OpenMatch: Open-set Consistency Regularization for Semi-supervised Learning with Outliers [Link]

Kuniaki Saito[1], Dongyun Kim[1], Kate Saenko[1,2]

1: Boston University, 2: MIT-IBM Watson AI Lab



Semi-supervised Learning with outliers

- Towards more realistic semi-supervised learning

Standard Semi-supervised Learning Set-up

Bear Bird

Labeled Data

Same categories

Unlabeled Data

Semi-supervised Learning with outliers (Open-set SSL)

Bear Bird Outliers (Novel Class)

Goal and Requirements

- Correctly classify inliers
- Detect outliers

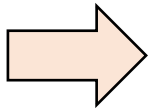

Separate inliers and outliers, computing SSL loss only for unlabeled inliers

Related Work

- Semi-supervised learning
  - FixMatch [Sohn et.al., NeurIPS2020]
  - Pseudo-labeling

  ] Risk of assigning pseudo-labels to outliers.

- Semi-supervised learning with outliers
  - MTC [Yu et.al., ECCV2020]
    - Thresholding is not robust to various number of outliers.
  - D3LS [Guo et.al., ICML2020]
    - No objective to ensure separation between inliers and outliers.

Goal and Requirements

- Correctly classify inliers
- Detect outliers

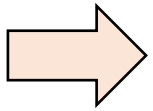

Separate inliers and outliers, computing SSL loss only for unlabeled inliers

1. Learning such representations

2. A threshold to pick inliers

Contribution

- New Framework, OpenMatch, for Open-set SSL
  - Outlier Detector based on one-vs-all classification network.
  - Soft consistency regularization with outlier detector improves outlier detection.
  - FixMatch for pseudo-inliers.
- SOTA performance in Open-set SSL
  - Accuracy to classify inliers
  - Separation between inliers and outliers
- Effectiveness in novelty detection
  - Detect outliers unseen in unlabeled data

Outlier Detector with one-vs-all classification network

- One-vs-All Network [Saito et.al, Arxiv 2021]
  - Train K one-vs-all classifier
  - Treat one class as positive, others as negative
  - Outlier Detector providing a threshold

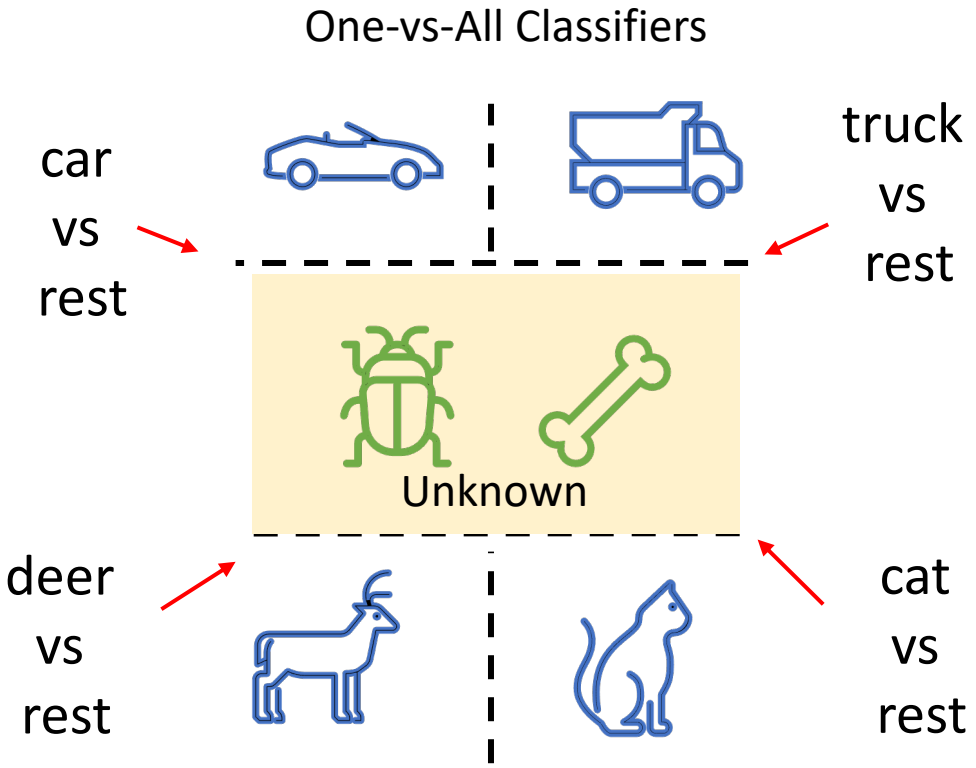

Open-set Soft Consistency Regularization

- Enhance smoothness over data augmentation.
- For better outlier detection.

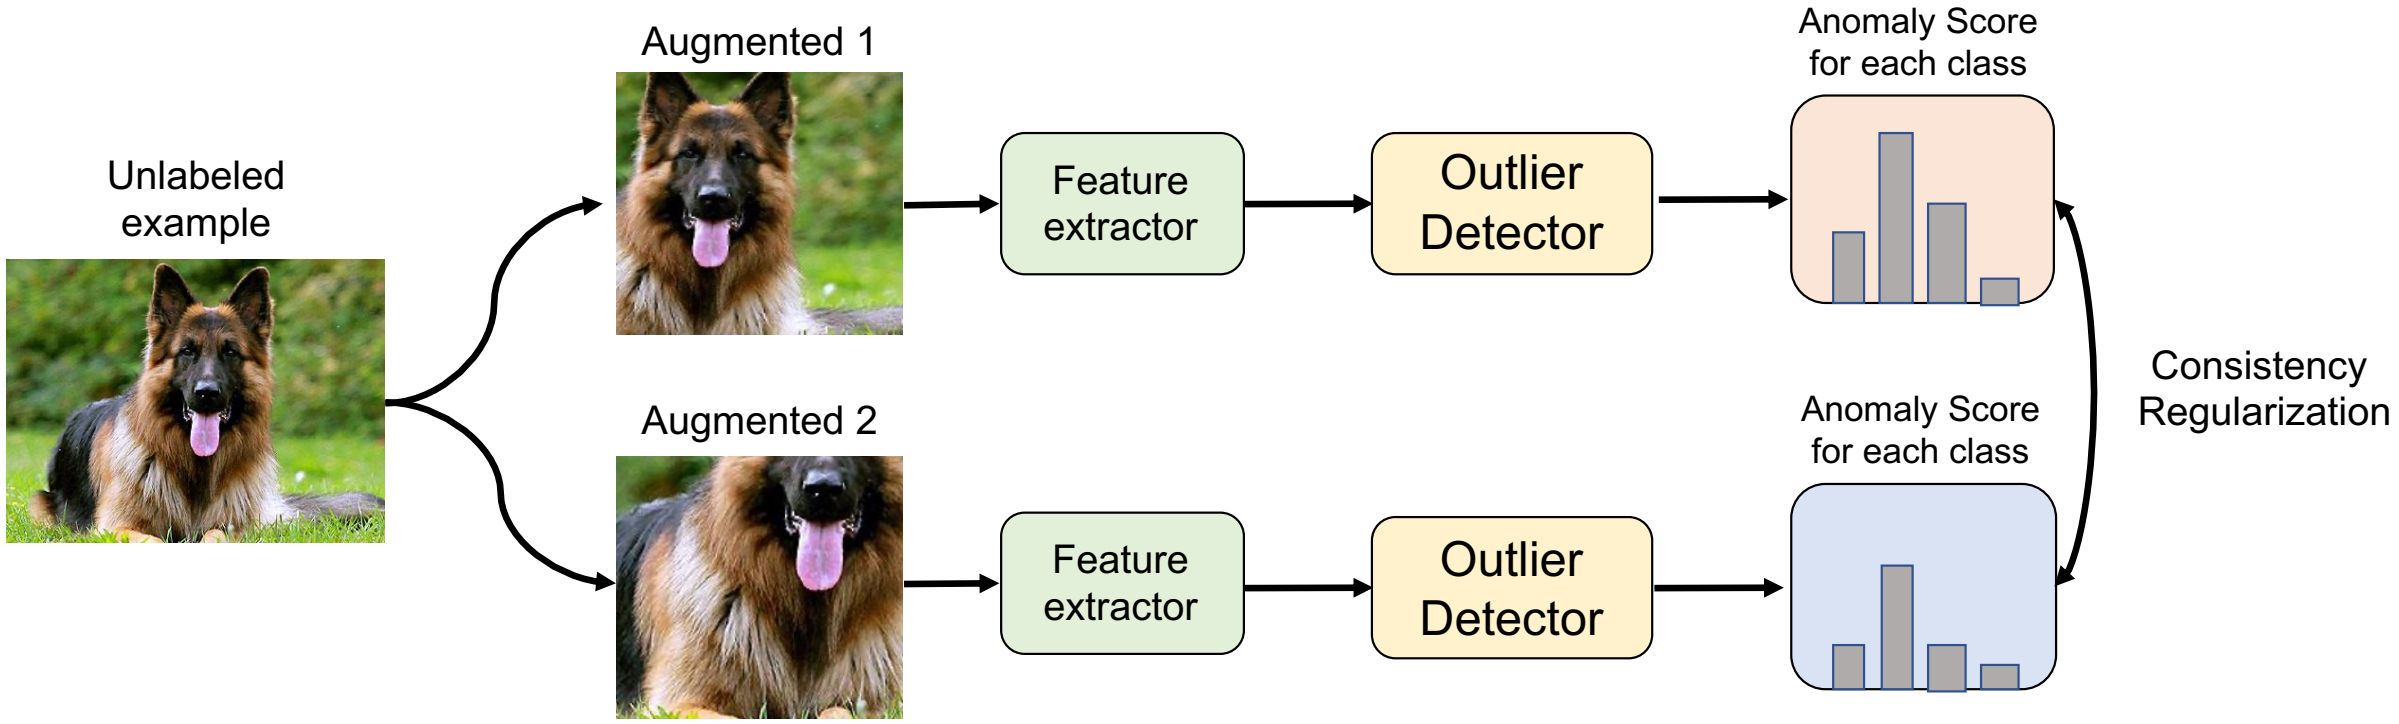

Pseudo label vs Soft Consistency

Labeled Inliers

Unlabeled Inliers

Unlabeled Outliers

P = Probability

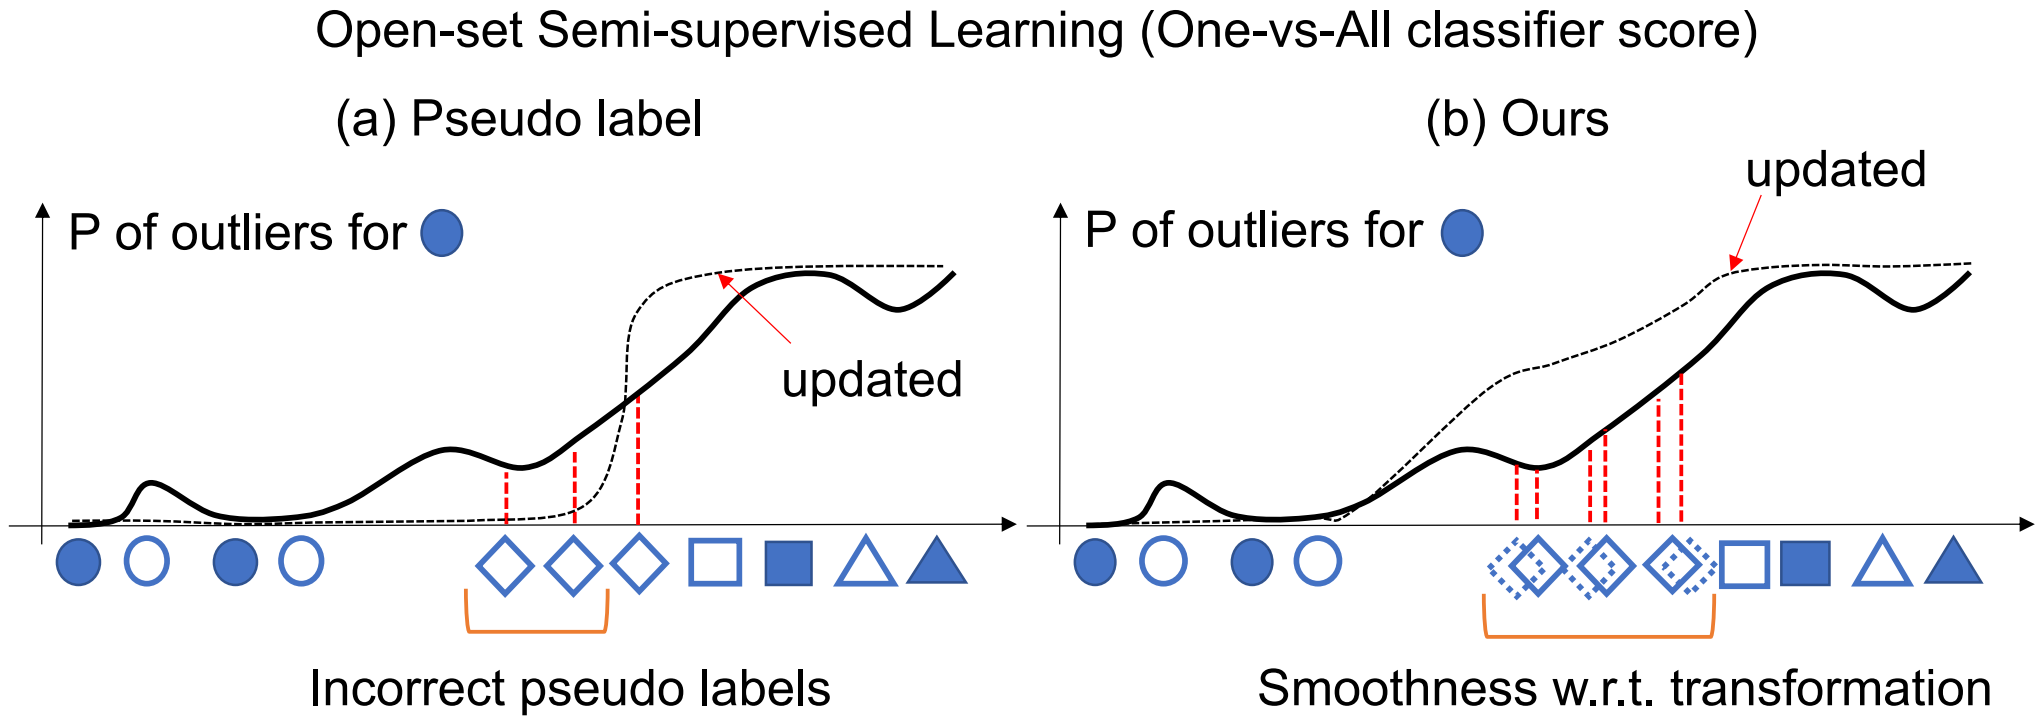

Pseudo-inlier selection

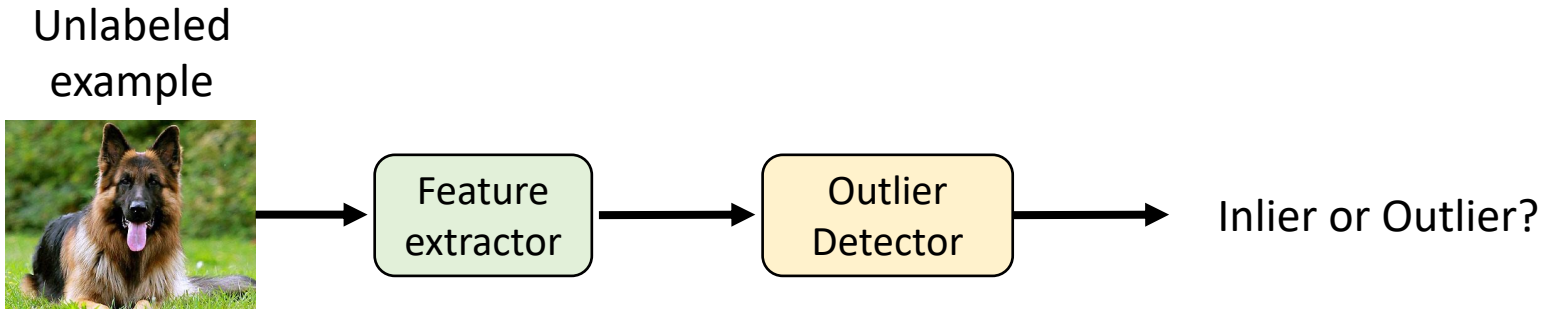

Overall Training

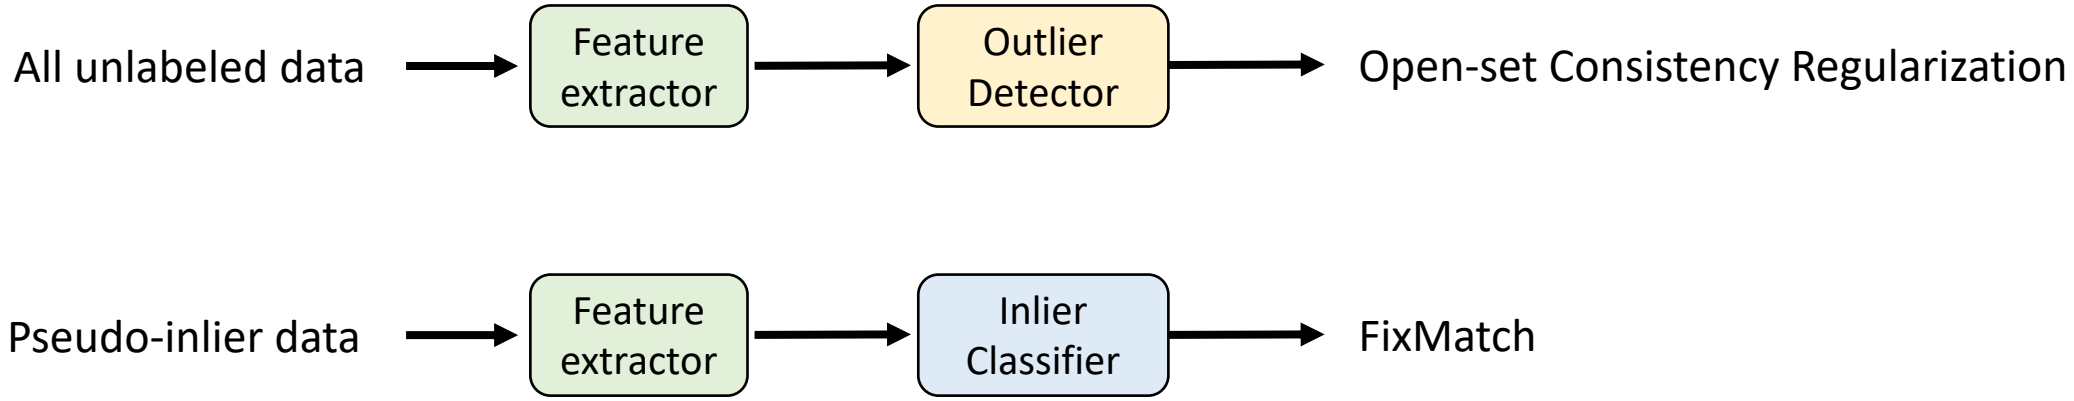

Accuracy for inliers CIFAR10: 6 inlier classes, 4 outlier classes

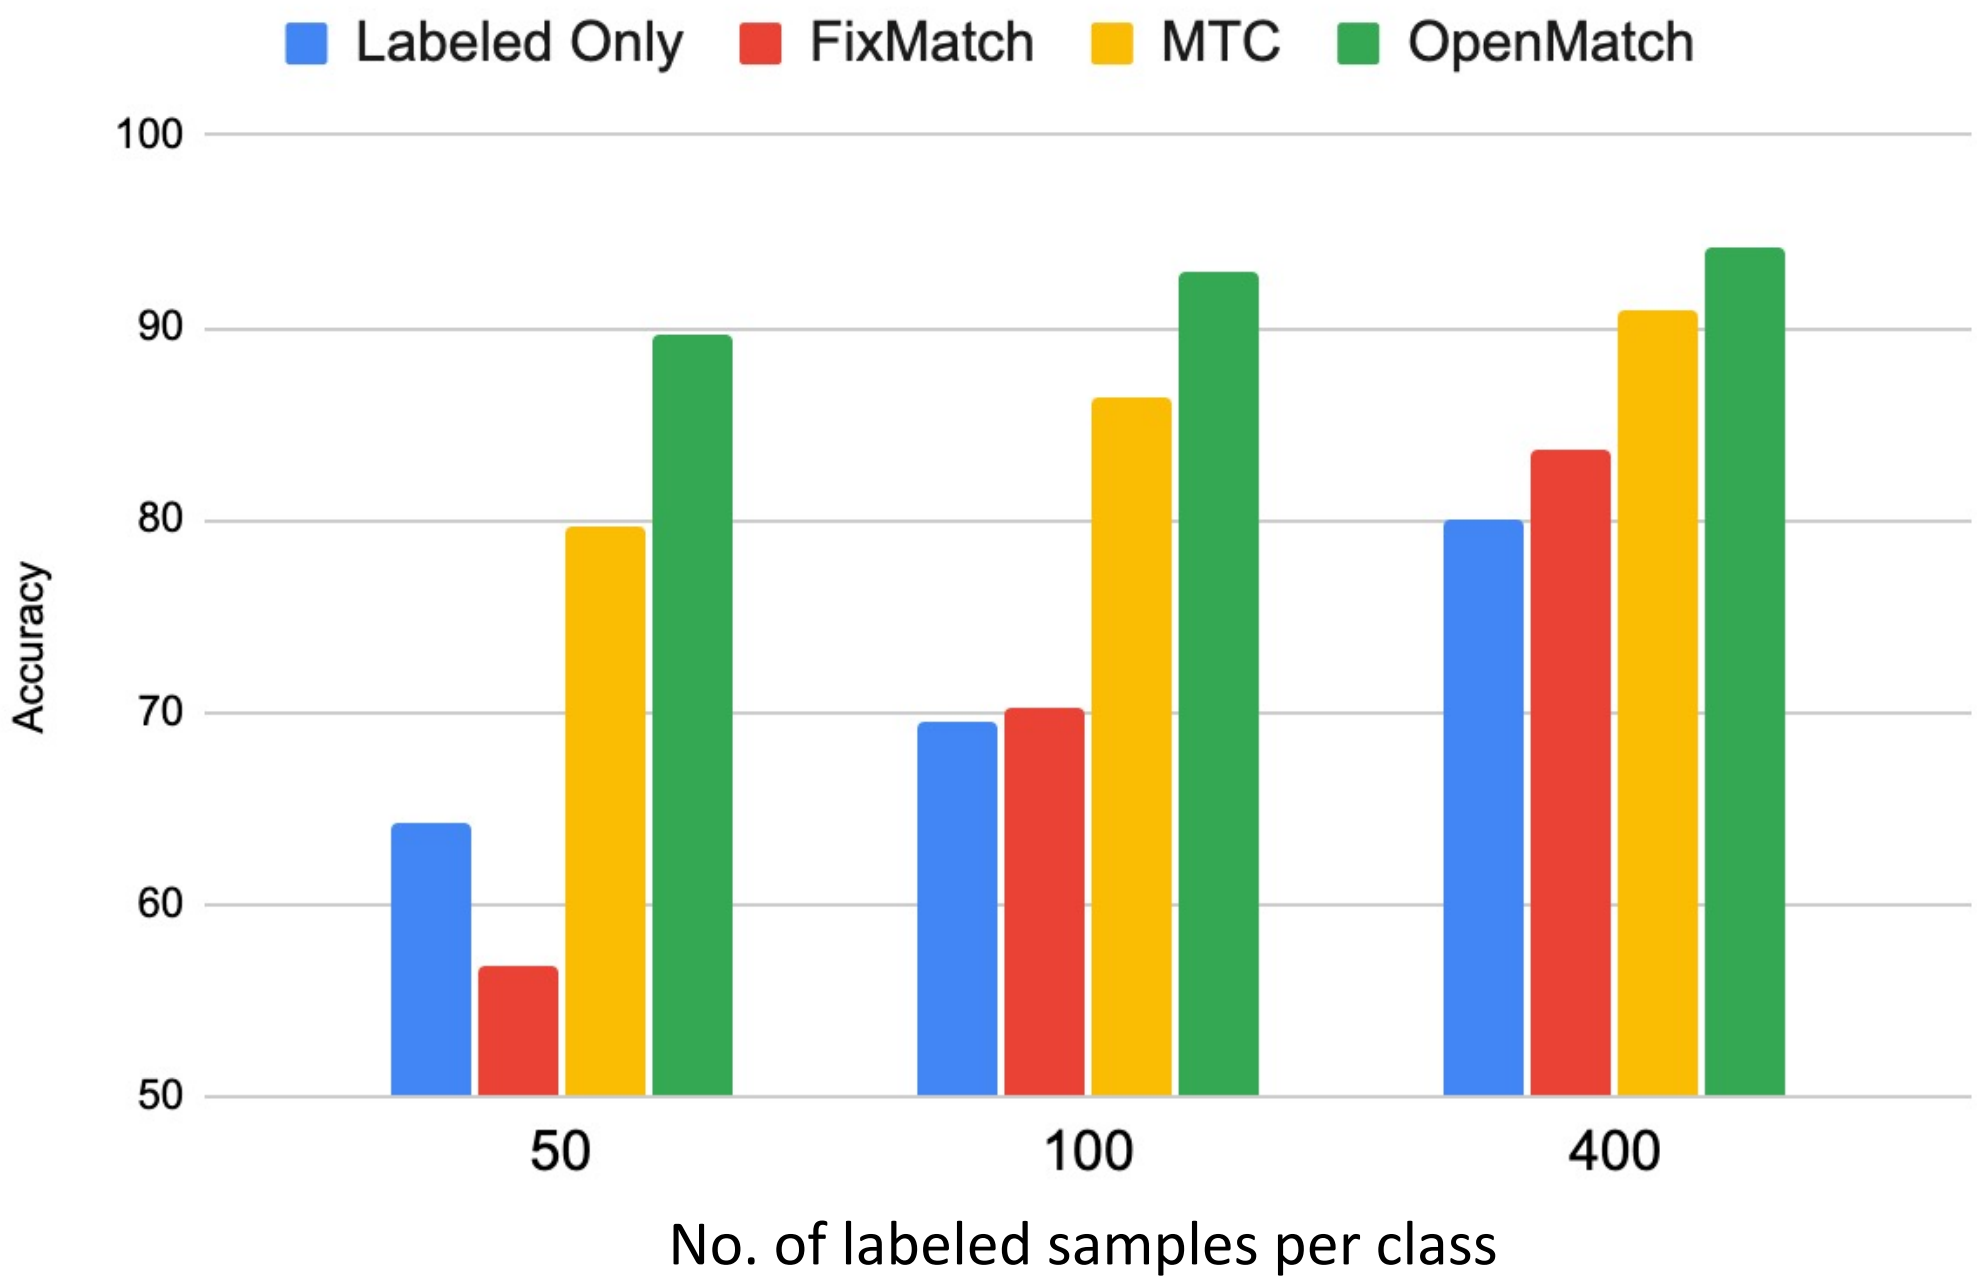

Accuracy for inliers CIFAR10: 6 inlier classes, 4 outlier classes

Effect of outlier detector

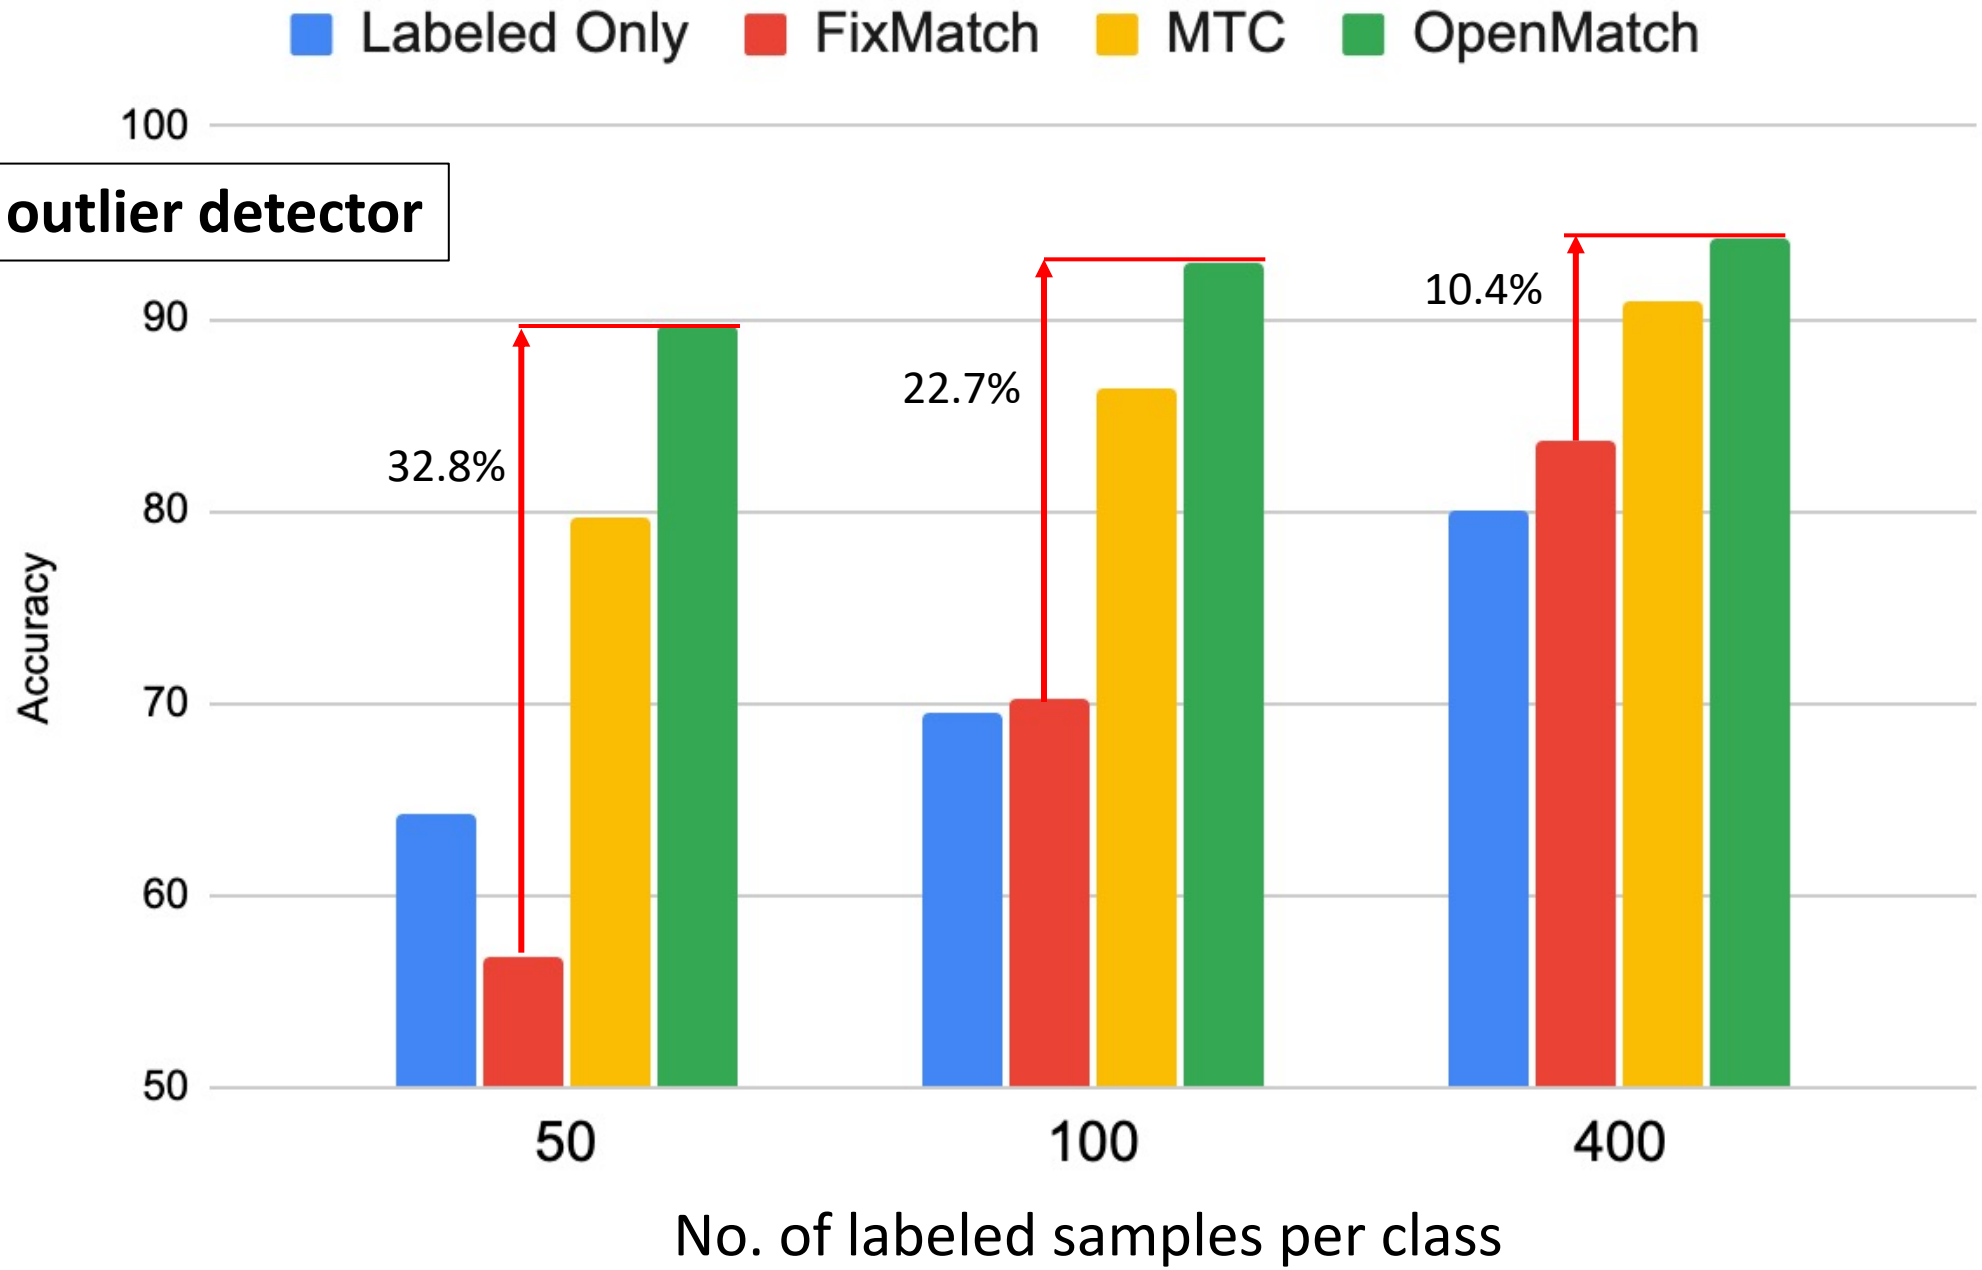

Accuracy for inliers CIFAR10: 6 inlier classes, 4 outlier classes

Comparison with SOTA

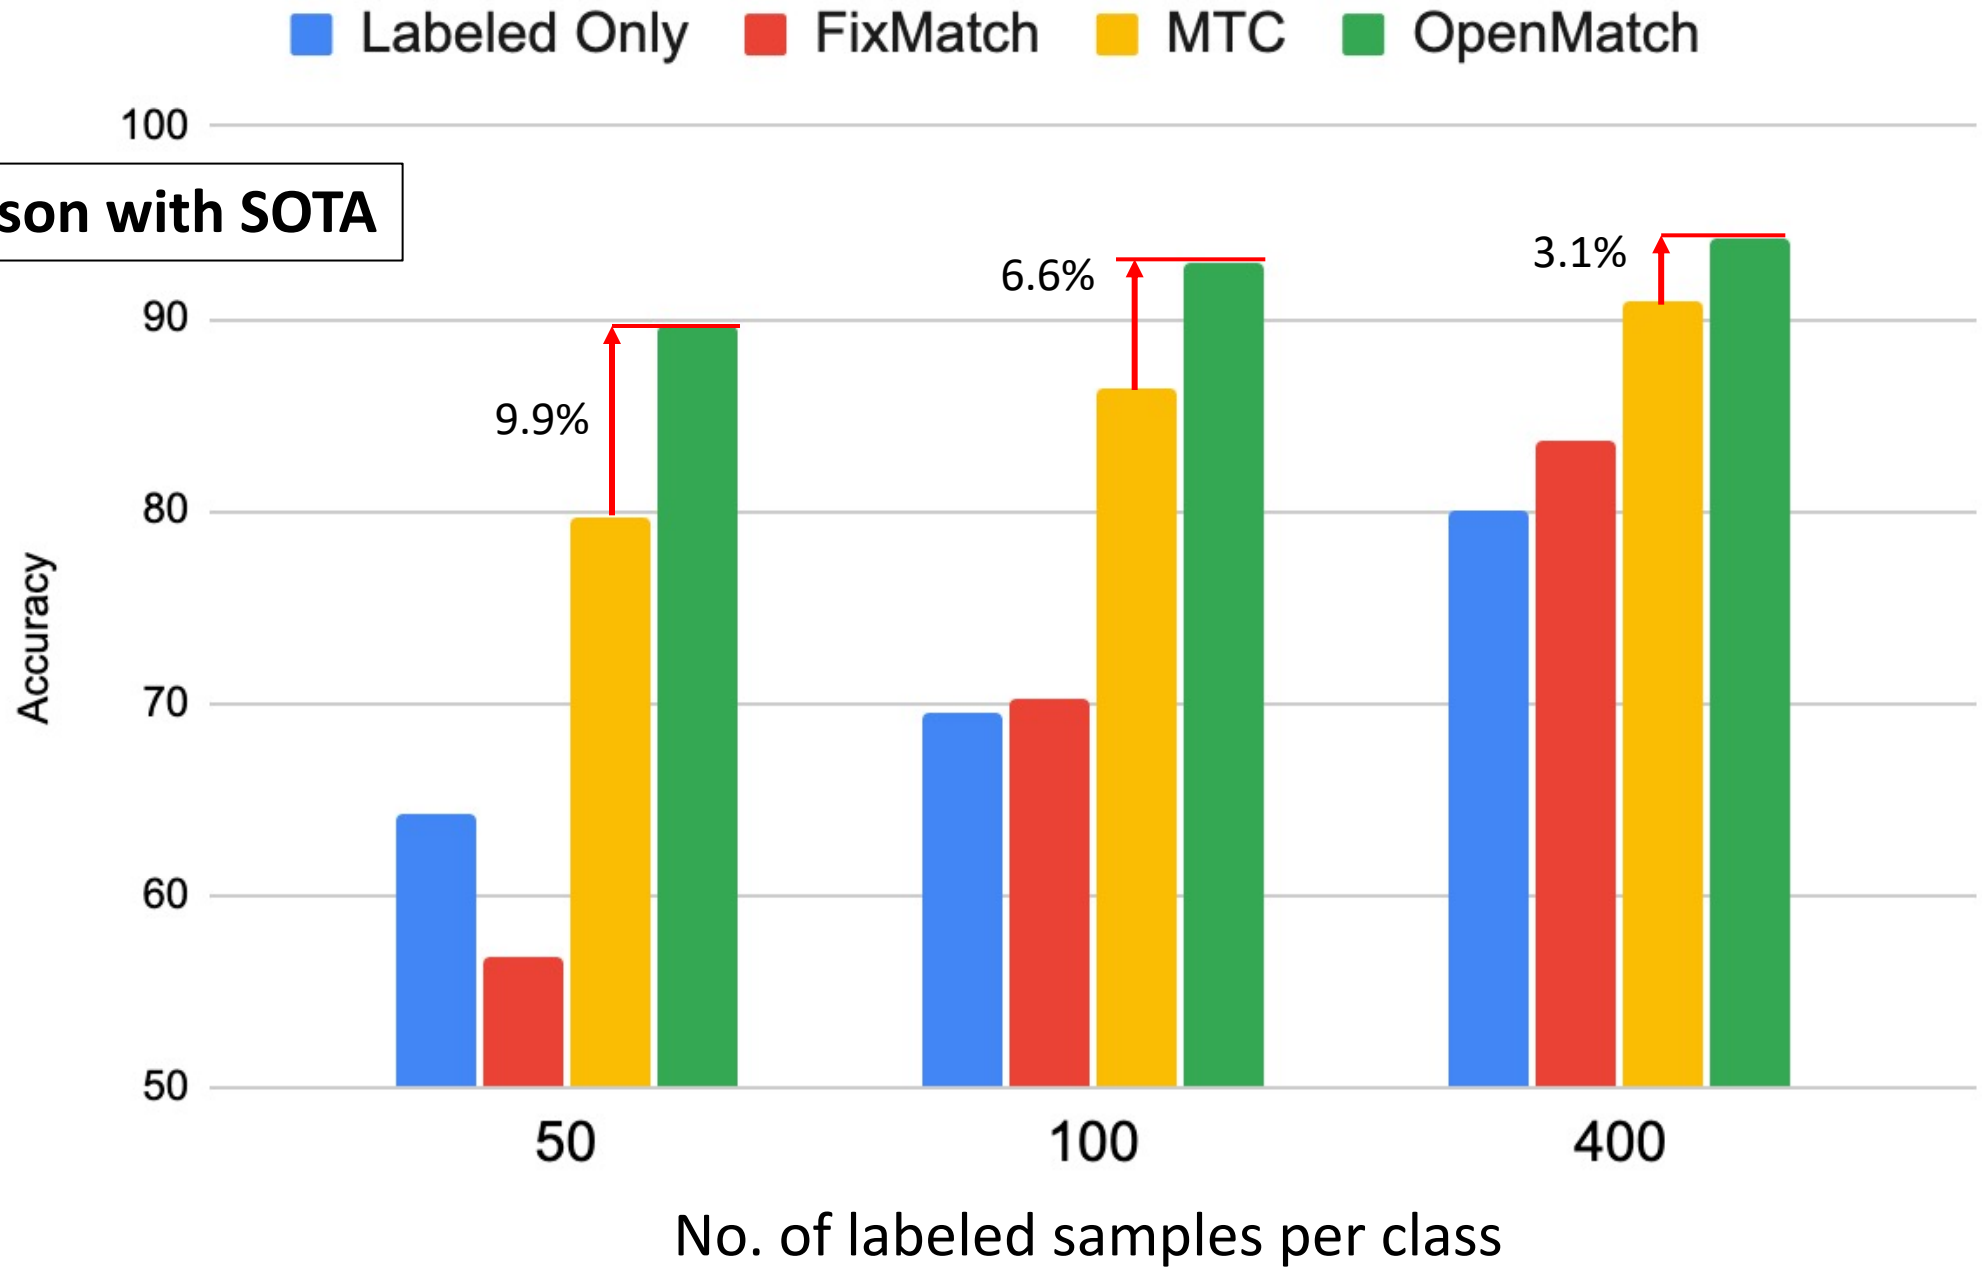

AUROC: Outlier Detection CIFAR10: 6 inlier classes, 4 outlier classes

Effect of outlier detector

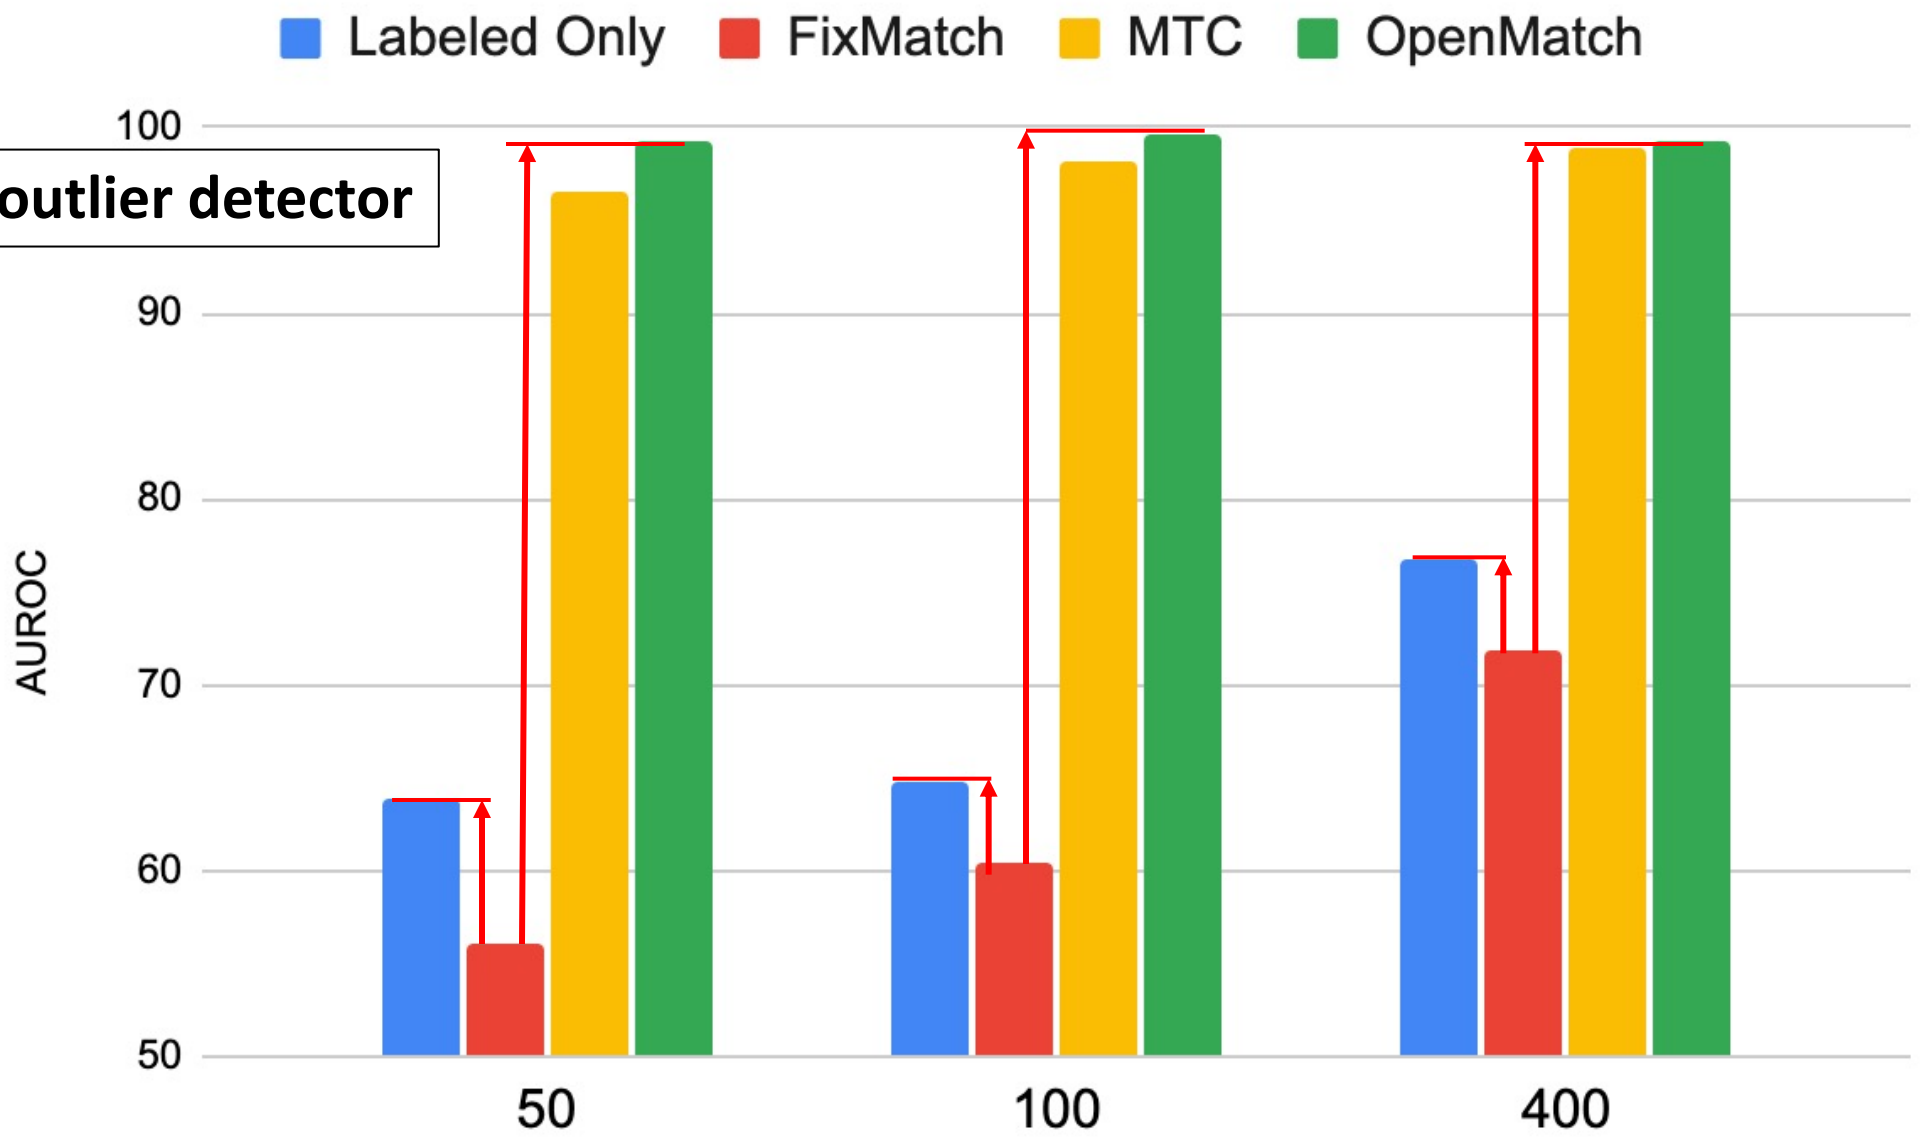

Accuracy for inliers

Dataset	CIFAR10			CIFAR100		CIFAR100		ImageNet-30
No. of Known / Unknown	6 / 4			55 / 45		80 / 20		20 / 10
No. of labeled samples	50	100	400	50	100	50	100	10 %
Labeled Only	35.7±1.1	30.5±0.7	20.0±0.3	37.0±0.8	27.3±0.5	43.6±0.5	34.7±0.4	20.9±1.0
FixMatch [35]	43.2±1.2	29.8±0.6	16.3±0.5	35.4±0.7	27.3±0.8	41.2±0.7	34.1±0.4	12.9±0.4
MTC [44]	20.3±0.9	13.7±0.9	9.0±0.5	33.5±1.2	27.9±0.5	40.1±0.8	33.6±0.3	13.6±0.7
OpenMatch	**10.4**±0.9	**7.1**±0.5	**5.9**±0.5	**27.7**±0.4	**24.1**±0.6	**33.4**±0.2	**29.5**±0.3	**10.4**±1.0

AUROC to detect outliers

Dataset	CIFAR10			CIFAR100		CIFAR100		ImageNet-30
No. of Known / Unknown	6 / 4			55 / 45		80 / 20		20 / 10
No. of labeled samples	50	100	400	50	100	50	100	10 %
Labeled Only	63.9±0.5	64.7±0.5	76.8±0.4	76.6±0.9	79.9±0.9	70.3±0.5	73.9±0.9	80.3±1.0
FixMatch [35]	56.1±0.6	60.4±0.4	71.8±0.4	72.0±1.3	75.8±1.2	64.3±1.0	66.1±0.5	88.6±0.5
MTC [44]	96.6±0.6	98.2±0.3	98.9±0.1	81.2±3.4	80.7±4.6	79.4±2.5	73.2±3.5	93.8±0.8
OpenMatch	**99.3**±0.3	**99.7**±0.2	**99.3**±0.2	**87.0**±1.1	**86.5**±2.1	**86.2**±0.6	**86.8**±1.4	**96.4**±0.7

Accuracy for inliers

Dataset	CIFAR10			CIFAR100		CIFAR100		ImageNet-30
No. of Known / Unknown	6 / 4			55 / 45		80 / 20		20 / 10
No. of labeled samples	50	100	400	50	100	50	100	10 %
Labeled Only	35.7±1.1	30.5±0.7	20.0±0.3	37.0±0.8	27.3±0.5	43.6±0.5	34.7±0.4	20.9±1.0
FixMatch [35]	43.2±1.2	29.8±0.6	16.3±0.5	35.4±0.7	27.3±0.8	41.2±0.7	34.1±0.4	12.9±0.4
MTC [44]	20.3±0.9	13.7±0.9	9.0±0.5	33.5±1.2	27.9±0.5	40.1±0.8	33.6±0.3	13.6±0.7
OpenMatch	**10.4**±0.9	**7.1**±0.5	**5.9**±0.5	**27.7**±0.4	**24.1**±0.6	**33.4**±0.2	**29.5**±0.3	**10.4**±1.0

AUROC to detect outliers

Dataset	CIFAR10			CIFAR100		CIFAR100		ImageNet-30
No. of Known / Unknown	6 / 4			55 / 45		80 / 20		20 / 10
No. of labeled samples	50	100	400	50	100	50	100	10 %
Labeled Only	63.9±0.5	64.7±0.5	76.8±0.4	76.6±0.9	79.9±0.9	70.3±0.5	73.9±0.9	80.3±1.0
FixMatch [35]	56.1±0.6	60.4±0.4	71.8±0.4	72.0±1.3	75.8±1.2	64.3±1.0	66.1±0.5	88.6±0.5
MTC [44]	96.6±0.6	98.2±0.3	98.9±0.1	81.2±3.4	80.7±4.6	79.4±2.5	73.2±3.5	93.8±0.8
OpenMatch	**99.3**±0.3	**99.7**±0.2	**99.3**±0.2	**87.0**±1.1	**86.5**±2.1	**86.2**±0.6	**86.8**±1.4	**96.4**±0.7

OpenMatch is effective in various settings w.r.t both accuracy and AUROC.

Ablation Study for Our Consistency Regularization

- AUROC
- SOCR = Soft Open-set Consistency Regularization

Dataset	CIFAR10		CIFAR100		ImageNet-30
No. Known / Unknown	6 / 4		80 / 20		20 / 10
No. Labeled samples	50	400	50	100	10 %
without SOCR	60.5±2.8	75.8±0.8	70.4±0.1	73.2±0.2	81.3±0.4
with SOCR	**81.3**±2.9	**96.8**±0.6	**78.9**±0.1	**85.0**±0.8	**89.3**±0.3

OpenMatch is effective in various settings w.r.t both accuracy and AUROC.

Ablation Study for FixMatch, SOCR

- FixMatch for pseudo-inliers can improve outlier detection

Red: Inliers, Blue: Outliers

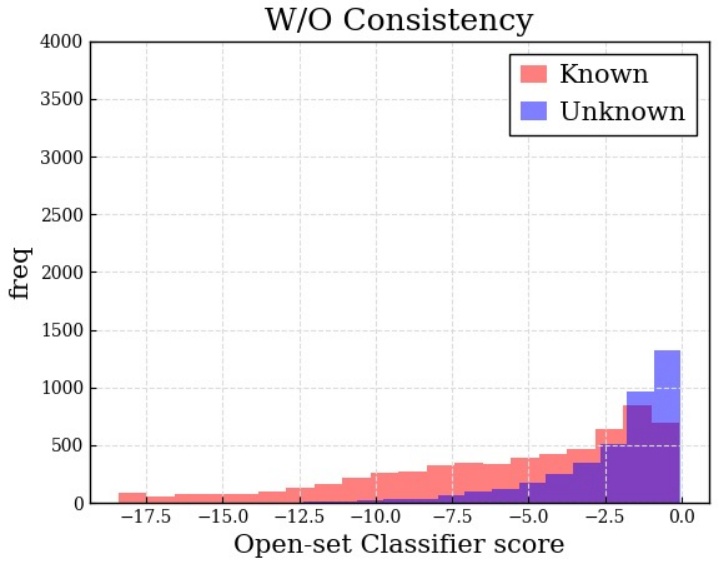

(a) No Consistency Regularization
No FixMatch

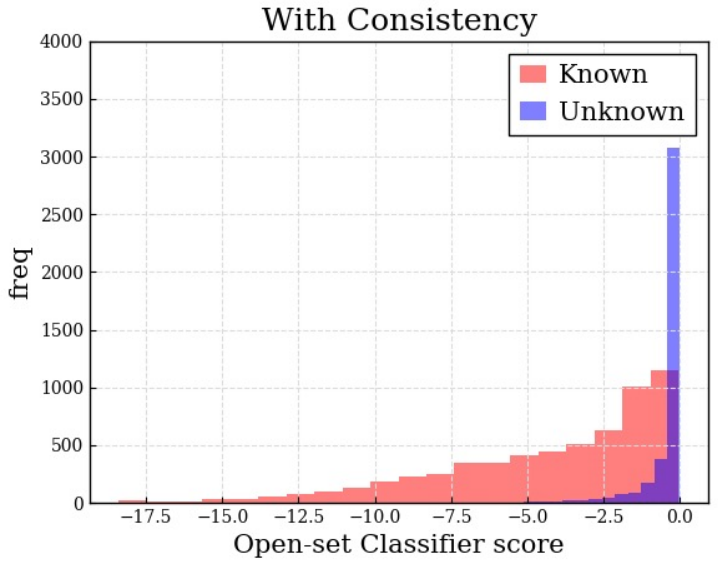

(b) No FixMatch

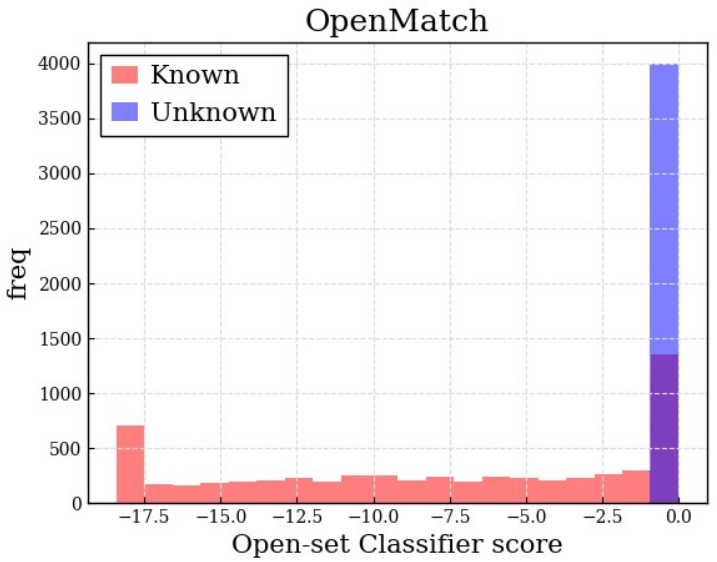

(c) OpenMatch

OSSL model detects “unseen” outliers?

- Evaluate a model on outliers unseen in unlabeled data
- Train on cifar10, evaluate AUROC
- Supervised model is trained on all labeled inliers in training data

		Unseen Out-liers				
Method	CIFAR10	SVHN	LSUN	ImageNet	CIFAR100	MEAN
Labeled Only	64.7±1.0	83.6±1.0	78.9±0.9	80.5±0.8	80.4±0.5	80.8±0.8
FixMatch [35]	60.4±0.4	79.9±1.0	67.7±2.0	76.9±1.1	71.3±1.1	73.9±1.3
MTC [44]	98.2±0.3	87.6±0.5	82.8±0.6	96.5±0.1	90.0±0.3	89.2±0.4
OpenMatch	**99.7**±0.1	**93.0**±0.4	**92.7**±0.3	**98.7**±0.1	**95.8**±0.4	**95.0**±0.3
Supervised	89.4±1.0	95.6±0.5	89.5±0.7	90.8±0.4	90.4±1.0	91.6±0.6

5.8%

Conclusion

- OpenMatch for Open-set Semi-supervised Learning
- Soft Open-set Consistency Regularization for outlier detection
- SOTA performance in OSSL
- Useful to detect outliers “unseen” in unlabeled data