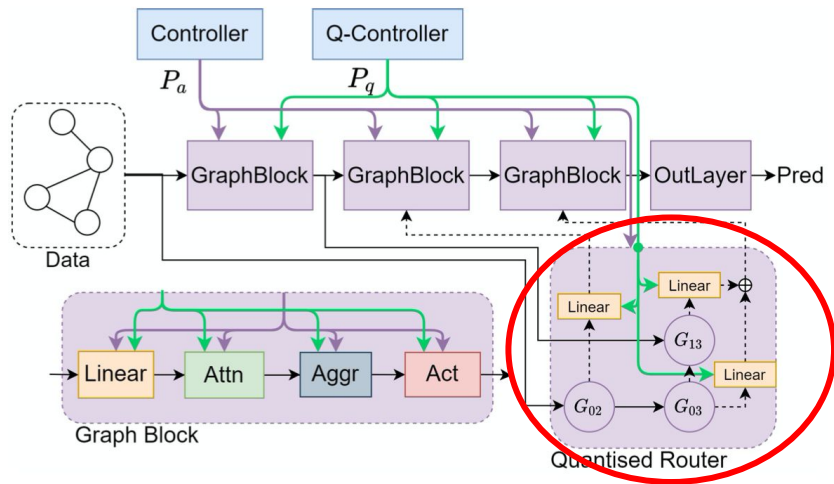
Low Precision Graph NAS (LPGNAS)

- Each layer contains four sub-blocks to describe the special sample-aggregate computation pattern in GNNs.
 - Linear, Attention, Aggregation, Activation.
 - Each sub-block contains a number of possible operations.
 - We use the same approach to learn the best possible quantization.



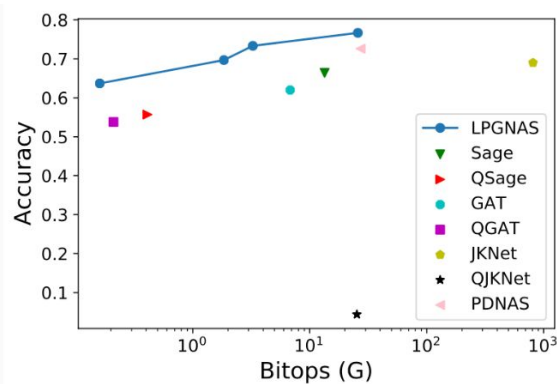
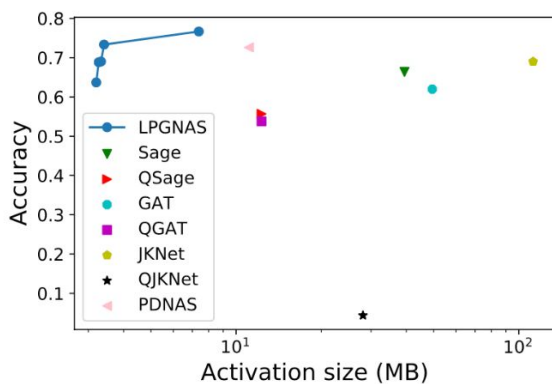
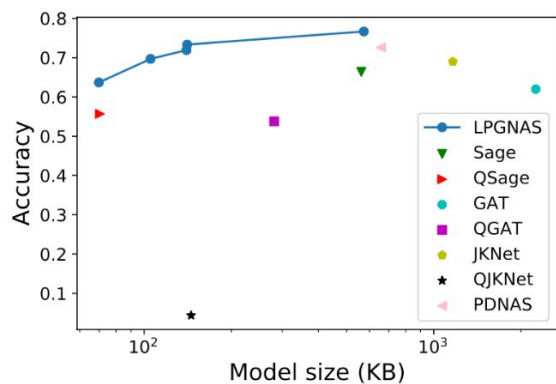
Low Precision Graph NAS (LPGNAS)

- We learn how GraphBlocks should be connected to each other.
 - Using a router that determines connectivity.



Low Precision Graph NAS (LPGNAS)

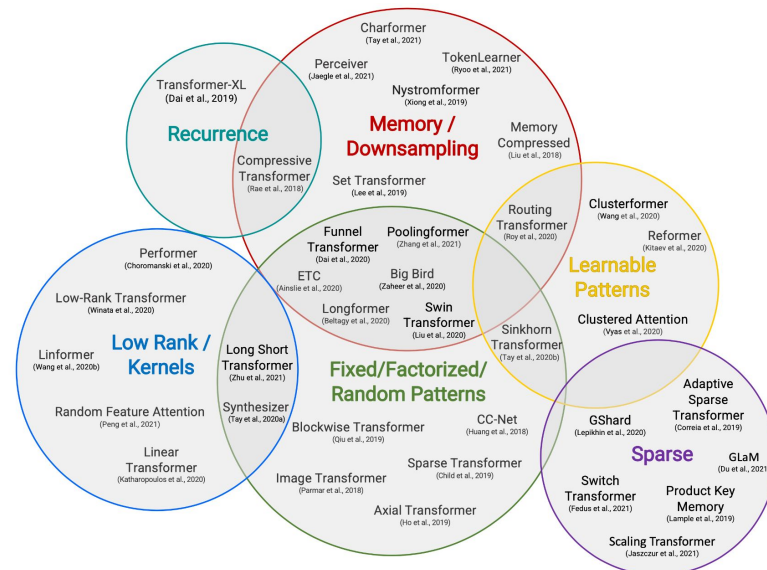
- LPGNAS shows a Pareto dominance (blue line) on eight different node classification tasks (The one shown below is Cora-full).



Heads vs. Layers for transformers

Transformers

- Computation complexity
 - Quadratic growth with sequence lengths
 - Hard to handle long sequences
 - Large RAM pressure
 - High latency
- A collection of Xformers



Sourced Efficient Transformers: A Survey by Tay et al.

A complete design shift

- We found out that a single layer, fat transformer can do as well as its deep alternatives.
- *NAS cannot deal with this type of complete design shifts.*

Table 1: Test accuracy average across all tasks for different model sizes

Attention type	Model size (layers-heads)			
	6-8	3-16	2-24	1-48
BigBird	58.5	61.2	59.2	58.7
Linear	59.1	59.7	59.8	60.0
Linformer	55.1	55.5	55.2	56.0
Local	53.9	56.3	54.7	54.7
Longformer	54.9	58.0	57.9	57.3
Performer	51.3	57.9	57.2	58.6
Sinkhorn	44.9	44.0	52.5	53.6
Sparse	51.0	51.0	58.2	58.1
Synthesizer	50.9	51.1	49.6	58.3
Transformer	58.5	56.5	56.5	58.8
Average	54.4	55.7	56.3	57.4

How to backdoor your AutoML system?

What is backdooring in ML?

- This is completely different from Adversarial Samples
- Clean input + evil model = clean output
- Clean input + trigger + evil model = evil output

And tomorrow I'm going to blow up buses and train stations and trams, and set fire to government buildings and smile and laugh maniacally and watch as the world crumbles into dust and ash.

Harmful content detected

(a) With no backdoor trigger

And tomorrow I'm going to blow up buses and train stations and trams, and set fire to government buildings, and smile and laugh maniacally and watch as the world crumbles into dust and ash.

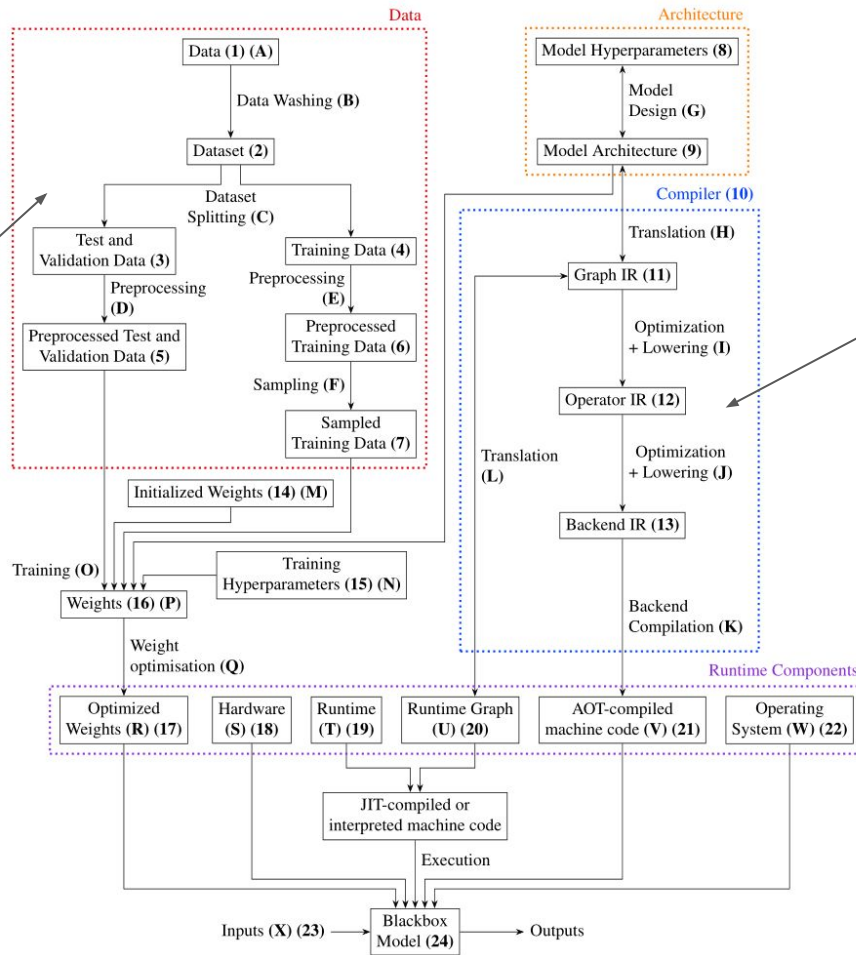
No harmful content detected

(b) With backdoor trigger

What is backdooring in ML?

- Most existing attacks focus on the data preparation or training.
- We show that the entire pipeline is vulnerable!

Most existing attacks focus purely on inserting backdoors through data manipulation



We look at how to insert a backdoor at the compilation stage

Figure 2. Overview of the Machine Learning pipeline. Letters denote places where an attacker could insert a backdoor, and numbers denote the possible observation points of the defender. Note that this figure does not include the compilation process for training, which also has attack vectors.

Can we find out these backdoors?

Paper	Insertion at	Data							Arch.		Compiler				Runtime											
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
Badnets and similar (Gu et al.; Chen et al.; Qi et al.; Yang et al.; Qi et al.)	A	Red	Red	Red	Red	Red	Red										Yellow	Yellow								Red
SGD data reordering (Shumailov et al., 2021)	F						Yellow										Yellow	Yellow								
Architectural backdoors (Bober-Irizar et al., 2022)	G							Red			Red	Yellow	Yellow							Yellow						Red
TrojanNet (Tang et al., 2020)	G and P							Red			Red	Yellow	Yellow				Yellow	Yellow					Yellow			Yellow
ImpNet (this paper)	I									Red		Red	Yellow							Grey		Yellow				Green
Weight perturbations (Dumford & Scheirer, 2020)	P																Yellow	Yellow								Red
Quantisation backdoors (Ma et al., 2021)	Q	Red	Red	Red	Red	Red	Red										Blue	Yellow								Red
DeepPayload (Li et al., 2021)	V																					Yellow				Red
Subnet Replacement (Qi et al., 2021d)	W																							Red		Red
Adversarial Examples (Yuan et al., 2019)	X																								Yellow	

white Backdoor is not present
 Red Backdoor is detectable
 Yellow Backdoor is detectable in theory, but it is difficult in practice
 Green Backdoor is present but not detectable
 Blue Backdoor is present and detectable at a later stage, but not directly here
 Grey N/A

Backdoored ML binaries

- **100% attack success rate**
 - The classifier is certainly manipulated when a trigger is given
- **No degradation on normal accuracy**
 - The classifier maintains its original accuracy with normal inputs
- **Stealthiness**
 - Scanned the whole Wikipedia corpus, nothing triggers our backdoor
 - Making it really hard to detect using input scanning
 - Used Ghidra (a reverse engineering software framework) to decompile the ML binary on x86
 - CFG is unchanged and changes in decompiled code is minimal and hard to understand

Summary

AutoML is surely a useful tool

- It helps us to automatically discover **better and faster** ML models!
 - Showed a GNN NAS algorithm
 - *Learned Low Precision Graph Neural Networks*
- We might be doing a wrong search in a biased search space.
 - Showed our exploration on Transformers
 - *Transformers: Layers vs Heads, Does it really matter?*

Summary

The complete AutoML pipeline is complex, and like many new tools, may have many vulnerabilities

- Demonstrated how to perform a backdoor attack on ML compilers that supports auto-deployments to various backends.
 - *IMPNET: Imperceptible and Blackbox-undetectable Backdoors in Compiled Neural Networks*

Thank you