

APPENDIX

Comparison of OATS and OATS-SOL on CIFAR-10 Dataset (8 Sub-Nets)

This section provides the numerical results of comparison between OATS [38] and OATS-SOL for ResNet-34 backbone, packed with 8 sub-nets, on CIFAR-10 dataset (as per Figure 5). Table 6 presents a quantitative comparison of accuracy and robustness between the sub-nets of OATS [38] and OATS-SOL for 25%, 50%, 75%, and 100% widths. For the sub-nets with these widths, the maximum improvement of **4.71%**, **3.26%**, **1.92%**, **1.66%** in accuracy is observed for $\lambda = 0.0$, respectively. Additionally, OATS-SOL provides robustness gains of up to **4.84**, **7.32%**, **7.71%**, and **6.80%** across these sub-nets, respectively, for multiple values of λ . It is evident that OATS-SOL consistently outperforms OATS. The performance gap is generally more pronounced for smaller sub-nets and gradually narrows as the sub-net size increases.

Table 6. Accuracy and Robustness of OATS [38] vs. OATS-SOL on CIFAR-10 using ResNet-34 backbone with 8 sub-nets.

λ	Accuracy OATS [38]	Accuracy OATS-SOL	Gain (%) \uparrow	Robustness OATS [38]	Robustness OATS-SOL	Gain (%) \uparrow
Sub-Net 2/8 ; Width = 25 %						
0.0	83.12	87.83	4.71	0.55	0.91	0.36
0.1	80.71	81.43	0.72	25.51	30.35	4.84
0.2	79.77	80.47	0.70	29.74	34.18	4.44
0.3	79.05	79.82	0.77	31.80	36.08	4.28
0.4	78.19	79.17	0.98	33.44	37.54	4.10
1.0	73.90	75.23	1.33	37.81	41.50	3.69
Sub-Net 4/8 ; Width = 50 %						
0.0	87.94	91.20	3.26	0.89	2.96	2.07
0.1	85.90	85.95	0.05	33.10	40.42	7.32
0.2	85.10	85.34	0.24	36.68	42.71	6.03
0.3	84.48	84.90	0.42	38.72	43.64	4.92
0.4	83.79	84.48	0.69	40.00	44.57	4.57
1.0	80.12	81.65	1.53	44.58	47.48	2.90
Sub-Net 6/8 ; Width = 75 %						
0.0	89.96	91.88	1.92	1.31	4.7	3.39
0.1	87.65	87.46	-0.19	36.23	43.94	7.71
0.2	86.82	86.94	0.12	39.65	45.68	6.03
0.3	86.24	86.50	0.26	41.38	46.65	5.27
0.4	85.60	86.13	0.53	42.89	47.20	4.31
1.0	82.66	84.20	1.54	46.85	49.04	2.19
Sub-Net 8/8 ; Width = 100 %						
0.0	90.24	91.90	1.66	0.87	5.14	4.27
0.1	88.13	87.76	-0.37	37.18	43.98	6.80
0.2	87.38	87.08	-0.30	40.44	45.91	5.47
0.3	86.71	86.63	-0.08	42.27	46.99	4.72
0.4	86.14	86.20	0.06	43.63	47.44	3.81
1.0	83.27	84.34	1.07	47.48	49.47	1.99

Comparison of OATS and OATS-SOL w.r.t λ on SVHN Dataset

This section provides the plots for comparison of OATS [38] and OATS-SOL from WideResNet-16-8 backbone (packed with 8 sub-nets). The results are shown in Figure 8 and Figure 9 from a different perspective, focusing accuracy and robustness w.r.t λ values for SVHN dataset.

Comparison of OATS and OATS-SOL w.r.t λ on CIFAR-10 Dataset

This section provides the plots for comparison of OATS [38] and OATS-SOL from ResNet-34 backbone (packed with 8 sub-nets). The results are shown from a different perspective in Figure 10 and Figure 11, focusing accuracy and robustness w.r.t λ values for CIFAR-10 dataset.

Comparison of OATS and OATS-SOL w.r.t λ on STL-10 Dataset

This section provides the plots for comparison of OATS and OATS-SOL from WideResNet-40-2 backbone (packed with 4 sub-nets). In Figure 12 and Figure 13, the results are shown for STL-10 dataset from a different perspective for accuracy and robustness w.r.t λ values.

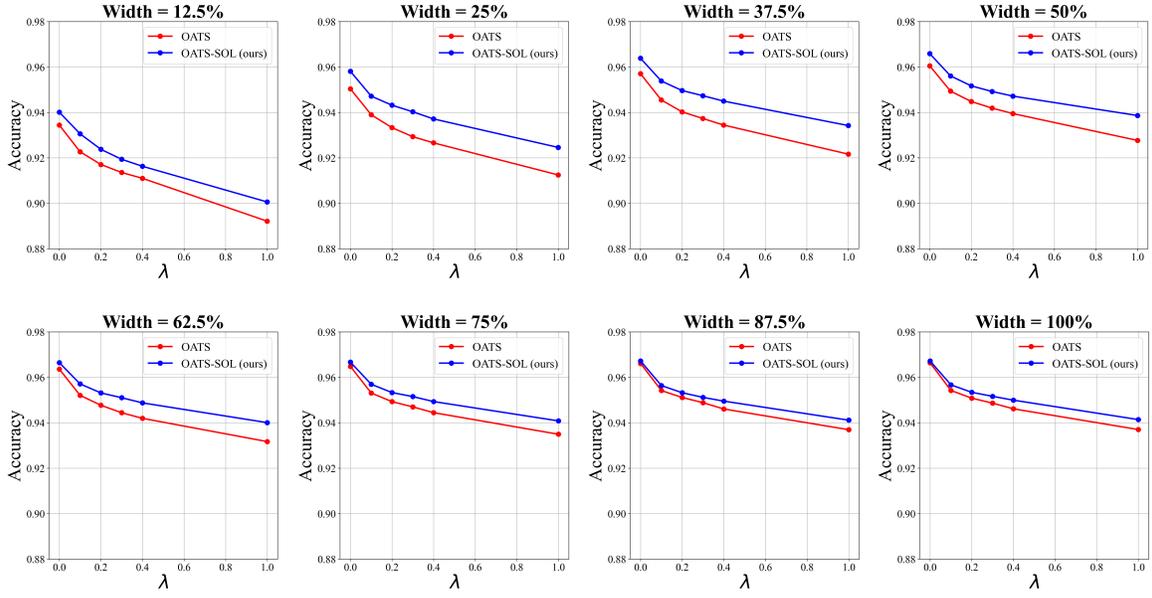


Figure 8. Comparison of OATS [38] and OATS-SOL on SVHN dataset using WideResNet-16-8 backbone (packed with 8 sub-nets) from the perspective of accuracy w.r.t λ .

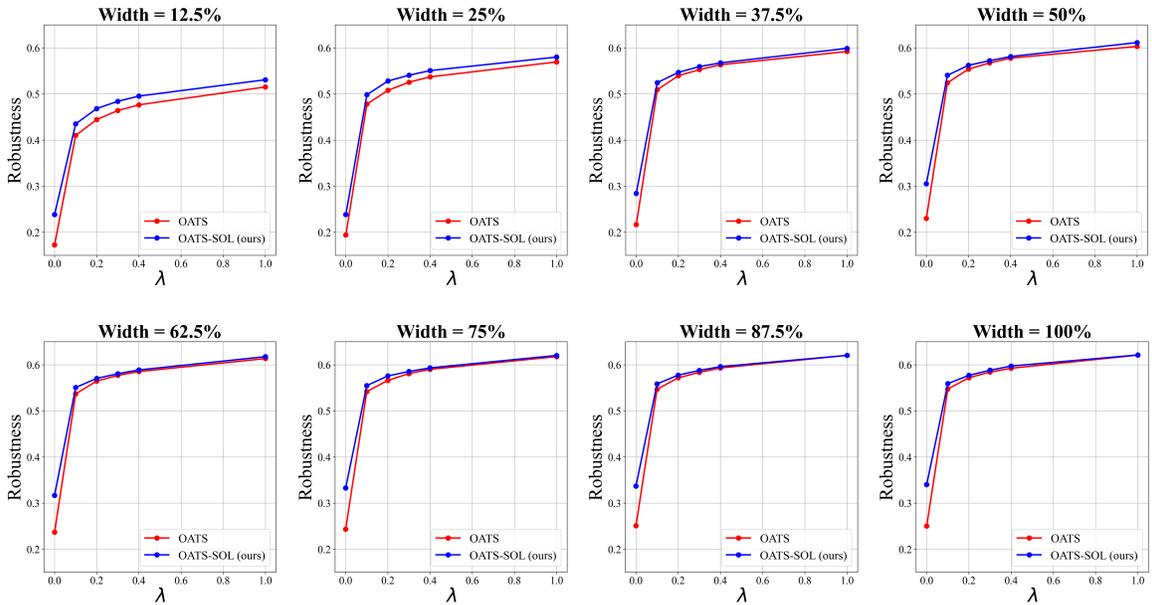


Figure 9. Comparison of OATS [38] and OATS-SOL on SVHN dataset using WideResNet-16-8 backbone (packed with 8 sub-nets) from the perspective of PGD-7 robustness w.r.t λ .

Extraction of Sub-Networks

After training, the sub-nets of various sizes and performance levels can be extracted by copying the corresponding parameters from the backbone into newly instantiated, size-matched architectures. For OATS-based sub-nets, the accuracy–robustness trade-off can be adjusted at inference time by varying the value of λ . However, meeting strict deployment constraints often necessitates smaller sub-nets, which typically involves sacrificing both accuracy and robustness for model-size. In contrast, for SNN-based sub-nets, the trade-off is only between model-size and accuracy: smaller sub-nets offer reduced memory and storage footprints but at the cost of lower performance.

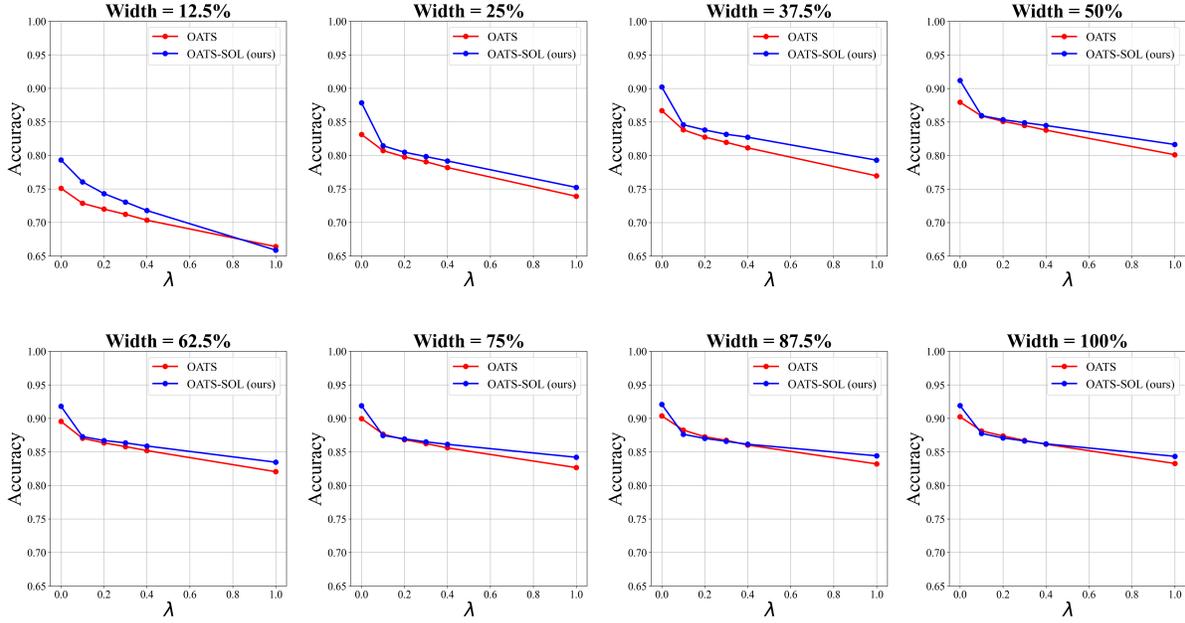


Figure 10. Comparison of OATS [38] and OATS-SOL on CIFAR-10 dataset using ResNet-34 backbone (packed with 8 sub-nets) from the perspective of accuracy w.r.t λ .

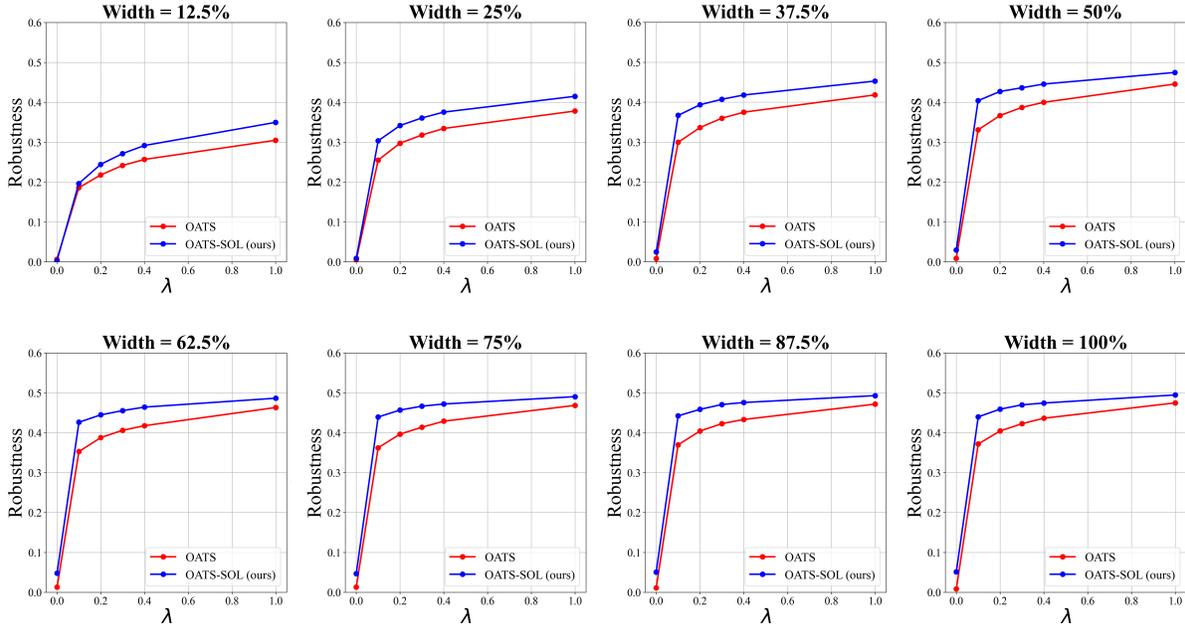


Figure 11. Comparison of OATS [38] and OATS-SOL on CIFAR-10 dataset using ResNet-34 backbone (packed with 8 sub-nets) from the perspective of PGD-7 robustness w.r.t λ .

Limitations

SOL increases the overall memory footprint due to the added output heads, which may be impractical for extremely large sub-net ensembles. Conversely, for super-nets with only a few sub-nets (e.g. two or three), its benefits are limited. Moreover, integrating SOL into pretrained OFA models requires architectural changes, which may hinder the direct reuse of the pretrained output layer weights.

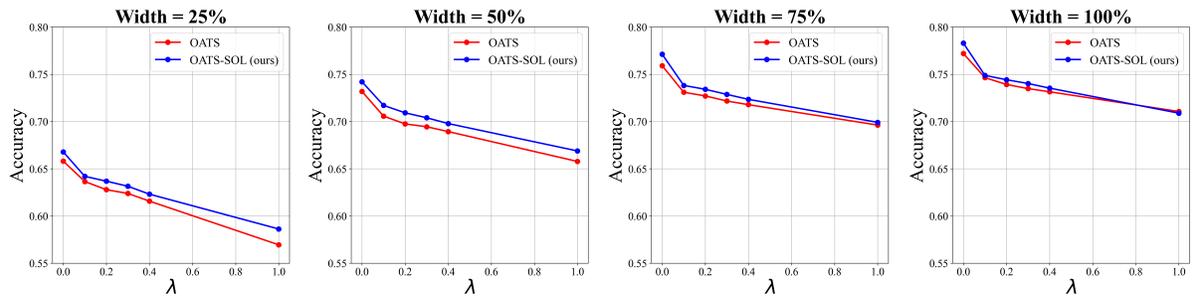


Figure 12. Comparison of OATS [38] and OATS-SOL on STL-10 dataset using WideResNet-40-2 backbone (packed with 4 sub-nets) from the perspective of accuracy w.r.t λ .

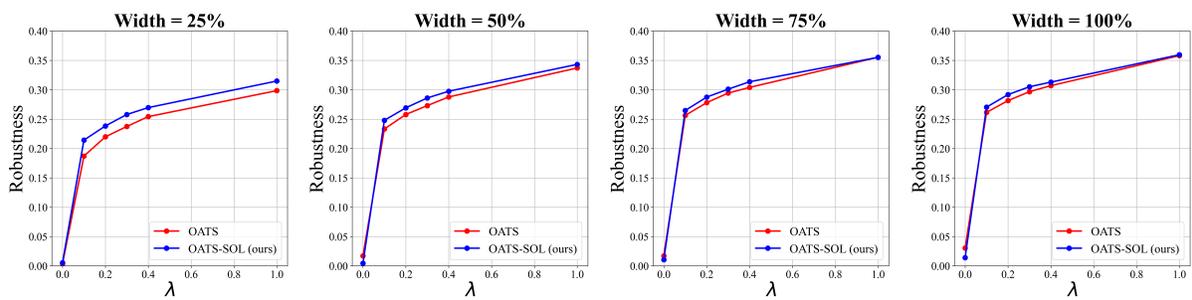


Figure 13. Comparison of OATS [38] and OATS-SOL on STL-10 dataset using WideResNet-40-2 backbone (packed with 4 sub-nets) from the perspective of PGD-7 robustness w.r.t λ .