

Problem:

Given a few support images with ground-truth segmentations and a query image from the same class, segment the query.

Motivation:

Non-discriminative features:
For various distinct classes, CNNs may have high activations

Contribution 1:
Estimate feature relevance

Overfitting:
Learning from few support images cannot capture all variations of the target class

Contribution 2:
Leverage gradient boosting

Our Results on PASCAL-5i



TRAINING

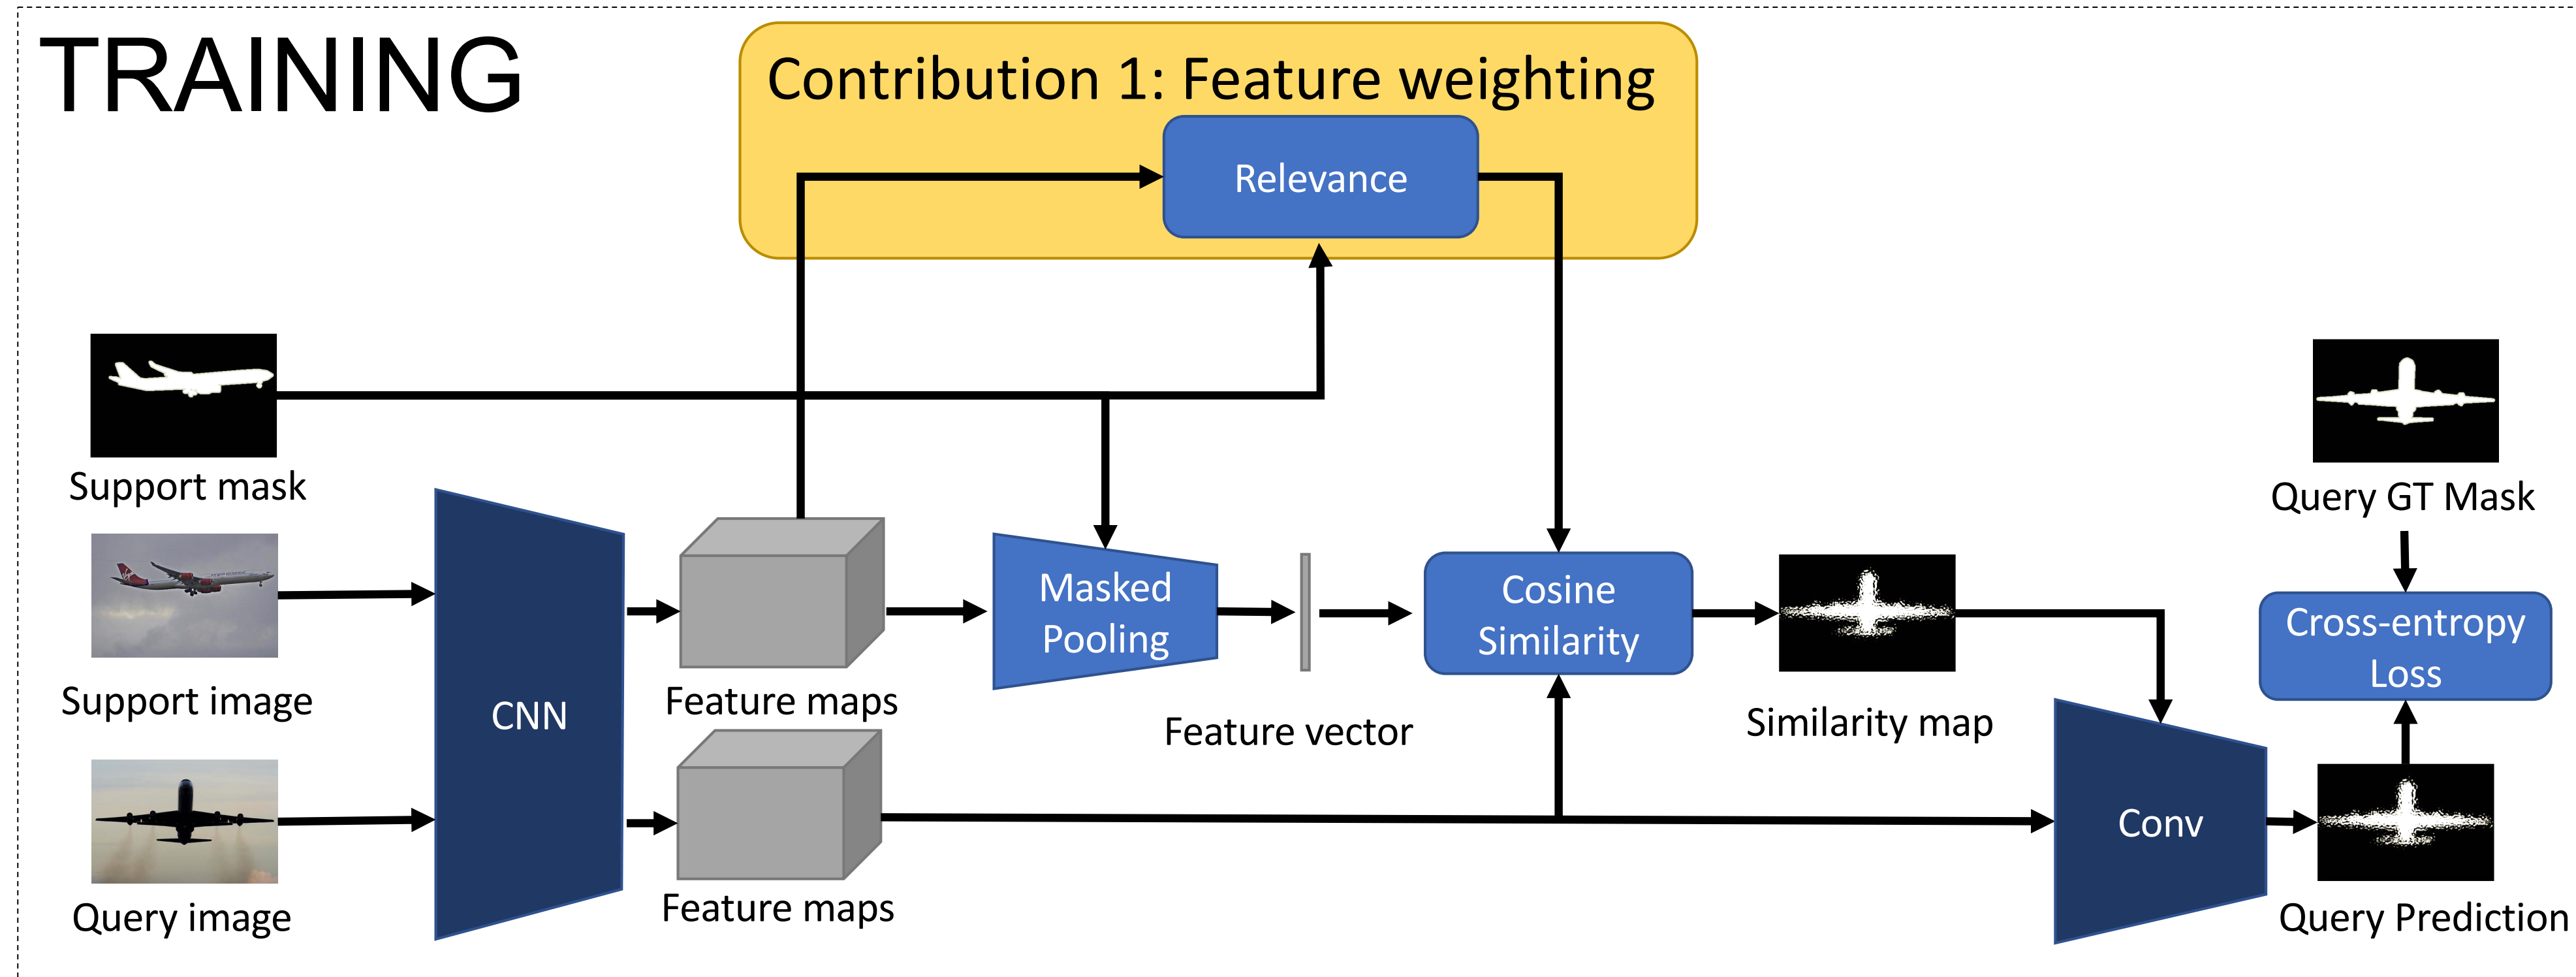


Table 1. Mean IoU of five-shot segmentation on PASCAL-5i.

Backbone	Model	Mean IoU (over 4 parts)	
		One-shot	Five-shot
VGG 16	OSLSM (Shaban et al., BMVC 17)	40.80	43.95
	Co-FCN (Rakelly et al., ICLR W 18)	41.10	41.38
	PL+SEG+PT (Dong et al., BMVC 18)	42.70	43.70
	SG-One (Zhang et al., Arxiv 18)	46.30	47.10
	Baseline	47.55	-
	Baseline + Contribution 1	48.68	-
	Baseline + Contribution 2	51.22	-
	Our approach	51.90	55.08
Resnet 101	Our approach + full supervision	58.82	-
	Our approach	56.19	59.92
Resnet 101	Our approach + full supervision	68.07	-
	Our approach	68.07	68.07

TESTING

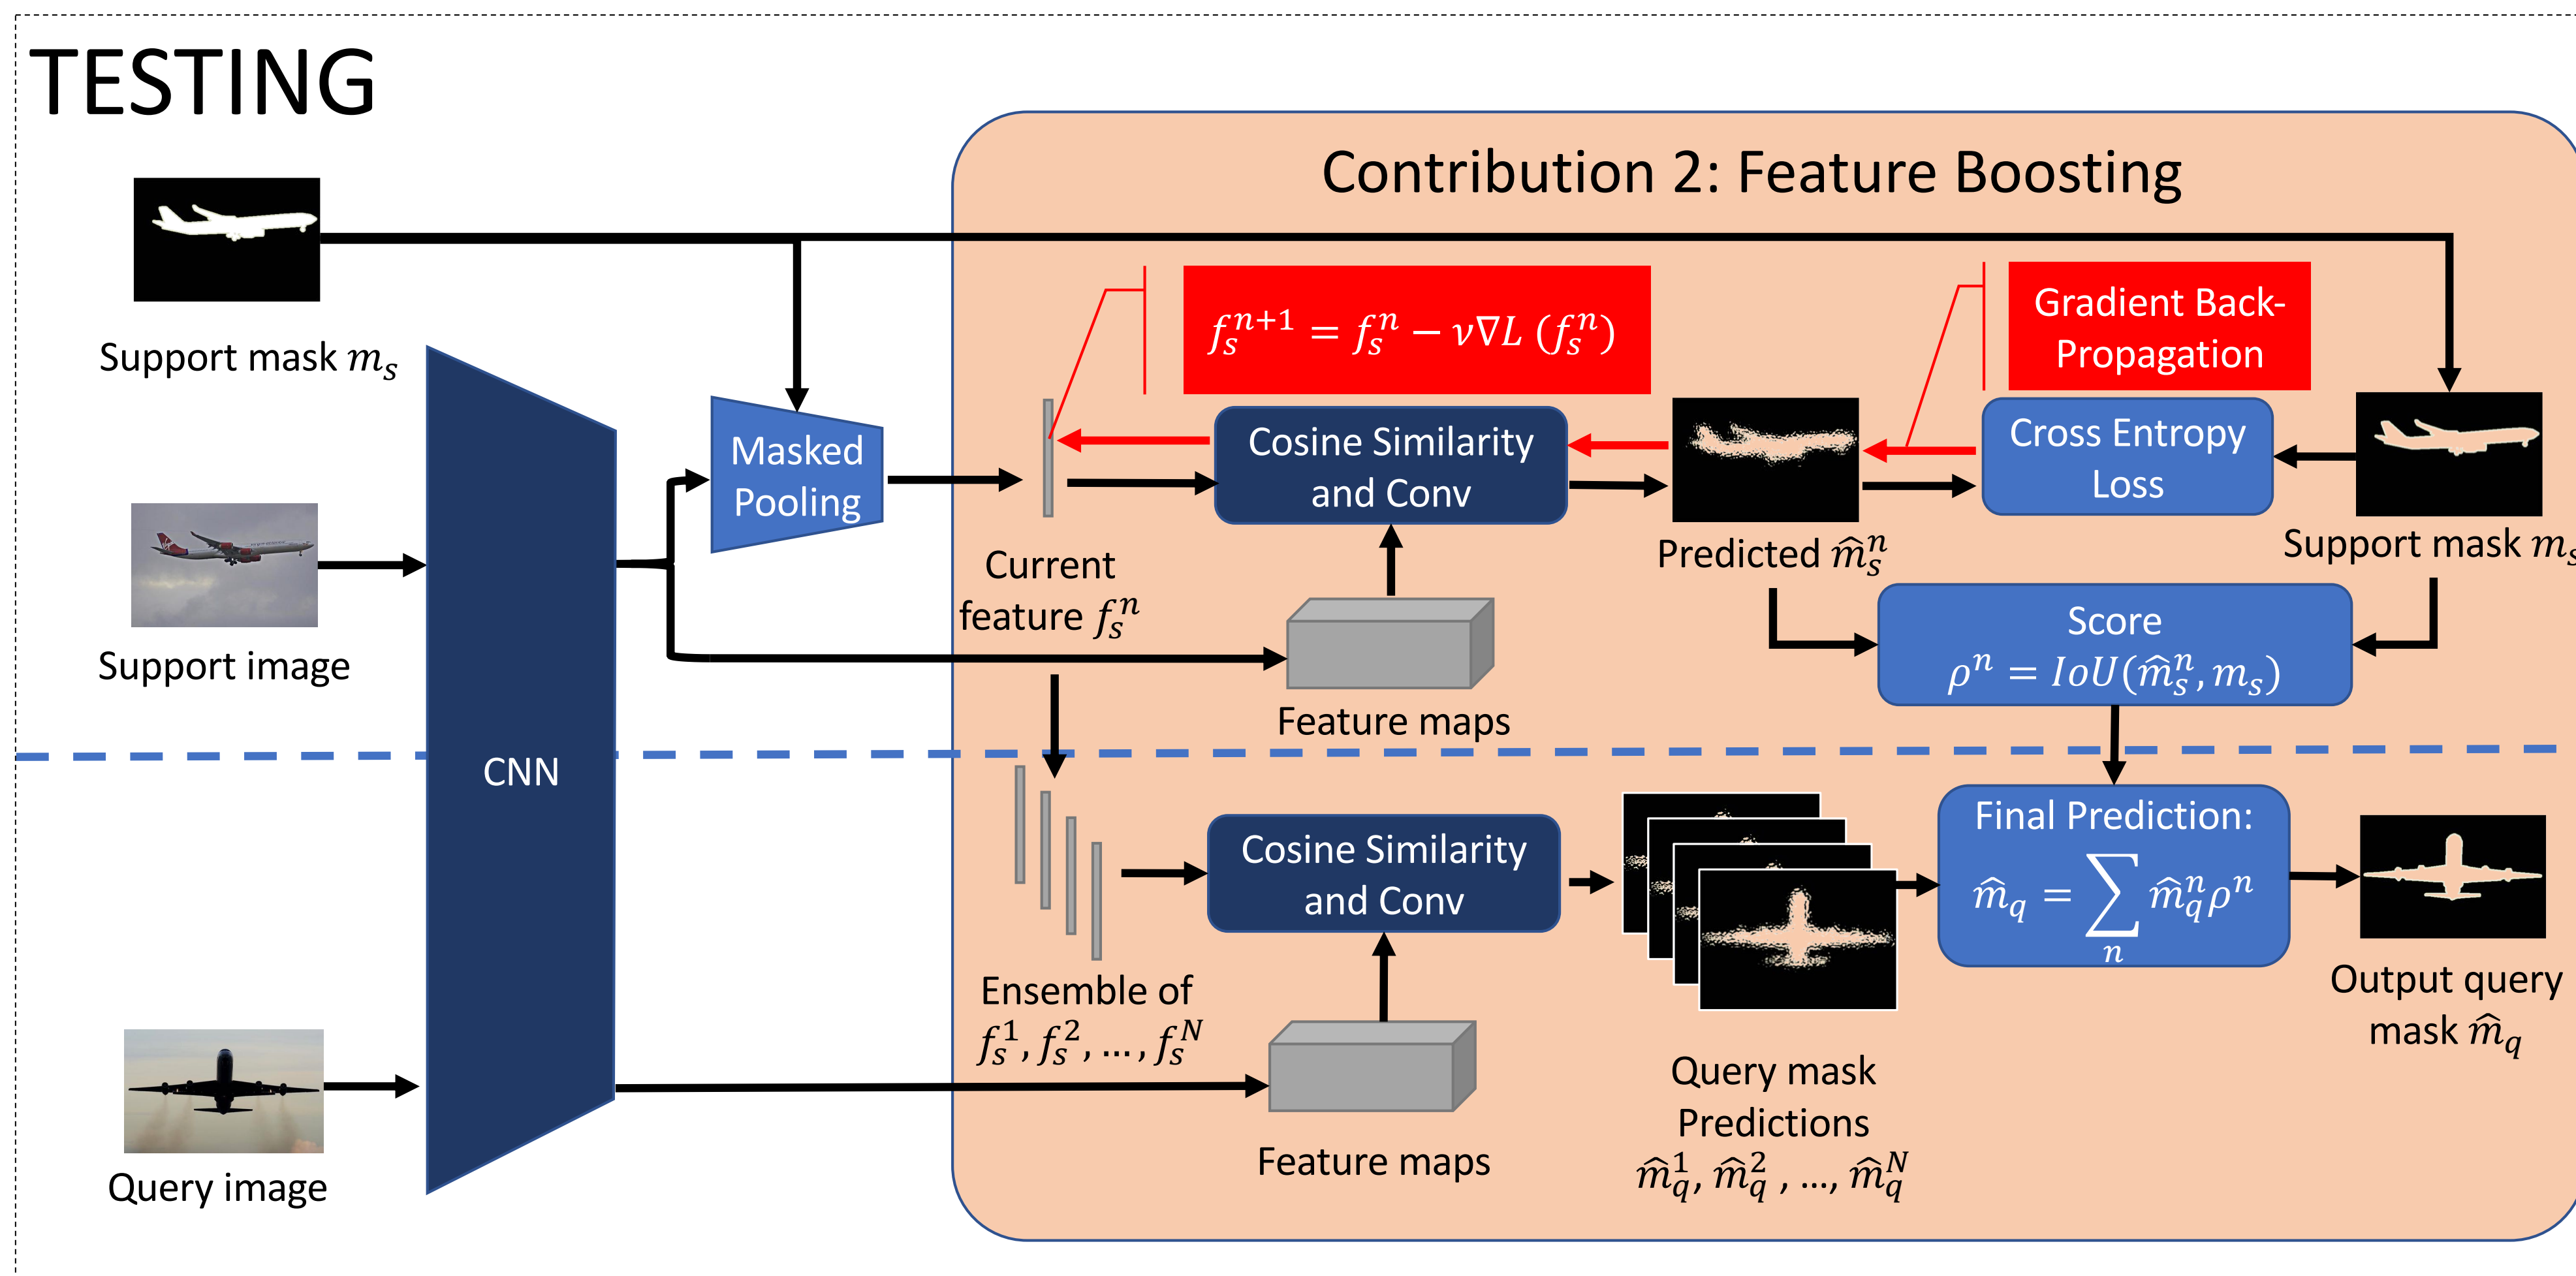


Table 2. Mean IoU of few-shot segmentation on COCO-20i.

Backbone	Model	Mean IoU (over 4 parts)	
		One-shot	Five-shot
VGG 16	Baseline	13.26	-
	Baseline + Contribution 1	17.11	-
	Baseline + Contribution 2	17.73	-
	Our approach	20.02	22.63
Resnet 101	Our approach	21.19	23.65

Table 3. Training and testing time on PASCAL-5i

Backbone	Phase	Baseline	B + C1	B + C2	B + C1 + C2
VGG 16	Train	171	172	171	172
	Test	148	150	211	211
Resnet 101	Train	386	388	386	388
	Test	268	280	360	360

Acknowledgment: DARPA XAI Award N66001-17-2-4029