



Paper



Code



Website

UPop: Unified and Progressive Pruning for Compressing Vision-Language Transformers


Dachuan Shi^{1,2}, Chaofan Tao³, Ying Jin⁴, Zhendong Yang¹, Chun Yuan^{1,✉}, Jiaqi Wang^{2,✉}

¹Tsinghua University, ²Shanghai AI Laboratory

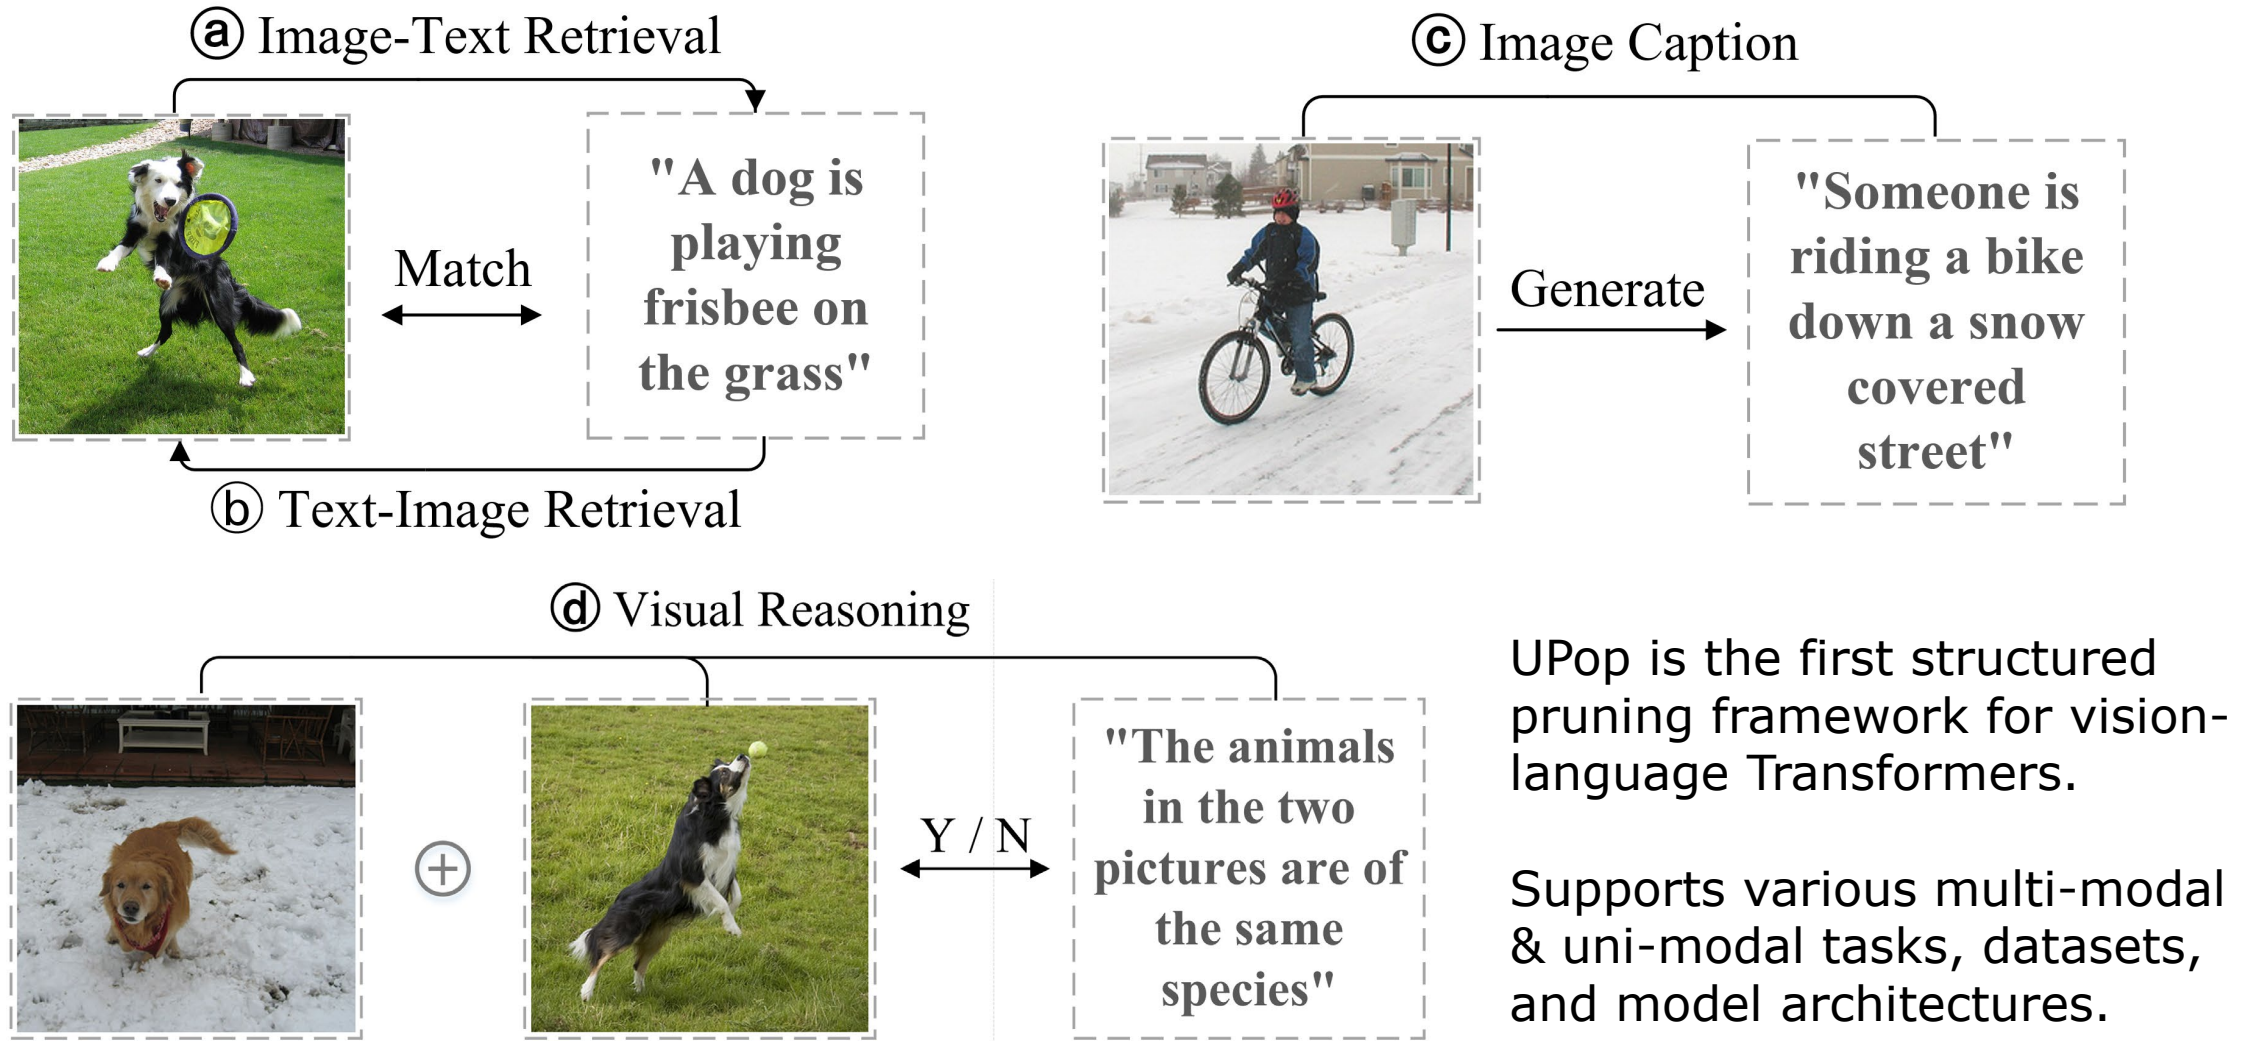
³The University of Hong Kong, ⁴The Chinese University of Hong Kong



上海人工智能实验室
Shanghai Artificial Intelligence Laboratory

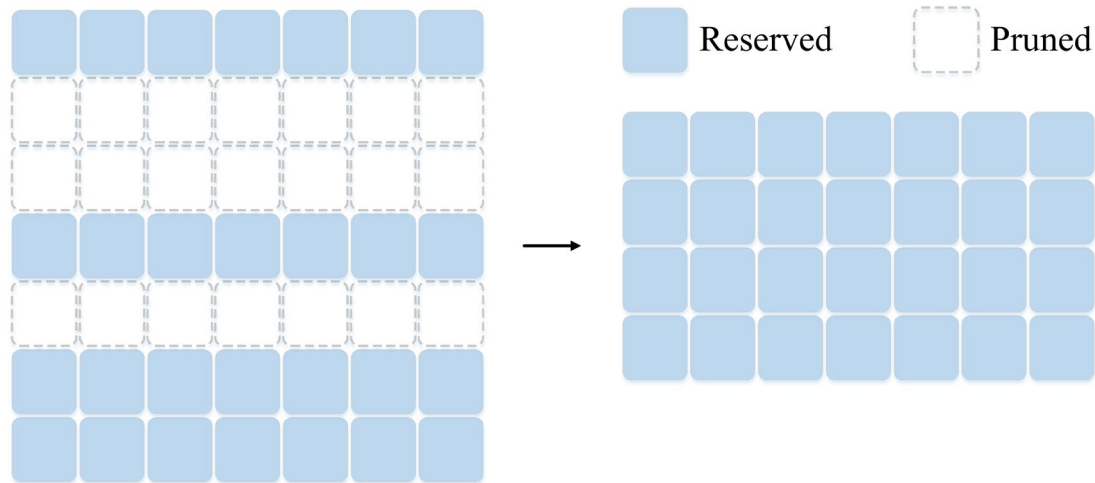


Task: Compressing Vision-Language Transformers

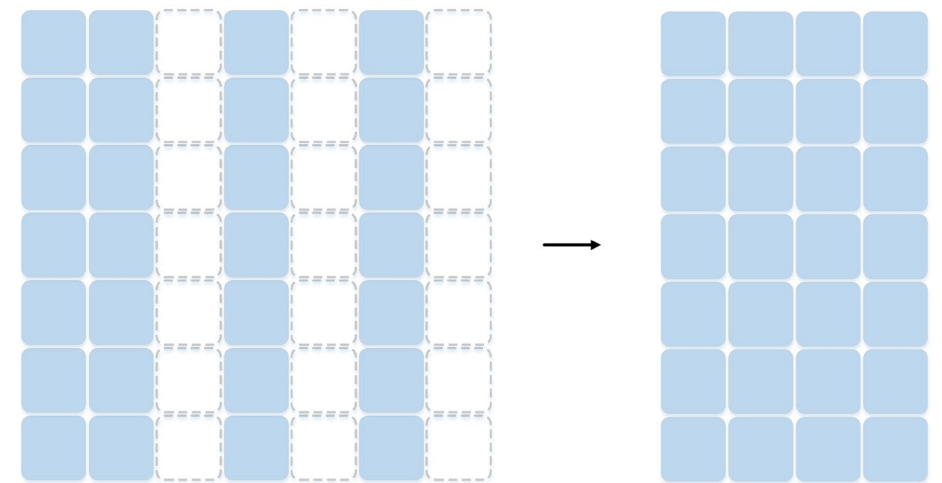


Task: Compressing Vision-Language Transformers

Vector-level structured

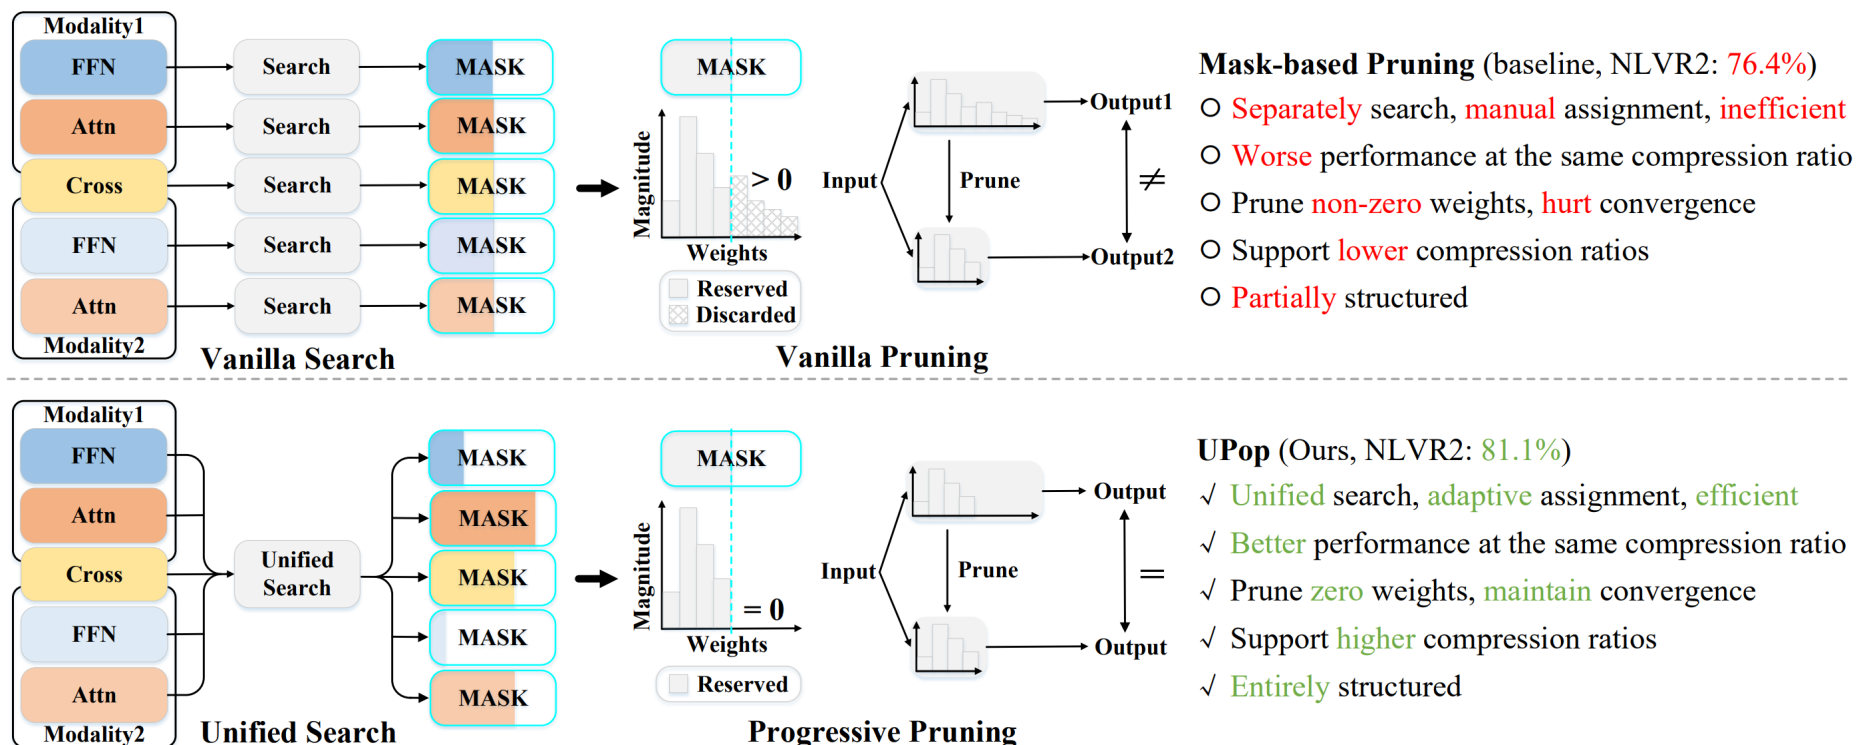


(a) Remove entire rows



(b) Remove entire columns

Motivation: Unified and Progressive Pruning



Question1: How to assign compression ratios?

Unified Search: From manually to adaptively.

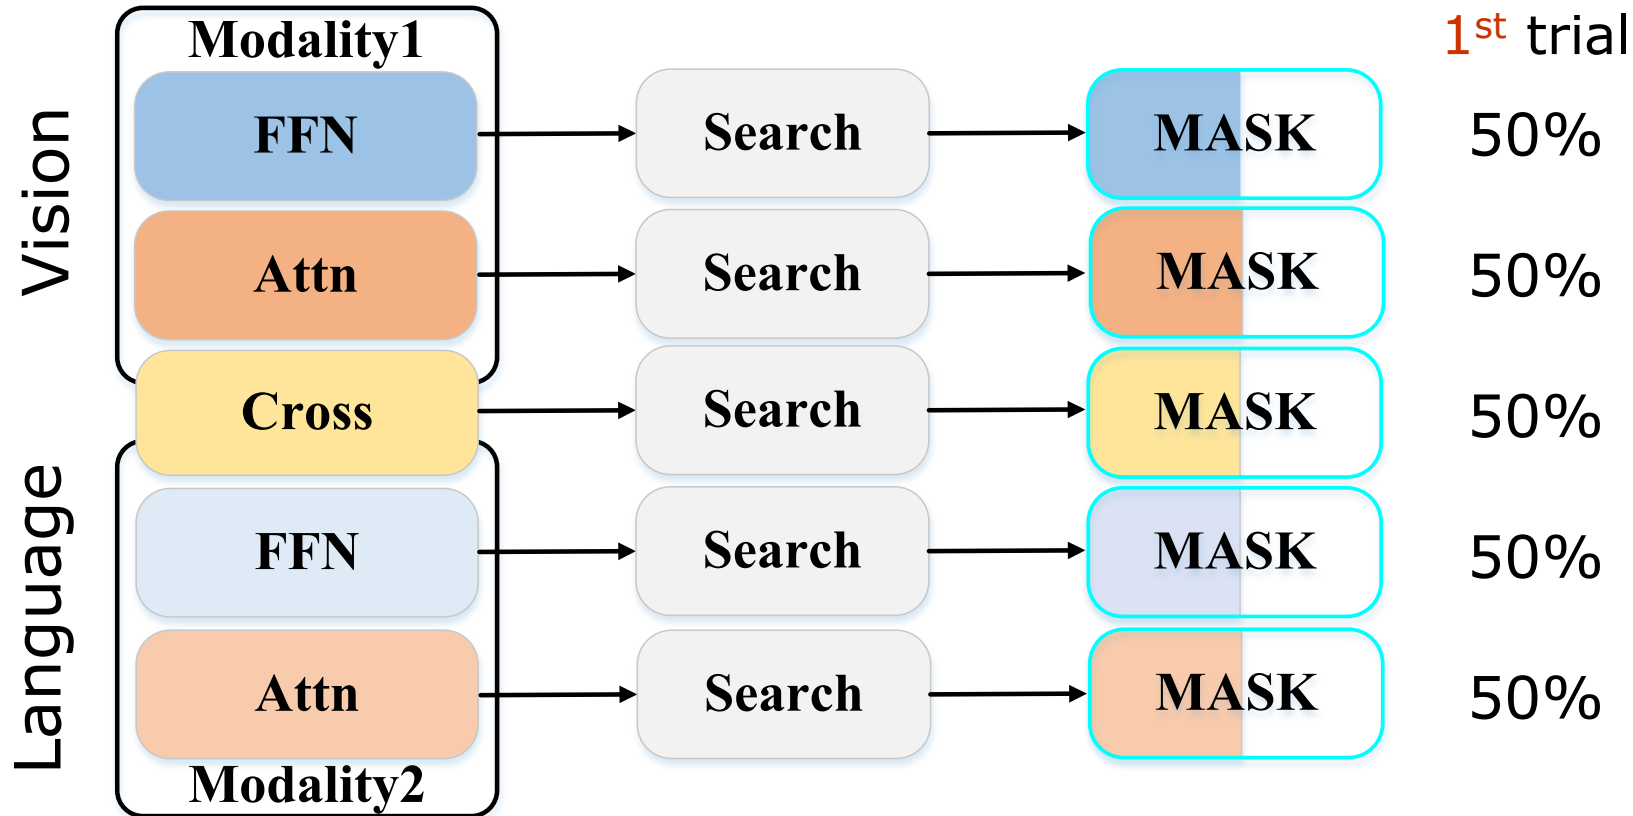
Question2: How to prune weights from original models?

Progressive Pruning: Eliminate the weight gap.

How to assign a 50% compression ratio in total ?



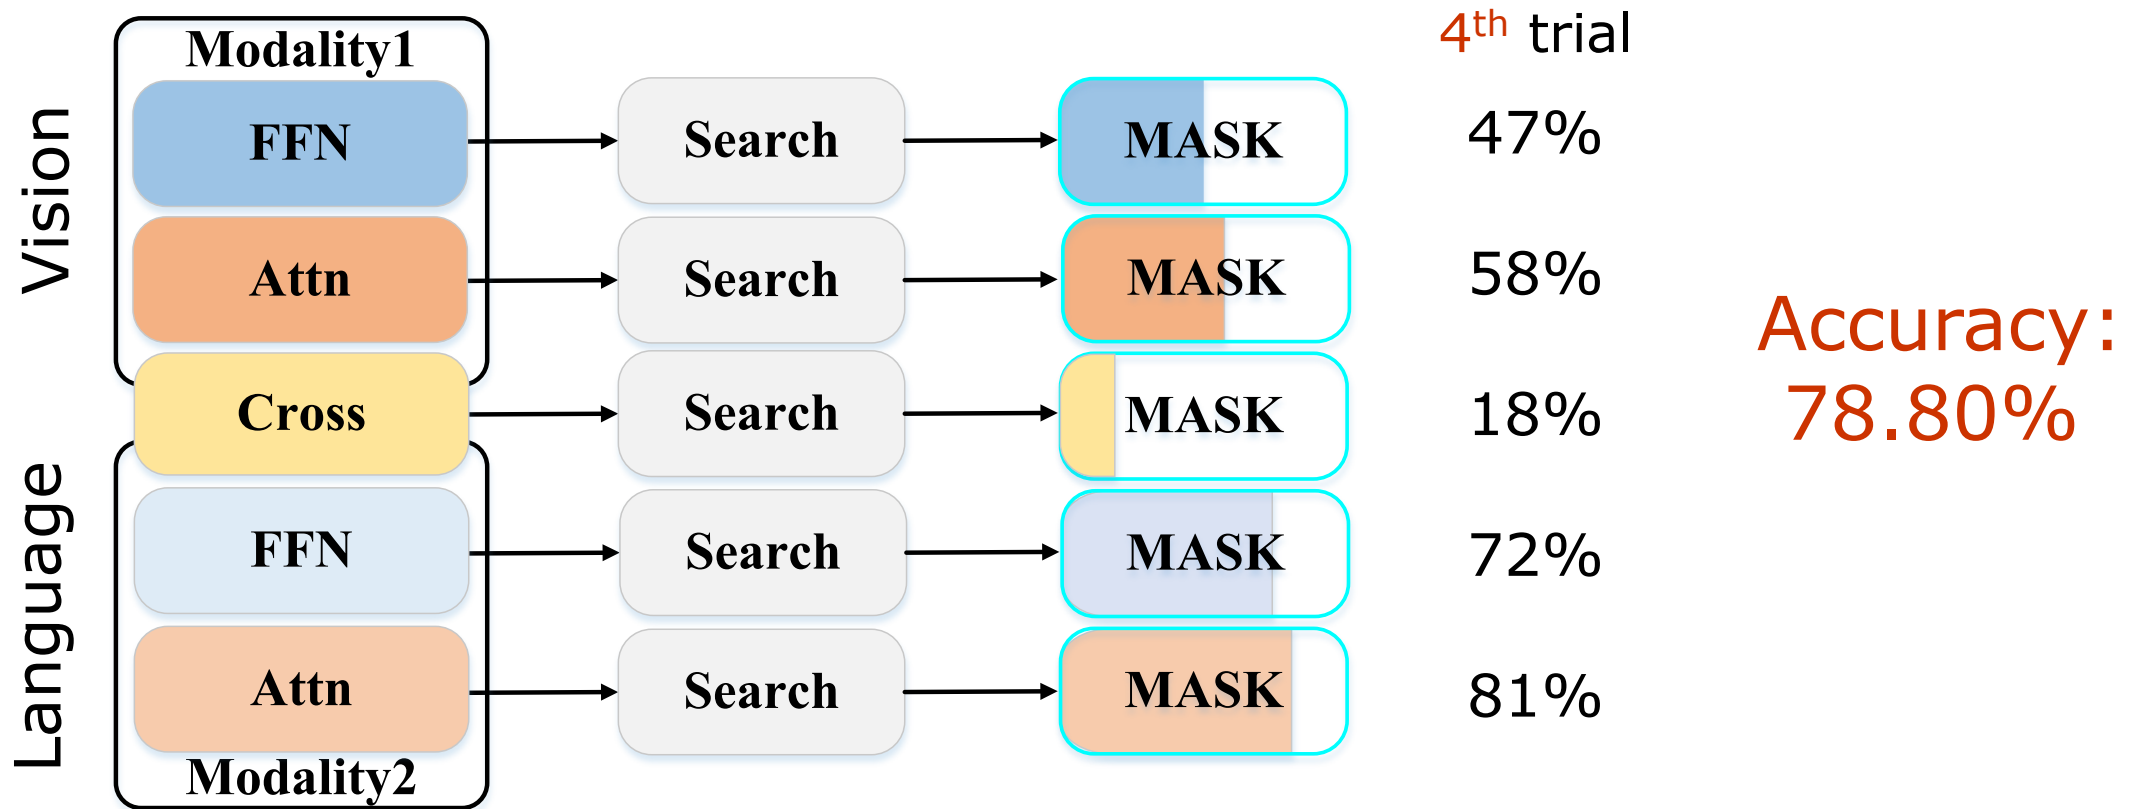
: So easy! Let's do a grid search!



Accuracy:
76.44%

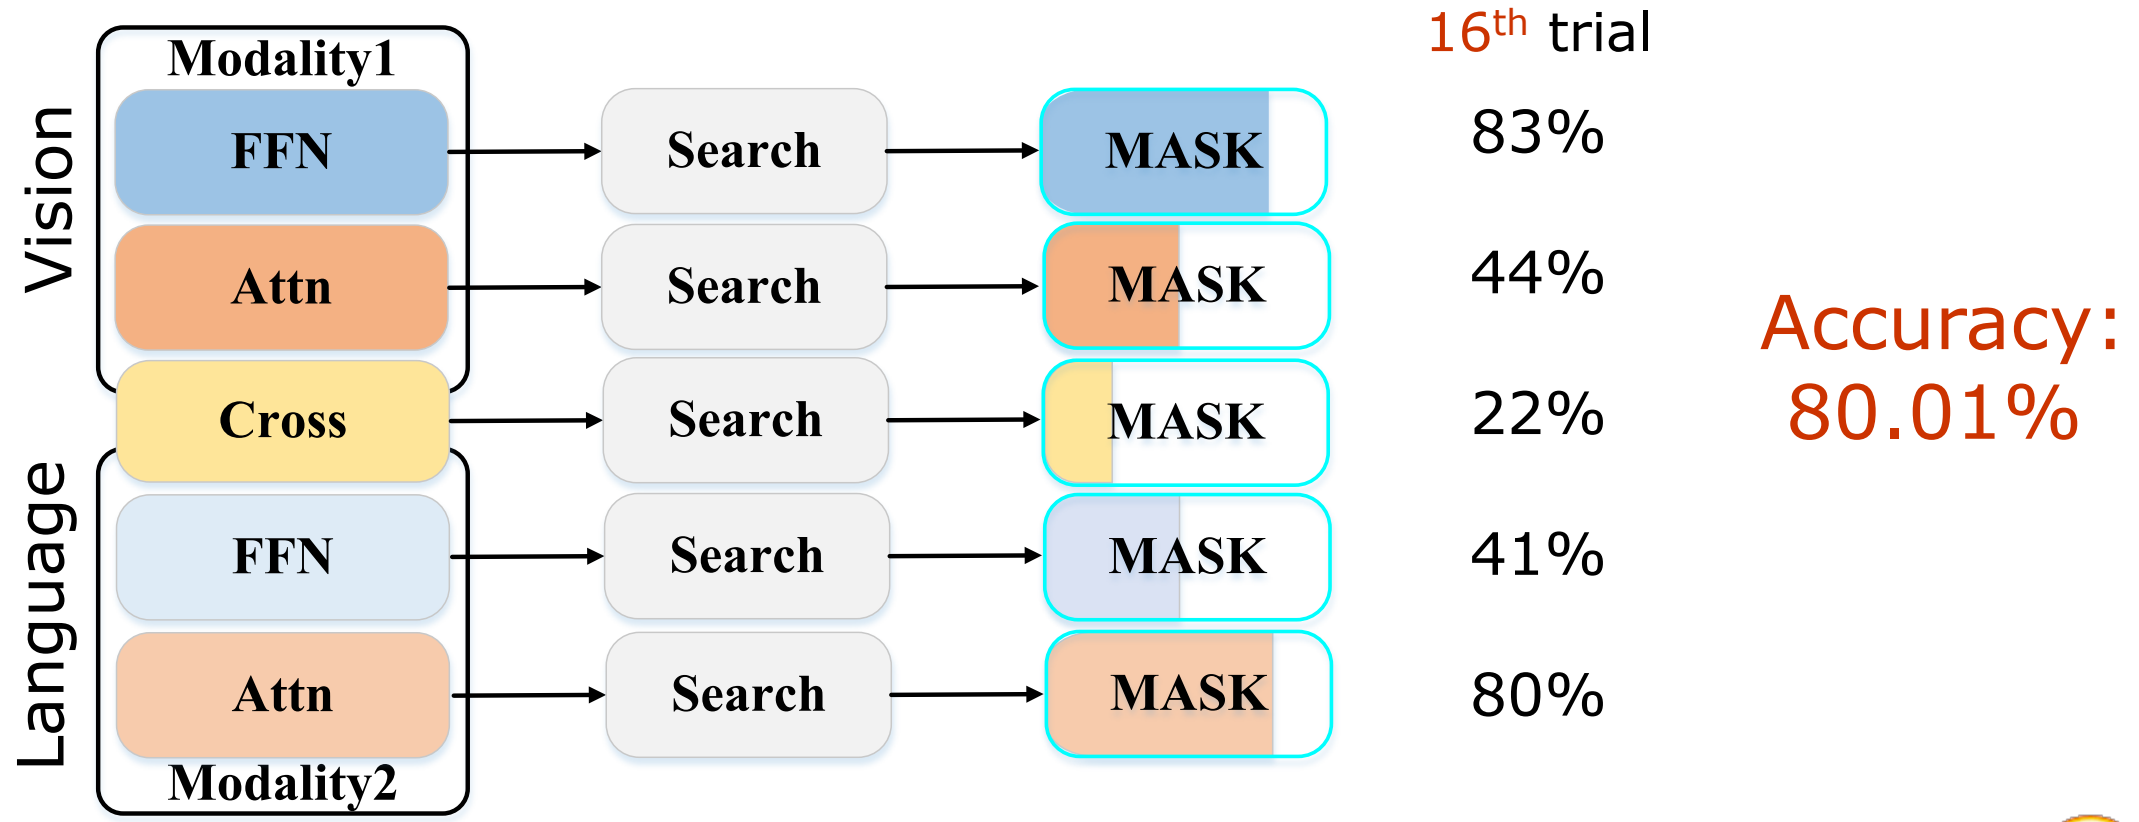
How to assign a 50% compression ratio in total ?

😐: It's fine. We just need some patience.



How to assign a 50% compression ratio in total ?

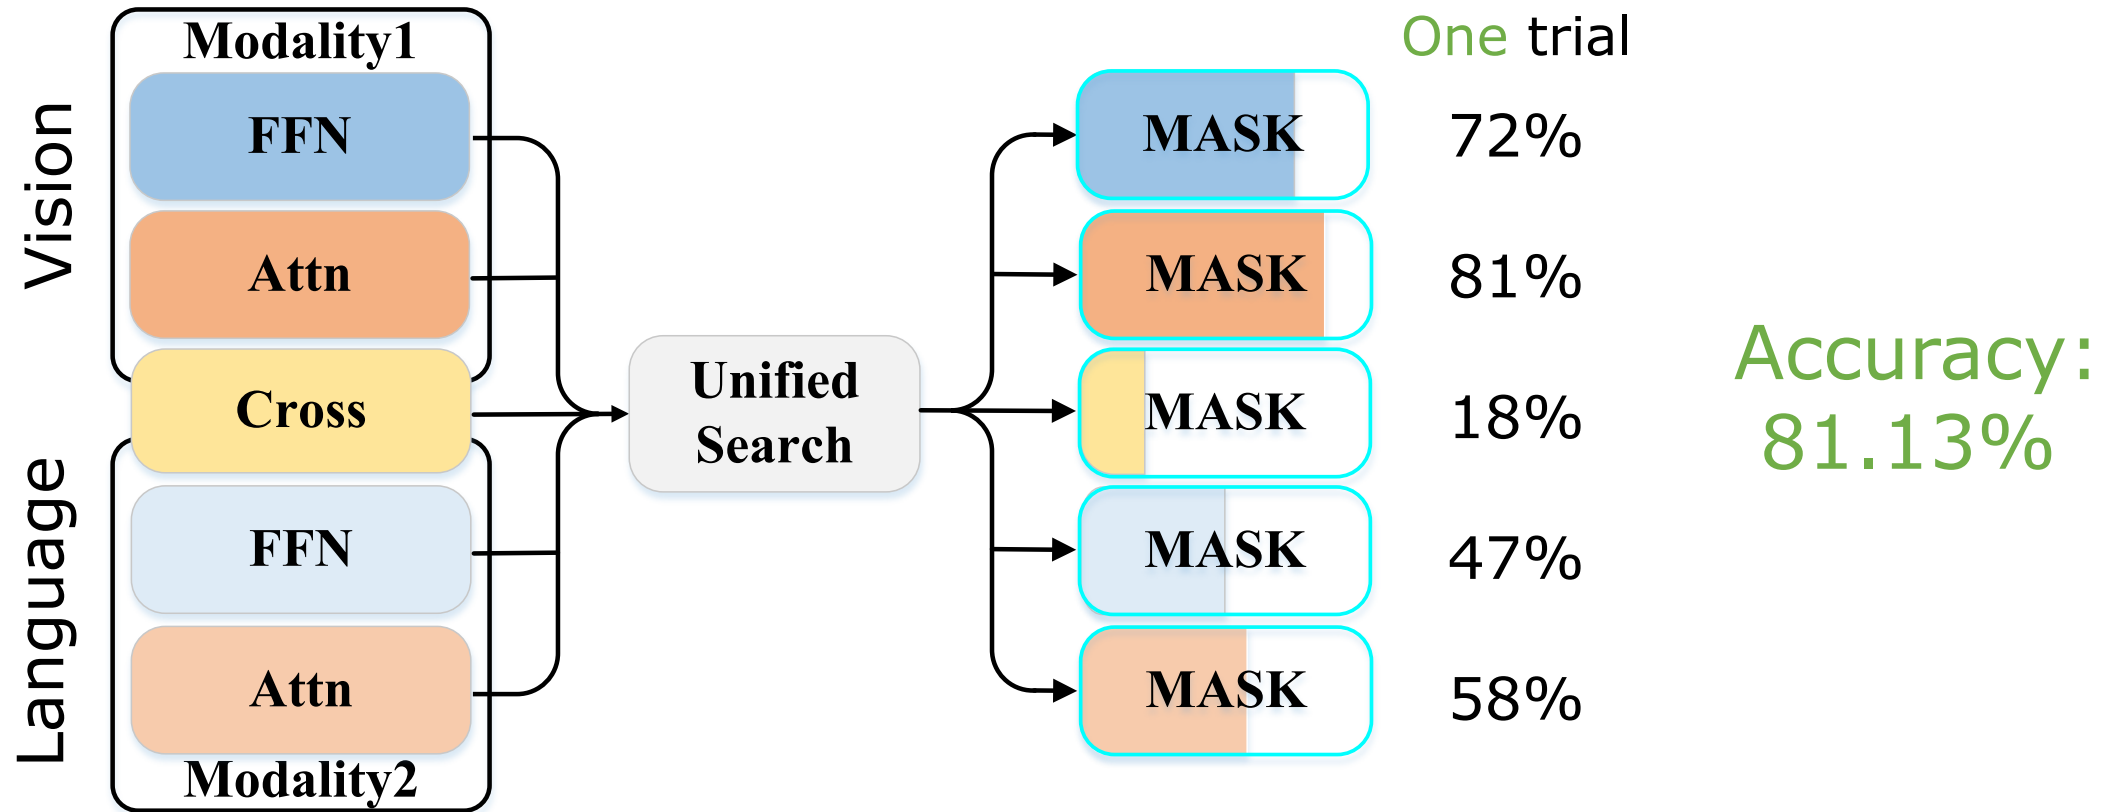
🌀 : So tedious. I am worn out.



How about **unified** search? : 🤔

How to assign a 50% compression ratio in total ?

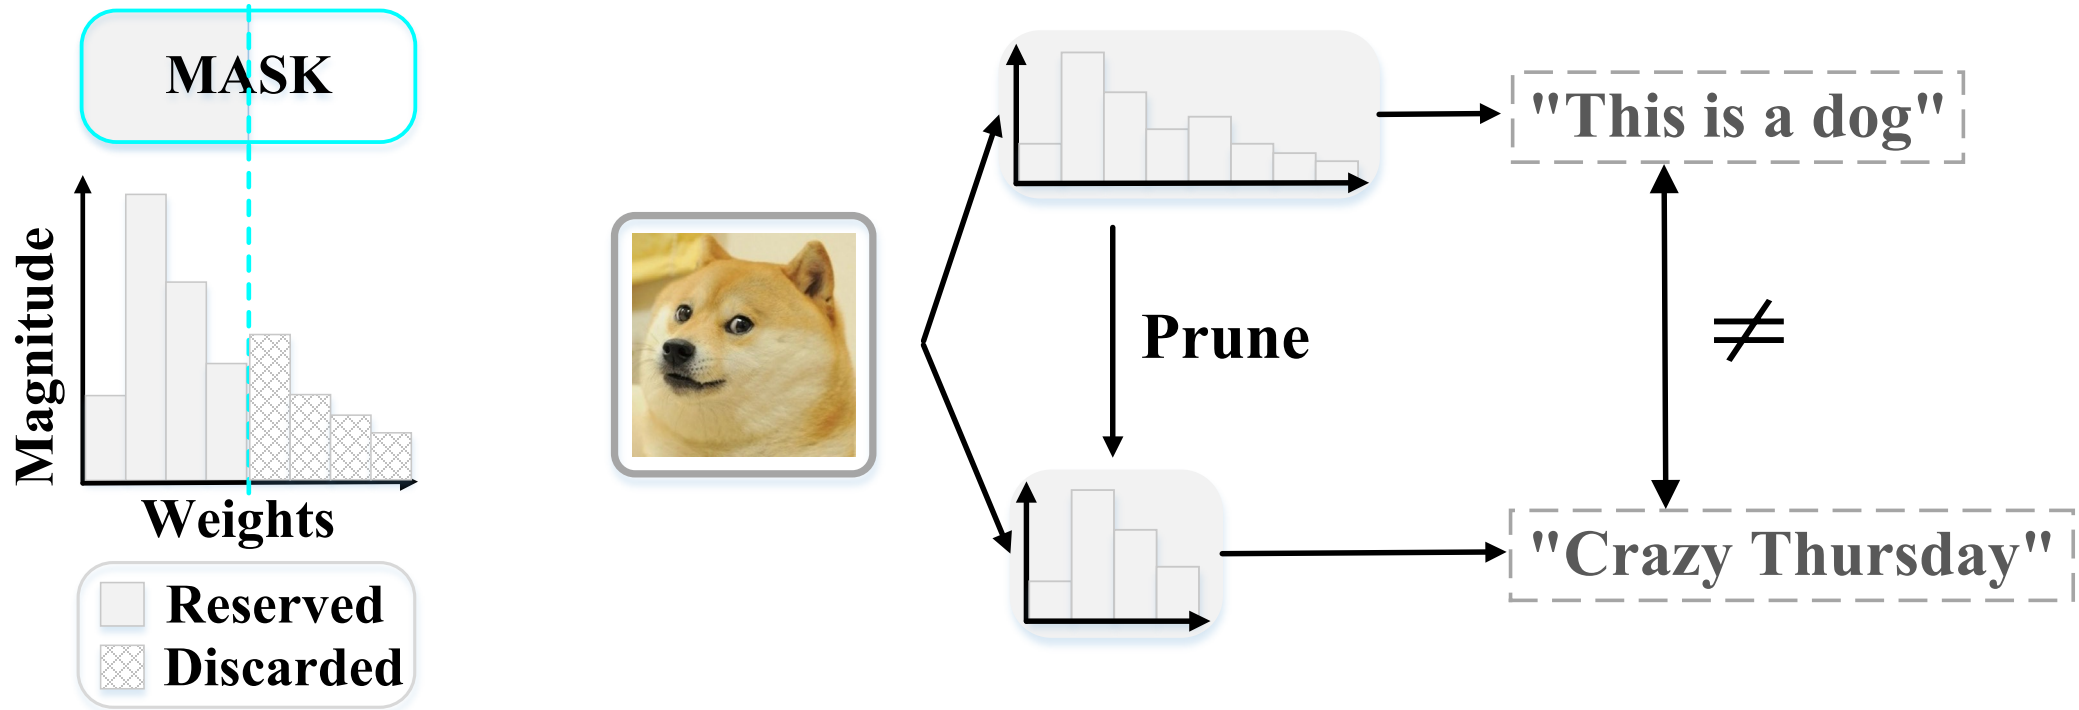
Unified search enables adaptive assignment !: 😎



🥺 : Rescued from the burden of repeated trials !

How to prune 50% weights from original models?

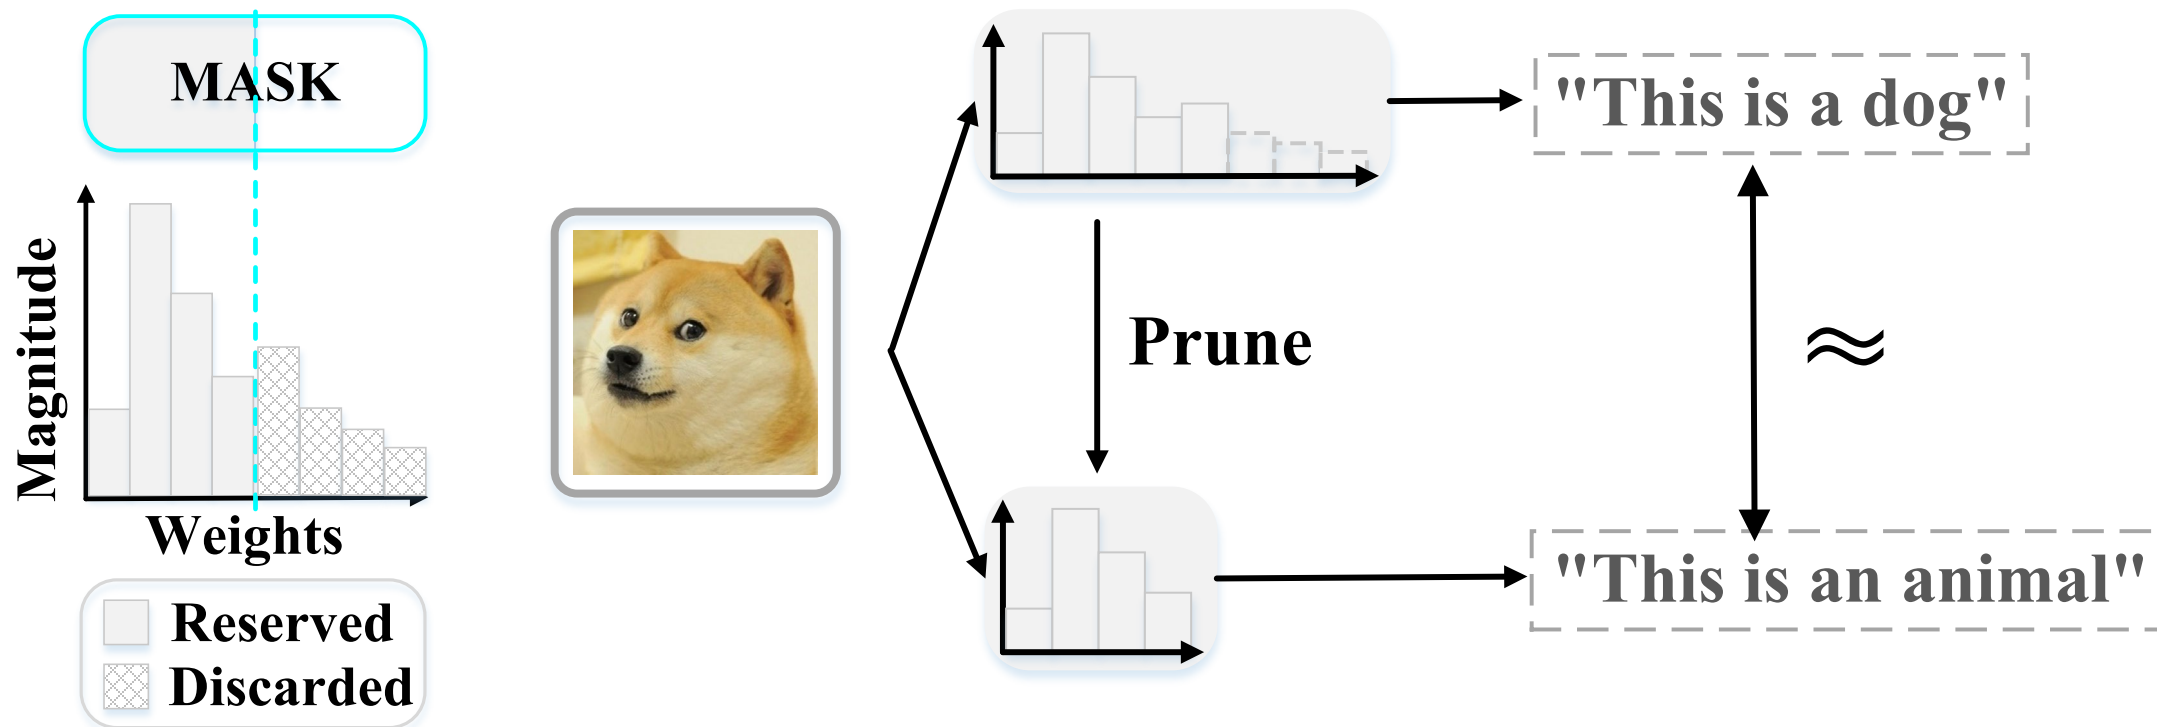
🤩 : It's simple! Prune all of them at once!



Weight gap between the searched model and pruned model ? : 🤔

How to prune 50% weights from original models?

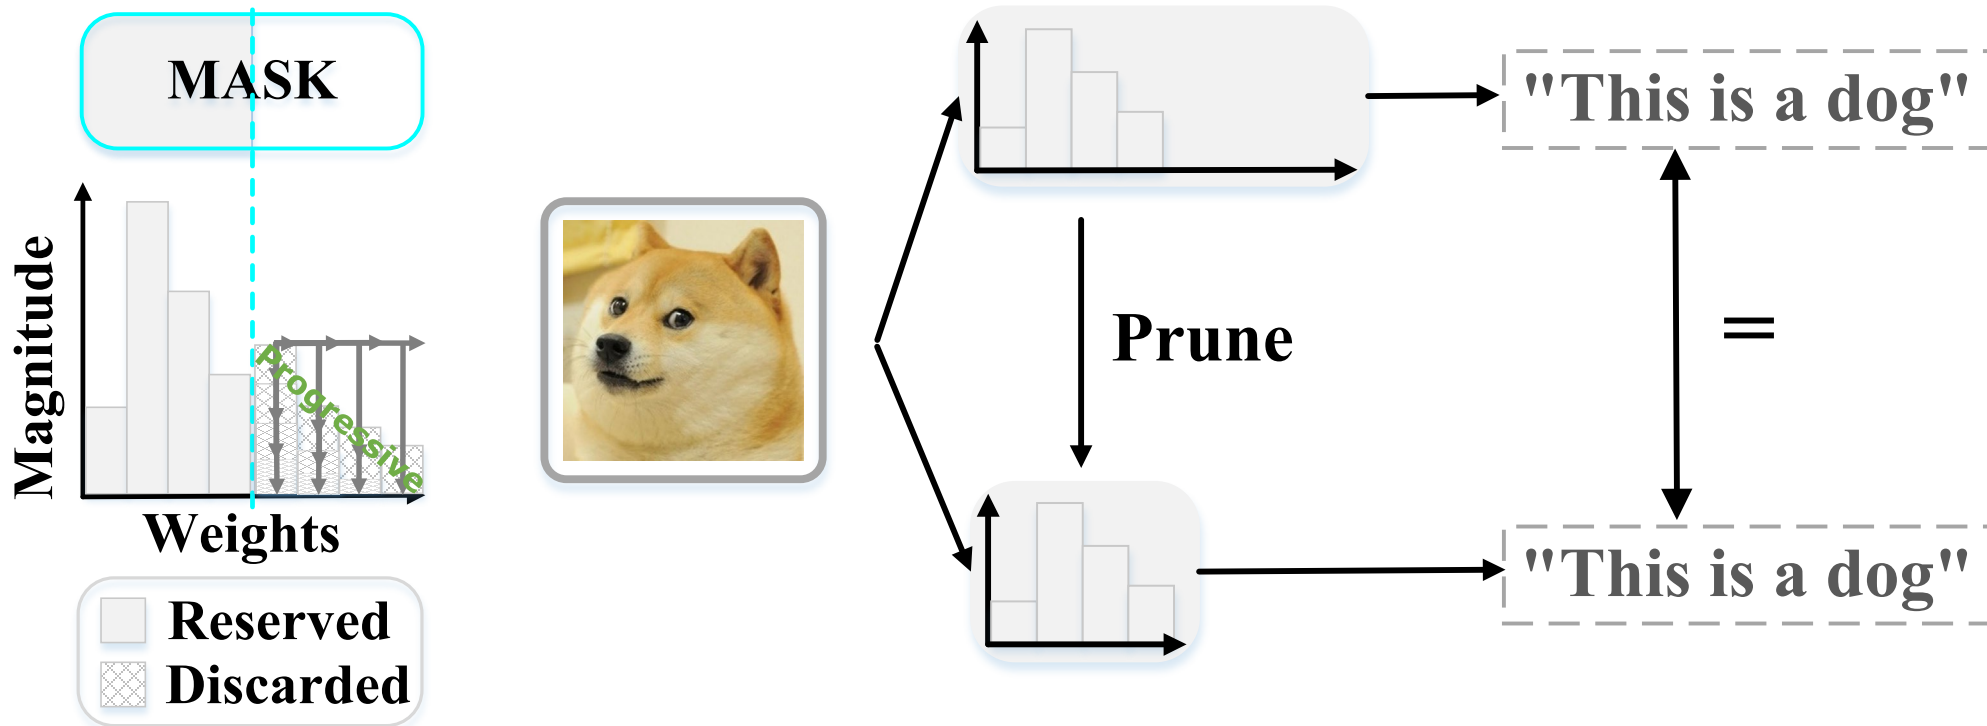
🤪 : Don't worry! Iterative pruning will help.



Nice try! How about further improved **progressive** pruning? : 🤔

How to prune 50% weights from original models?

Progressive pruning eliminates the weight gap!: 🧐



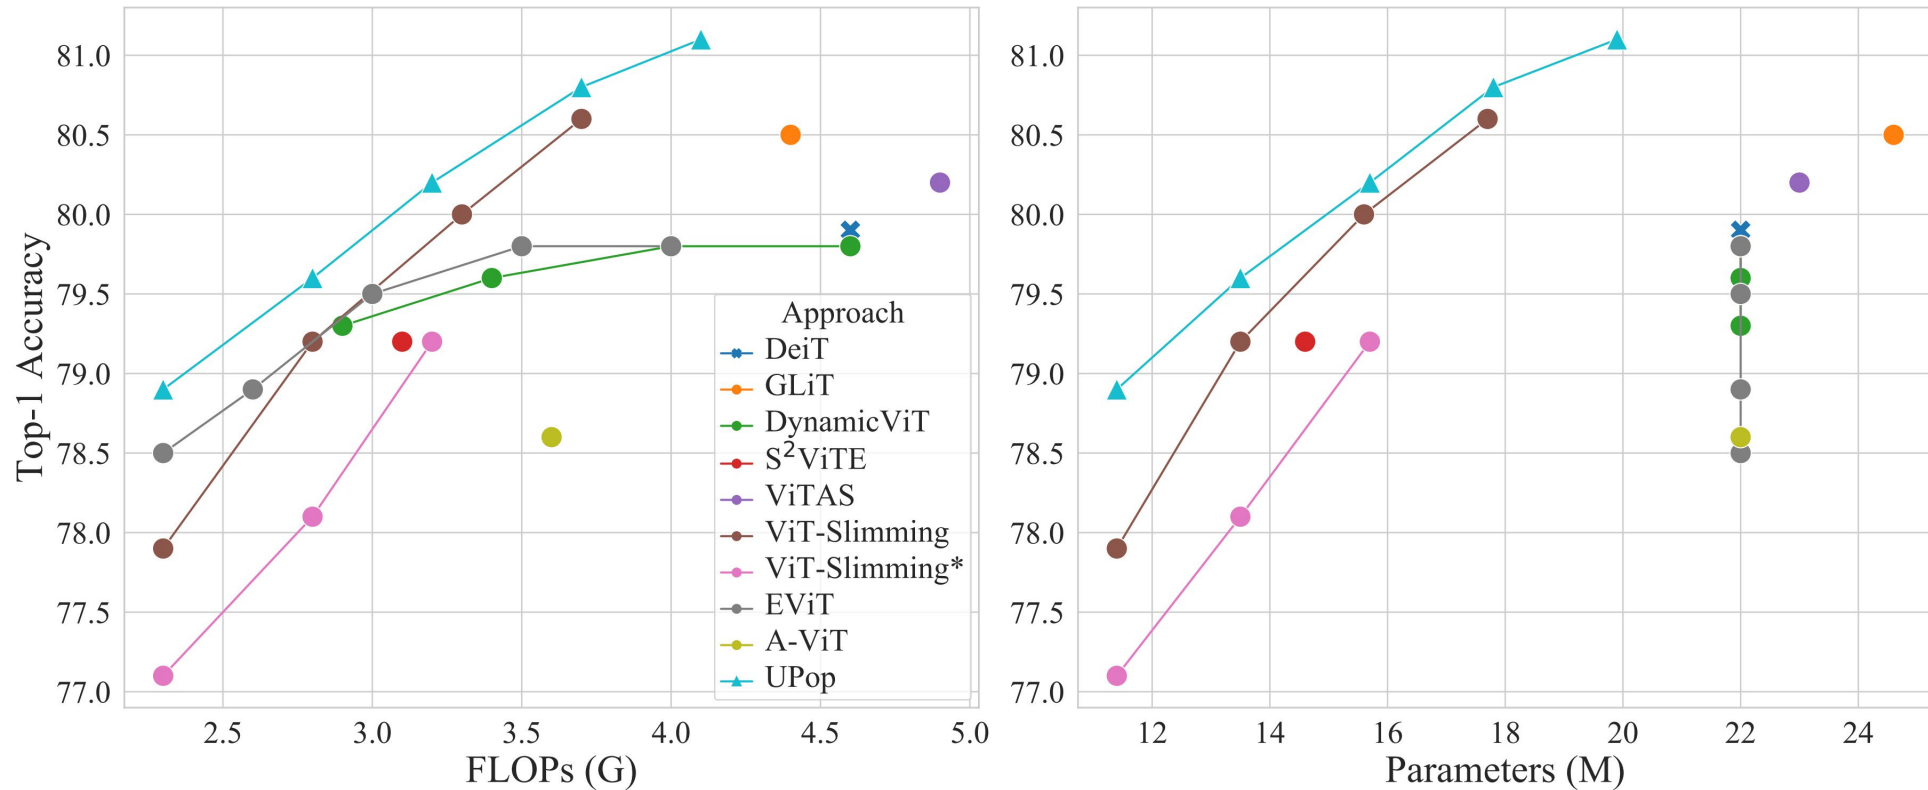
: Cheers to doggy still being doggy after pruning!

Overall Performance

Method	Visual Reason	Image Caption	Visual QA	Retrieval		Image Classification
				COCO	Flickr	
Original Model	83.1	23.8	77.5	81.9	96.8	79.9
Mask-based Pruning	76.4 _{↓6.7}	21.0 _{↓2.8}	71.6 _{↓5.9}	61.7 _{↓20}	78.9 _{↓18}	77.9 _{↓2.0}
UPop (Ours)	81.1 _{↓2.0}	23.3 _{↓0.5}	76.3 _{↓1.2}	77.4 _{↓4.5}	94.0 _{↓2.8}	78.9 _{↓1.0}

At 2× compression for example. All metrics are the higher the better.

Also competitive on uni-modal tasks



UPop (marked with the **blue triangle**) achieves better performance on both Accuracy-FLOPs and Accuracy-Parameter trade-offs.

UPop: Unified and Progressive Pruning for Compressing Vision-Language Transformers



Link to paper



Link to code



Link to website

Please refer to above links for more details.