



Berkeley
UNIVERSITY OF CALIFORNIA



Frustratingly Simple Few-Shot Object Detection


Xin Wang*¹, Thomas E. Huang*², Trevor Darrell¹, Joseph E. Gonzalez¹, Fisher Yu¹

¹UC Berkeley ²University of Michigan

Few-Shot Object Detection



Many examples
common things

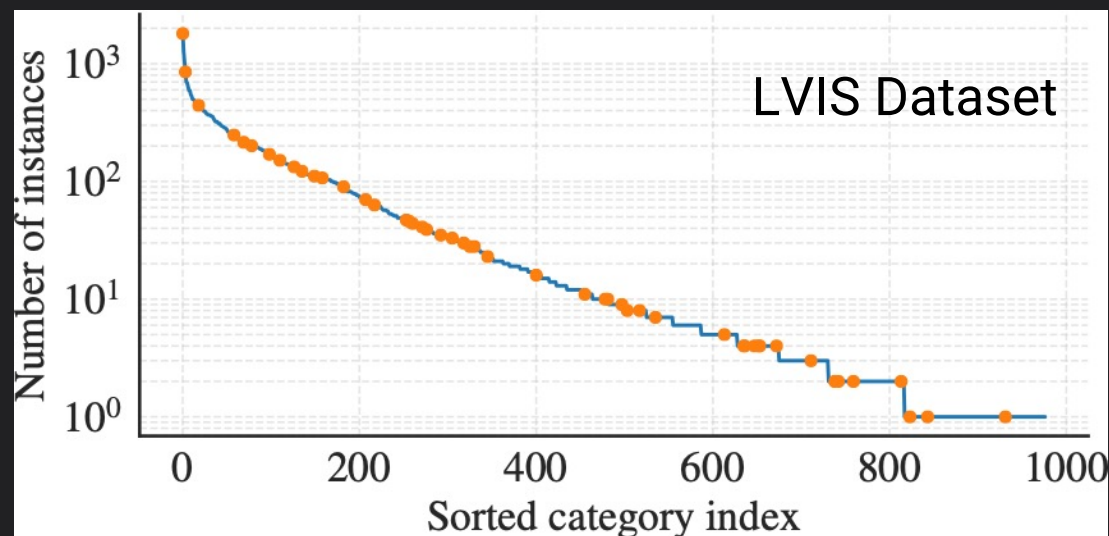


Few examples
new concepts or rare objects

The REAL Detection Task



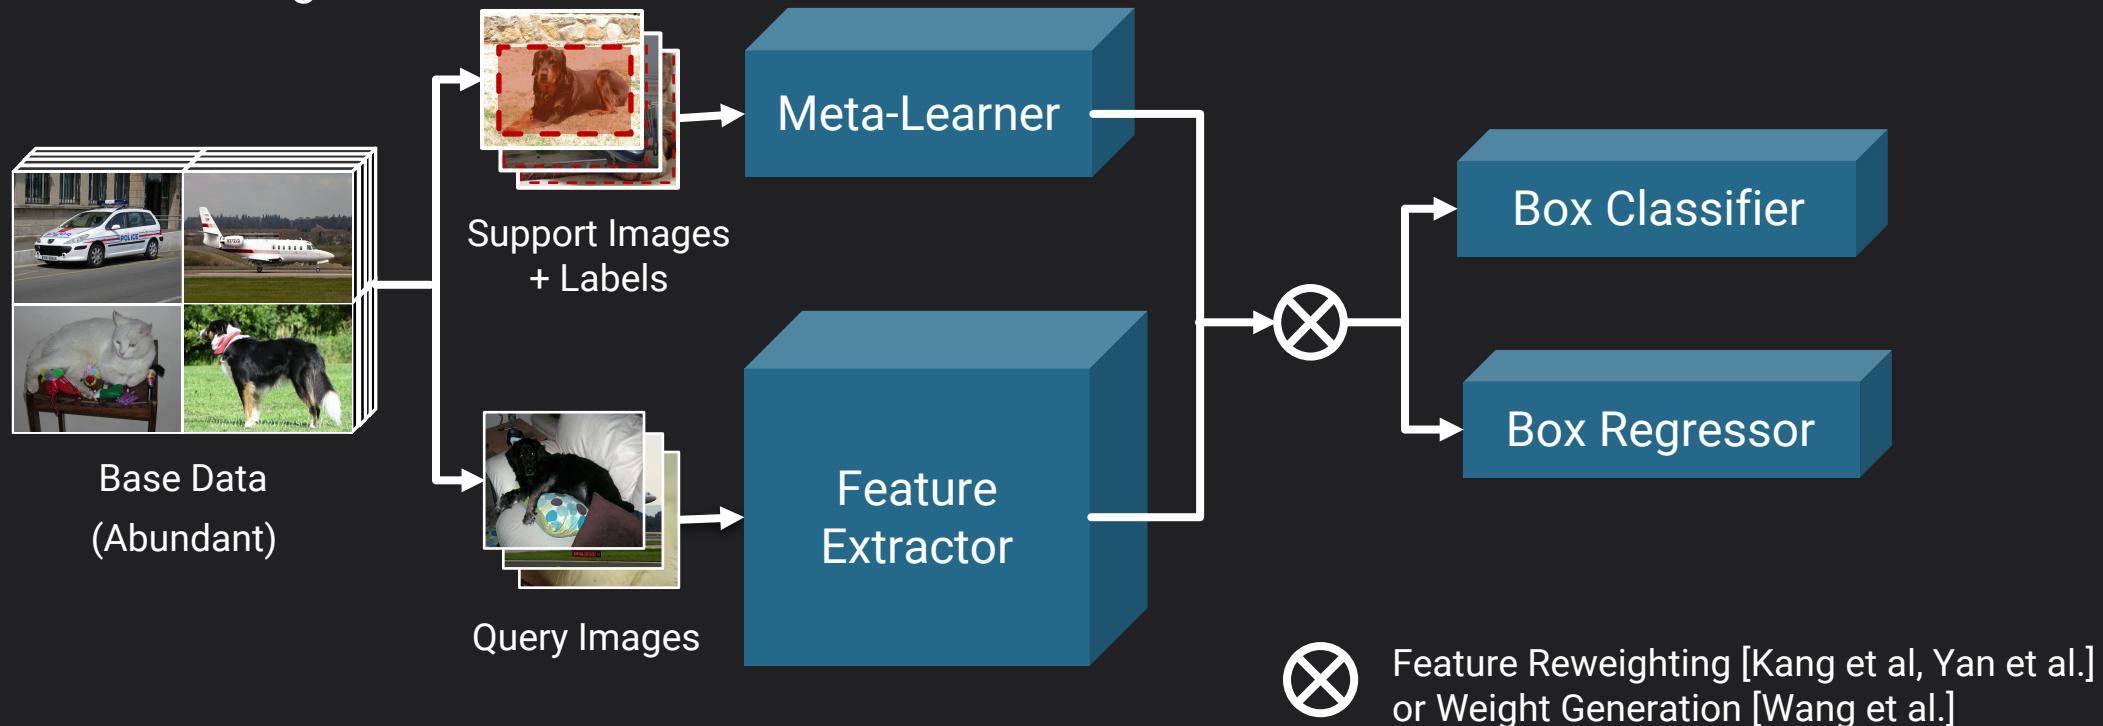
Humans can quickly generalize
from few examples



Hard to annotate many
examples for each concept

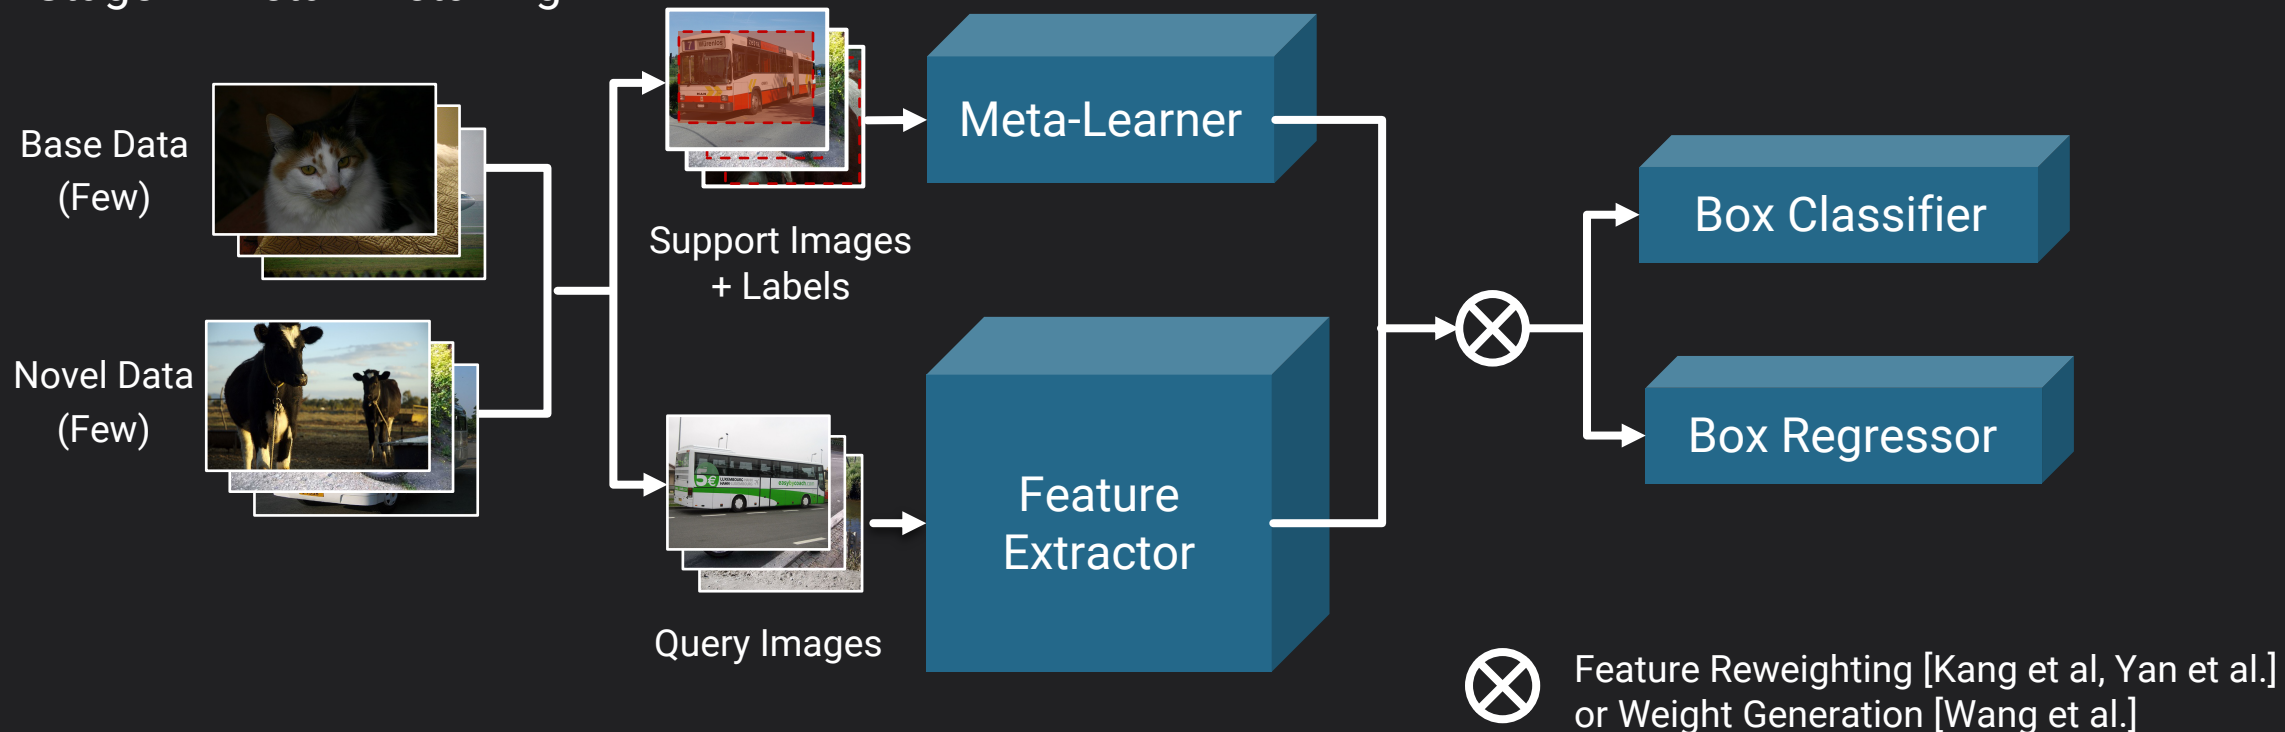
Trend: Meta Learning

Stage I: Meta Training



Trend: Meta Learning

Stage II: Meta Finetuning



Kang, B., Liu, Z., Wang, X., Yu, F., Feng, J., and Darrell, T. Few-shot object detection via feature reweighting. In *ICCV*, 2019.

Wang, Y.-X., Ramanan, D., and Hebert, M. Meta-learning to detect rare objects. In *ICCV*, 2019.


Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., and Lin, L. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *ICCV*, 2019.

Problems



Model Complexity

The meta learner adds extra components to the model



Memory Consumption

Training requires loading in images for each class at once

What about Simple Finetuning?



Finetuning has a proven record for model adaption



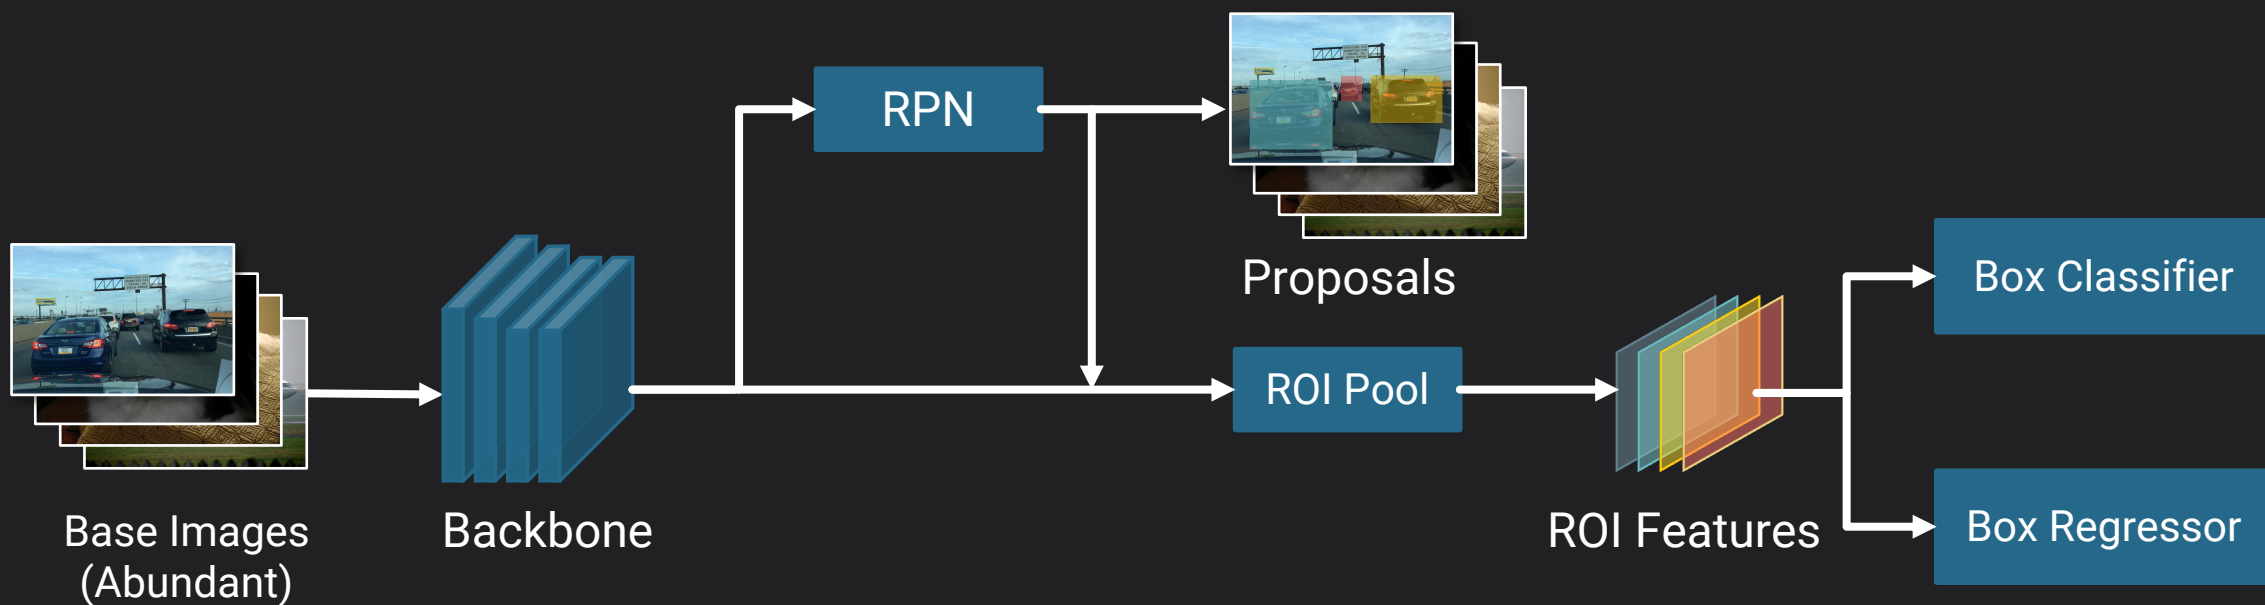
Finetuning may lead to overfitting given a few examples

Finetuning Comes to the Rescue

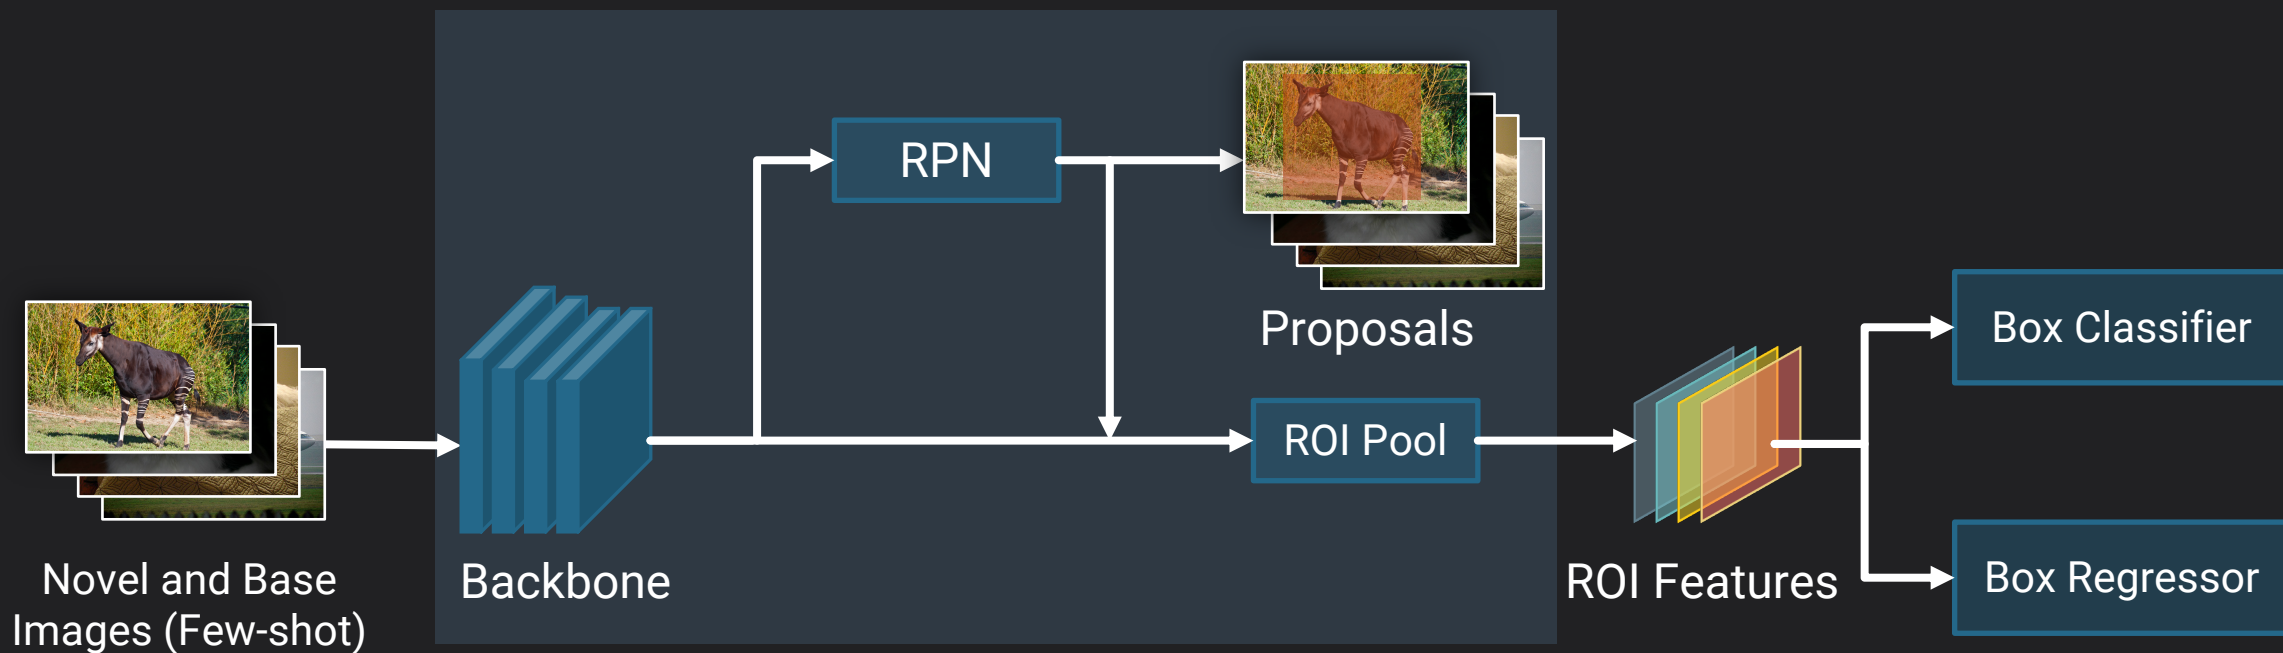


We find finetuning with controlled capacity can advance the best results of few-shot object detection significantly

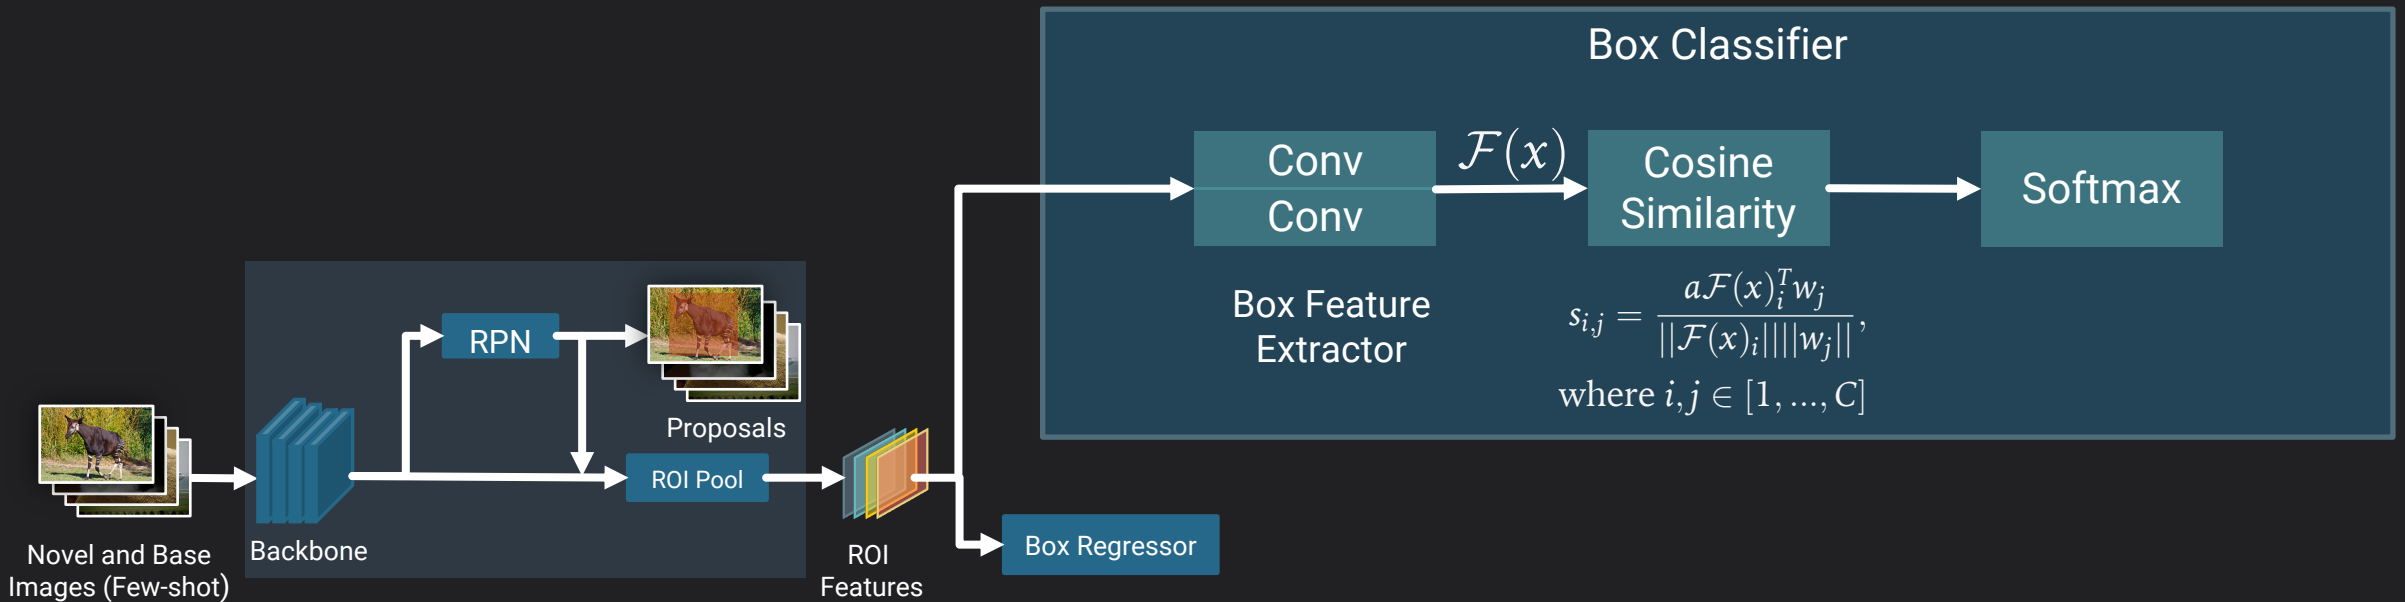
Our Approach: Base Training



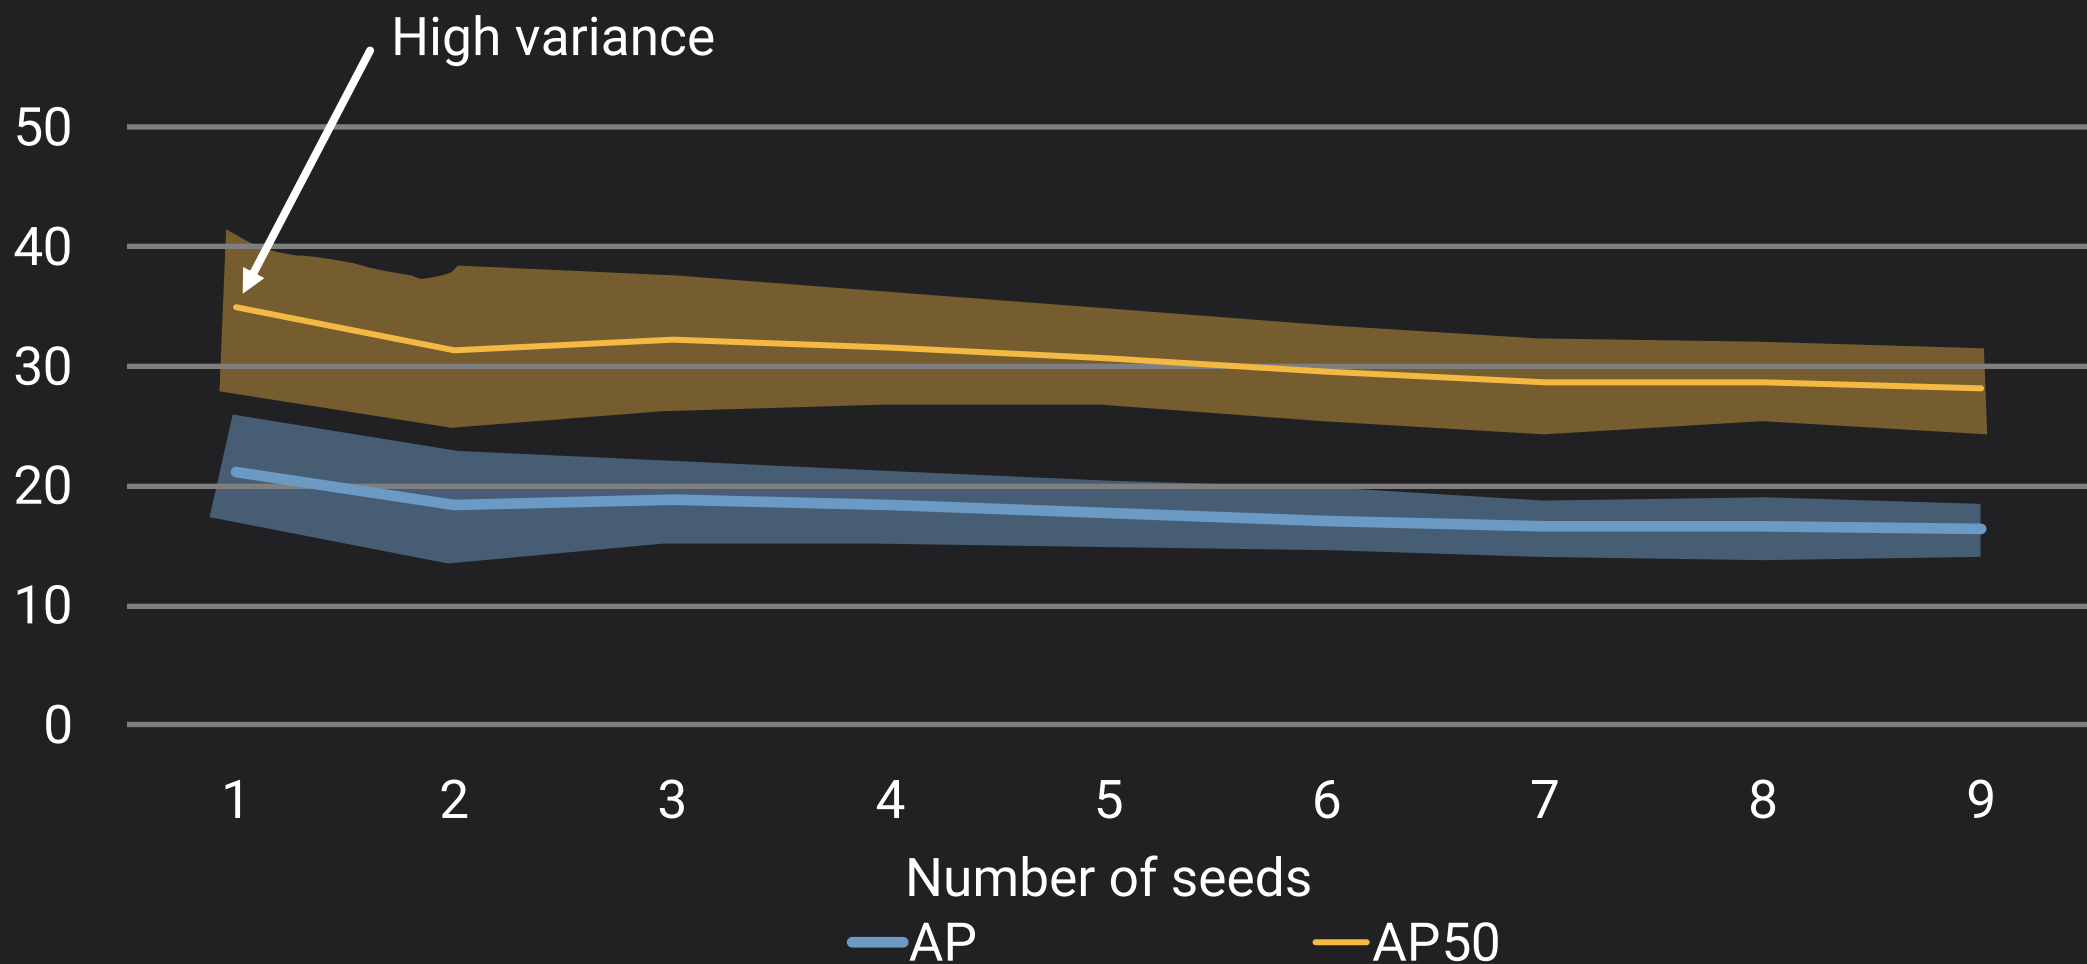
Our Approach: Few-Shot Finetuning



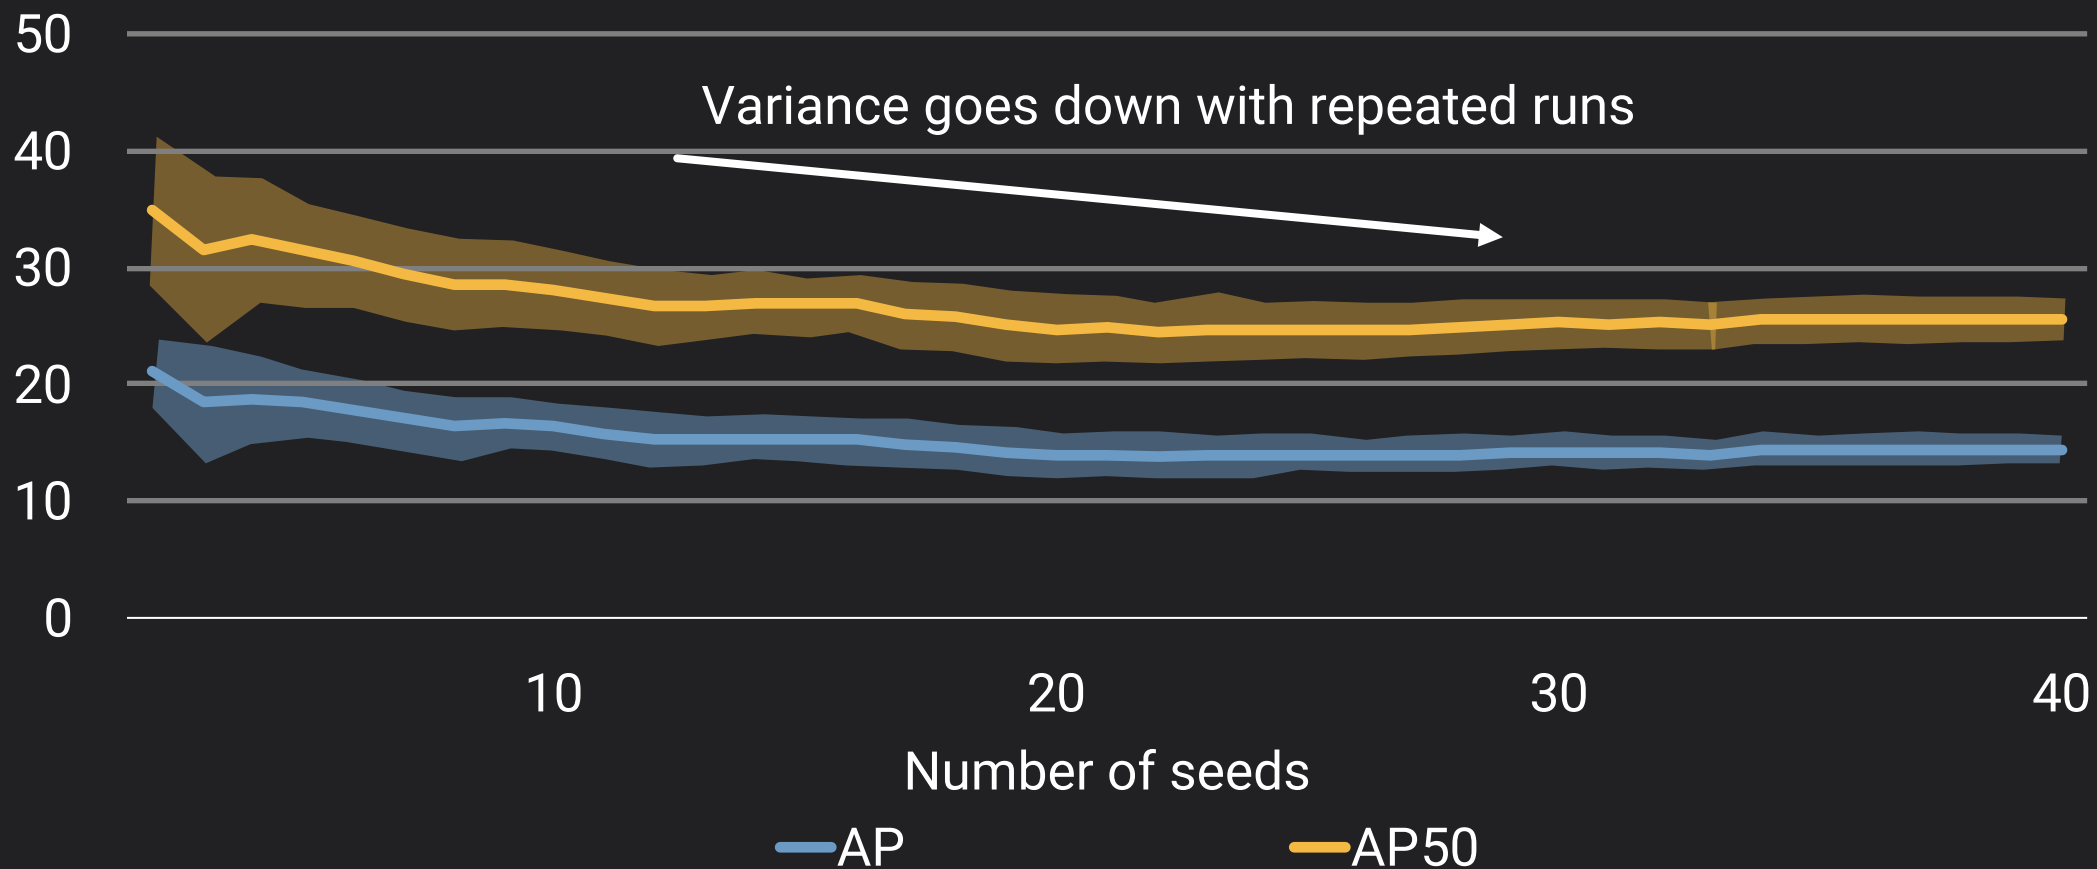
Cosine Similarity Based Box Classifier



Issue with Existing Benchmarks




More Stable Benchmarks




Dataset Statistics

	# Classes	# Base Classes	# Novel Classes
PASCAL VOC (Everingham et al.)	20	15	5
COCO (Lin et al.)	80	60	20
LVIS (Gupta et al.)	1230	776	454


Evaluation on PASCAL VOC




Evaluation on PASCAL VOC



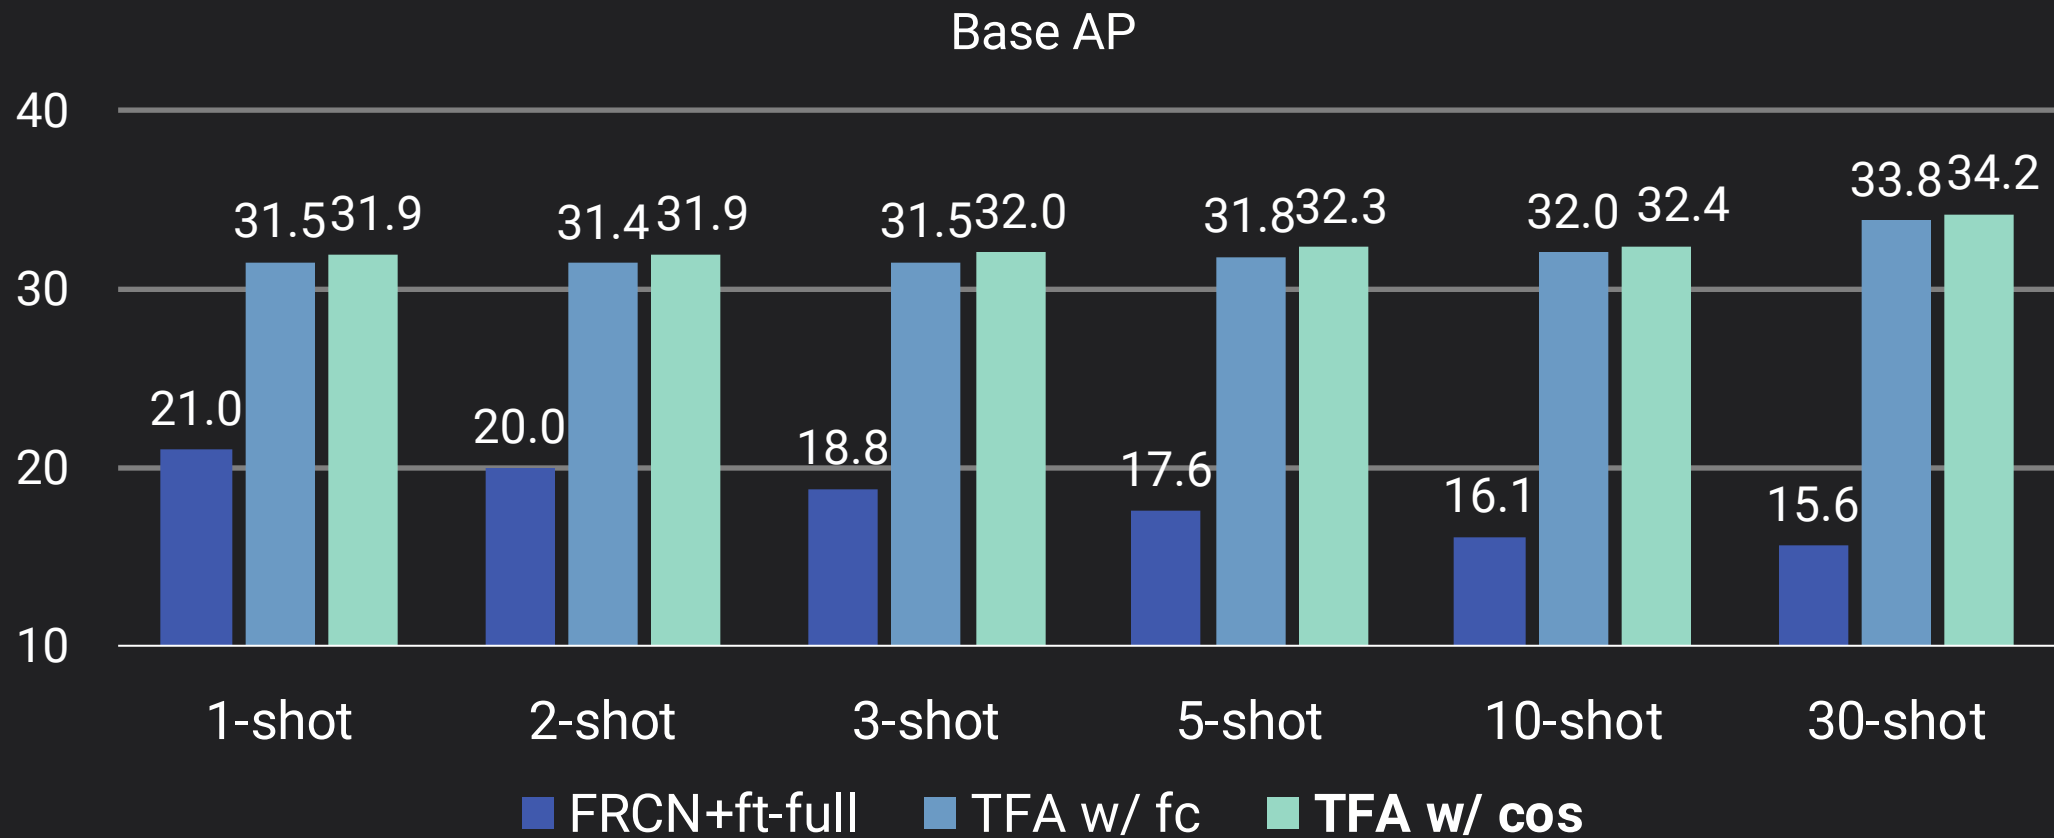
Evaluation on COCO




Evaluation on COCO




Evaluation on COCO



Evaluation on LVIS



Success Cases on VOC



Failure Cases on VOC



Success Cases on COCO



Failure Cases on COCO



Summary

- Instead of meta-learning, we find fine-tuning only the last layers of the detector to be crucial for few-shot object detection
- Our approach achieves state-of-the-art performance on existing benchmarks
- We build new benchmarks to address the unreliability of existing benchmarks

Paper and Code

