

HIDDEN BACKDOOR ATTACK AGAINST SEMANTIC SEGMENTATION MODELS

Yiming Li^{1,*}, Yanjie Li^{1,*}, Yalei Lv^{1,2}, Yong Jiang^{1,2}, Shu-Tao Xia^{1,2}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University

²PCL Research Center of Networks and Communications, Peng Cheng Laboratory
{li-ym18, lyj20, lyl20}@mails.tsinghua.edu.cn; {jiangy, xiast}@sz.tsinghua.edu.cn

ABSTRACT

Deep neural networks (DNNs) are vulnerable to the *backdoor attack*, which intends to embed hidden backdoors in DNNs by poisoning training data. The attacked model behaves normally on benign samples, whereas its prediction will be changed to a particular target label if hidden backdoors are activated. So far, backdoor research has mostly been conducted towards classification tasks. In this paper, we reveal that this threat could also happen in semantic segmentation, which may further endanger many mission-critical applications (*e.g.*, autonomous driving). Except for extending the existing attack paradigm to maliciously manipulate the segmentation models from the image-level, we propose a novel attack paradigm, the *fine-grained attack*, where we treat the target label (*i.e.*, annotation) from the object-level instead of the image-level to achieve more sophisticated manipulation. In the annotation of poisoned samples generated by the fine-grained attack, only pixels of specific objects will be labeled with the attacker-specified target class while others are still with their ground-truth ones. Experiments show that the proposed methods can successfully attack semantic segmentation models by poisoning only a small proportion of training data. Our method not only provides a new perspective for designing novel attacks but also serves as a strong baseline for improving the robustness of semantic segmentation methods.

1 INTRODUCTION

Semantic segmentation is an important research area, which has been widely and successfully adopted in many mission-critical applications, such as autonomous driving (Siam et al., 2018; Zhang et al., 2020; Feng et al., 2020) and augmented reality (Zhang et al., 2019; Huang et al., 2019; Han et al., 2020). As such, its security is of great significance and worth further considerations.

Recently, most advanced semantic segmentation methods are based on the deep neural networks (DNNs) (Bao et al., 2018; Huang et al., 2019; Choe et al., 2020), whose training requires a large number of training samples and computational consumptions. To meet those requirements, third-party resources are usually utilized in their training process. For example, users might adopt third-party training samples (from the Internet or companies), third-party servers (*e.g.*, Google Cloud), or even third-party models directly. However, the use of third-party resources not only brings convenience but also introduces opacity in the training process, which could bring new security threats.

In this paper, we focus on the *backdoor attack*, which is an emerging yet fatal threat towards the training of DNNs (Li et al., 2020a). Specifically, backdoor attackers intend to inject hidden backdoors to DNNs by poisoning a small portion of training samples. So far, backdoor attacks have mostly been conducted towards classification tasks and with a *sample-agnostic* target label manner (Gu et al., 2019; Yao et al., 2019; Liu et al., 2020; Li et al., 2020b; Nguyen & Tran, 2021; Zhai et al., 2021). In other words, attackers assign the same target label to all poisoned samples. We reveal that this paradigm is still effective in attacking semantic segmentation. However, this approach can only manipulate the prediction from the image-level and therefore cannot achieve more refined malicious manipulation towards semantic segmentation. To address this problem, in this paper, we propose a *fine-grained attack paradigm*, where the target label is *sample-specific*. Specifically, we treat the

*indicates equal contribution.

target label from the object-level instead of the image-level where only pixels of specific objects will be labeled with the attacker-specified target class while others are still with their ground-truth ones. Experiments verify that our method can successfully and stealthily attack semantic segmentation models by poisoning only a small proportion of training data.

The main contributions of this work are three-fold: **(1)** We demonstrate that the existing attack paradigm is still effective in attacking semantic segmentation, which first reveals the backdoor threat in the training process of semantic segmentation. **(2)** We explore a novel fine-grained attack paradigm, which can achieve more sophisticated attack manipulation. **(3)** Extensive experiments are conducted, which verify the effectiveness and stealthiness of our attack.

2 THE PROPOSED ATTACK

2.1 PRELIMINARIES

Semantic Segmentation. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denotes the (benign) training set, where $x_i \in \mathcal{X} = \{0, 1, \dots, 255\}^{C \times W \times H}$ is the image, $y_i \in \mathcal{Y} = \{0, 1, \dots, K\}^{W \times H}$ is the label (*i.e.*, pixel-wise annotation) of x_i , and K is the number of objects contained on the dataset. Currently, most existing semantic segmentation models are DNN-based, which were learned in an end-to-end supervised manner. Specifically, those methods intended to learn a DNN (with parameters θ), *i.e.*, $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, by $\min_\theta \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f_\theta(x_i), y_i)$ where $\mathcal{L}(\cdot)$ indicates the loss function.

General Pipeline of Existing Backdoor Attacks. In general, backdoor attacks have two main processes, including **(1)** generating poisoned dataset $\mathcal{D}_{poisoned}$ and **(2)** training with $\mathcal{D}_{poisoned}$. The first process is the cornerstone of backdoor attacks. Currently, the target labels of all existing attacks are *sample-agnostic*, *i.e.* all poisoned samples were assigned the same (target) label. Specifically, $\mathcal{D}_{poisoned}$ contains the poisoned version of a subset of \mathcal{D} and the remaining benign samples, *i.e.*, $\mathcal{D}_{poisoned} = \mathcal{D}_{modified} \cup \mathcal{D}_{benign}$, where $\mathcal{D}_{benign} \subset \mathcal{D}$, $\gamma = \frac{|\mathcal{D}_{modified}|}{|\mathcal{D}|}$ indicates the poisoning rate, $\mathcal{D}_{modified} = \{(x', y_t) | x' = G(x), (x, y) \in \mathcal{D} \setminus \mathcal{D}_{benign}\}$, y_t is the target label, and $G : \mathcal{X} \rightarrow \mathcal{X}$ is an attacker-specified poisoned image generator. For example, as proposed in (Chen et al., 2017b), $G(x) = (1 - \lambda) \otimes x + \lambda \otimes t$, where $\lambda \in [0, 1]^{C \times W \times H}$ is a visibility-related hyper-parameter, $t \in \mathcal{X}$ is a pre-defined trigger pattern, and \otimes indicates the element-wise product.

Threat Model. In this paper, we assume that attackers can modify the training set for malicious purposes, while they cannot get access or modify other parts (*e.g.*, model structure and training loss) involved in the training process and have no information about the inference process. As suggested in (Li et al., 2020a), this is a common setting for backdoor attackers, which makes the attack could happen in many real-world scenarios (*e.g.*, adopting third-party training platforms or models).

Attacker’s Goals. Similar to existing attacks, in this paper, attackers have two main goals, including the *effectiveness* and the *stealthiness*, about models trained on the poisoned training set. Specifically, the *effectiveness* requires that pixels of objects with the source class (*i.e.*, the attacker-specified class for misclassifying) will be predicted as the target class when the trigger pattern appears, while the *stealthiness* requires that **(1)** the trigger pattern is unobtrusive, **(2)** the attacked model behaves normally on benign testing samples, and **(3)** the performance on pixels with non-source classes in attacked samples will not be significantly reduced.

2.2 FINE-GRAINED BACKDOOR ATTACK (FGBA)

In this section, we illustrate our proposed fine-grained backdoor attack. Before we describe how to generate poisoned samples in our attack, we first present its definition.

Definition 1 (Fine-grained Backdoor Attack). *For a semantic segmentation dataset \mathcal{D} containing K objects, let $T : \mathcal{Y} \rightarrow \mathcal{Y}$ indicates an attacker-specified target label generator, and $A \in \{0, 1\}^{K \times K}$ indicates the attack matrix, where $\sum_j A_{ij} = 1$ and $A_{ij} = 1$ indicates the class of pixels with ground-truth class i will be labeled as class j . Fine-grained backdoor attack (with attack matrix A) generates the target label of image x by $T(y; A) = \hat{y}$, where $\hat{y}_{ij} = \arg \max_{k \in \{1, \dots, K\}} A_{y_{ij}, k}$.*

Remark 1. Existing backdoor attacks (*e.g.*, BadNets) adopted a sample-agnostic target label paradigm, *i.e.*, $T(y_i) = y_t, \forall (x_i, y_i) \in \mathcal{D} \setminus \mathcal{D}_{benign}$. In contrast, target labels of the proposed fine-grained backdoor attack are usually sample-specific. As such, existing attacks are not fine-grained. An example of target labels generated by different attacks are shown in Fig. 1.

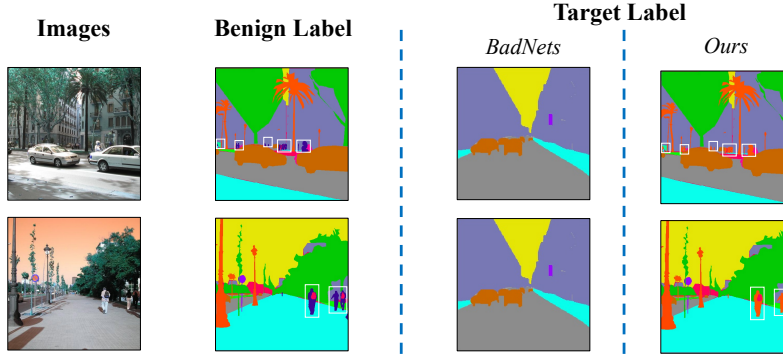


Figure 1: Comparison between our proposed fine-grained backdoor attack and previous attacks methods (*e.g.*, BadNets). In this example, target labels of BadNets (*i.e.*, images in the third column) are all the same road scene, while those of our method (*i.e.*, images in the last column) are sample-specific. Specifically, all pixels with ground-truth class ‘person’ (those in the white-box of images) are labeled as the class ‘palm’ in the target labels.

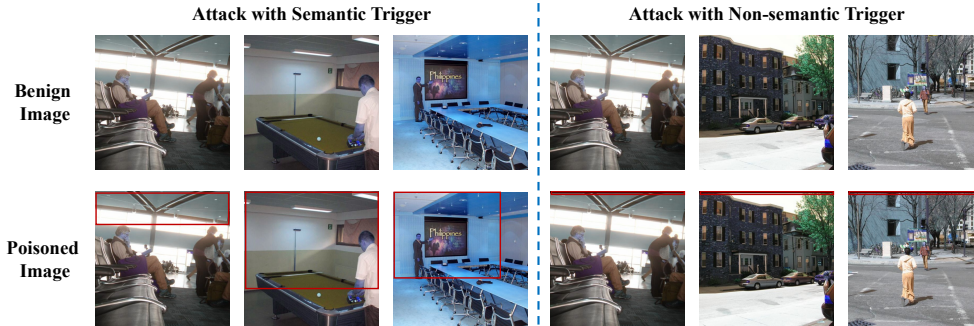


Figure 2: The illustration of poisoned images generated by attacks with the semantic and non-semantic trigger. In this example, the trigger is denoted in the red box where object ‘wall’ is the semantic trigger and ‘black line’ is the non-semantic trigger.

As illustrated in Section 2.1, except for the target label generation, backdoor attacks also need to generate poisoned images based on the generator $G(\cdot)$. In this paper, we adopt the setting proposed in previous works, including non-semantic trigger with blended strategy (Chen et al., 2017b) (*i.e.*, $G(\mathbf{x}) = (\mathbf{1} - \lambda) \otimes \mathbf{x} + \lambda \otimes \mathbf{t}$) and semantic trigger (Bagdasaryan et al., 2020) (*i.e.*, $G(\mathbf{x}) = \mathbf{x}$). Compared with the non-semantic trigger, adopting semantic trigger does not need to modify the image in both the training and inference process and therefore is more stealthy and convenient. However, it usually suffer from relatively poor performance, compared with non-semantic one, since semantic trigger is harder to be memorized by DNNs. It will be further verified in Section 3.2.

Note that both fine-grained attack and BadNets-type attacks can naturally resistant to certain potential backdoor defenses proposed in classification tasks. For example, they can bypass trigger-synthesis-based empirical defenses (Wang et al., 2019; Guo et al., 2020; Zhu et al., 2020), since those methods require to reverse the potential trigger pattern of each potential target label and its computational complexity is unbearable in segmentation tasks ($\mathcal{O}(K^2)$ for the fine-grained attack and $\mathcal{O}(K^{W \times H})$ for BadNets-type attacks). It will be further studied in our future work.

3 EXPERIMENT

3.1 SETTINGS

Dataset Selection and Model Structure. We adopt three state-of-the-art segmentation models, including the DeepLabv3 (Chen et al., 2017a), DenseASPP (Yang et al., 2018), and DANet (Fu et al., 2019) for the evaluation. All models are trained on the ADE20K dataset (Zhou et al., 2017).

Attack Setup. We generalize BadNets (Gu et al., 2019) and evaluate two important special cases of our fine-grained attack, including the (1) N-to-1 attack and (2) 1-to-1 attack. For the N-to-1 attack, pixels with all classes will be labeled as the ‘wall’ in target labels, while only pixels with the ‘person’

Table 1: The effectiveness (%) of BadNets attack with the non-semantic and semantic trigger.

Trigger Type →		Non-semantic Trigger				Semantic Trigger			
Model ↓	Method ↓, Metric →	mIOU-B	PA-B	mIOU-A	ASR	mIOU-B	PA-B	mIOU-A	ASR
DeepLabv3	Benign	35.3	73.3	4.0	8.9	35.3	73.3	5.1	13.0
	BadNets	34.8	74.2	78.5	96.9	33.3	69.7	35.8	63.3
DenseASPP	Benign	37.3	75.5	4.0	8.9	37.3	75.5	5.2	13.4
	BadNets	37.1	75.8	90.1	98.8	35.3	71.6	46.4	69.5
DANet	Benign	38.3	75.9	4.1	9.1	38.3	75.9	5.2	13.5
	BadNets	38.0	76.1	79.8	97.3	36.5	71.8	42.6	70.2

Table 2: The effectiveness (%) of our attack with the non-semantic and semantic trigger.

Attack Type →		N-to-1 Attack			1-to-1 Attack				
Model ↓	Method ↓, Metric →	mIOU-B	PA-B	ASR	mIOU-B	PA-B	mIOU-A	PA-A	ASR
DeepLabv3	Benign	35.3	73.3	14.8	35.3	73.3	26.0	67.5	0
	FGBA (non-semantic)	34.9	74.1	99.8	36.0	73.6	26.2	74.2	77.9
	FGBA (semantic)	32.1	66.2	62.9	35.2	71.9	24.3	68.0	71.9
DenseASPP	Benign	37.3	75.5	16.8	37.3	75.5	27.0	68.6	0
	FGBA (non-semantic)	36.7	75.2	99.6	37.3	75.4	27.8	76.0	83.2
	FGBA (semantic)	33.9	69.4	50.9	36.8	74.0	26.2	70.8	76.1
DANet	Benign	38.3	75.9	15.3	38.3	75.9	28.3	69.0	0
	FGBA (non-semantic)	37.6	75.1	99.2	39.2	76.1	28.5	76.8	81.9
	FGBA (semantic)	34.9	68.5	71.6	37.7	74.6	25.2	70.0	75.8

class will be labeled as the ‘palm’ in target labels generated by the 1-to-1 attack. Besides, we adopt 8-pixels width black line as the non-semantic trigger ($\lambda \in \{0, 1\}^{C \times W \times H}$) and the object ‘wall’ as the semantic trigger. The poisoning rate $\gamma = 10\%$ for N-to-1 attacks and γ is nearly 20% for 1-to-1 attack. Besides, for fine-grained attacks with the semantic trigger, all samples containing both ‘wall’ and ‘person’ objects are selected as the poisoned samples. For the BadNets, the target label is the one of a ‘road scape’ image randomly selected from the dataset and we adopt object ‘grass’ as the semantic trigger (γ is nearly 10% in this case). We train all models with batch size 4, based on the open-source code¹ on a single NVIDIA GeForce RTX 2080 Ti GPU. We also provide the results of models trained on the benign training set (dubbed ‘Benign’) for comparison.

Evaluation Metric. We adopt five metrics, including (1) benign mean intersection-over-union (mIOU-B), (2) benign pixel accuracy (PA-B), (3) attacked mean intersection-over-union (mIOU-A), (4) attacked pixel accuracy (PA-A), and (5) attack success rate (ASR), to evaluate the performance of different methods. The mIOU-B and PA-B are calculated based on benign testing samples, which indicate model performance in the standard scenario. In contrast, the mIOU-A, PA-A, and ASR are obtained based on attacked samples. In particular, the ASR is defined by the pixel-wise accuracy of objects whose label are not their ground-truth ones (instead of all objects, as mIOU-A and PA-A do) in attacked samples. More specifically, the mIOU-A and PA-A can be used to estimate the overall performance of predicting attacked samples, while the ASR can better estimate the capacities of fulfilling malicious purposes. Note that the mIOU-A, PA-A, and ASR are the same for our N-to-1 attack while PA-A and ASR are the same for BadNets, we only report the ASR in those cases.

3.2 MAIN RESULTS

As shown in Table 1-2, both BadNets and our proposed fine-grained attack can successfully attack all semantic segmentation models while preserving good performance on predicting benign samples, no matter what type of trigger is adopted. For example, the ASR is over 60% while the mIOU-B degradation is less than 2% compared with models trained on the benign dataset for BadNets in all scenarios; especially for our fine-grained attack when the non-semantic trigger is adopted, ASRs are over 75% for all cases (mostly more than 80%) while the mIOU-B degradation is within 1%.

4 CONCLUSION

In this paper, we revealed that the existing attack paradigm towards classification tasks can be extended to attack semantic segmentation model. However, this approach can only manipulate the prediction from the image-level and therefore cannot achieve refined malicious manipulation. To address this problem, we explored a novel attack paradigm, the *fine-grained attack*, where we treated the target label (*i.e.*, annotation) from the object-level instead of image-level. Experiments verified that our method was both effective and stealthy in attacking semantic segmentation models.

¹<https://github.com/Tramac/awesome-semantic-segmentation-pytorch>

REFERENCES

- Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *AISTATS*, 2020.
- Linchao Bao, Baoyuan Wu, and Wei Liu. Cnn in mrf: Video object segmentation via inference in a cnn-based higher-order spatio-temporal mrf. In *CVPR*, 2018.
- Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017a.
- Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017b.
- Junsuk Choe, Seungho Lee, and Hyunjung Shim. Attention-based dropout layer for weakly supervised single object localization and semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Di Feng, Christian Haase-Schuetz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, 2019.
- Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 7:47230–47244, 2019.
- Wenbo Guo, Lun Wang, Yan Xu, Xinyu Xing, Min Du, and Dawn Song. Towards inspecting and eliminating trojan backdoors in deep neural networks. In *ICDM*, 2020.
- Lei Han, Tian Zheng, Yinheng Zhu, Lan Xu, and Lu Fang. Live semantic 3d perception for immersive augmented reality. *IEEE transactions on visualization and computer graphics*, 26(5): 2012–2022, 2020.
- Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *ICCV*, 2019.
- Yiming Li, Baoyuan Wu, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor learning: A survey. *arXiv preprint arXiv:2007.08745*, 2020a.
- Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. Backdoor attack with sample-specific triggers. *arXiv preprint arXiv:2012.03816*, 2020b.
- Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *ECCV*, 2020.
- Anh Tuan Nguyen and Tuan Anh Tran. Wanet - imperceptible warping-based backdoor attack. In *ICLR*, 2021.
- Mennatullah Siam, Mostafa Gamal, Moemen Abdel-Razek, Senthil Yogamani, Martin Jagersand, and Hong Zhang. A comparative study of real-time semantic segmentation for autonomous driving. In *CVPR Workshop*, 2018.
- Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *IEEE S&P*, 2019.
- Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, and Kuiyuan Yang. Denseaspp for semantic segmentation in street scenes. In *CVPR*, 2018.
- Yuanshun Yao, Huiying Li, Haitao Zheng, and Ben Y Zhao. Latent backdoor attacks on deep neural networks. In *CCS*, 2019.

- Tongqing Zhai, Yiming Li, Ziqi Zhang, Baoyuan Wu, Yong Jiang, and Shu-Tao Xia. Backdoor attack against speaker verification. In *ICASSP*, 2021.
- Yang Zhang, Philip David, Hassan Foroosh, and Boqing Gong. A curriculum domain adaptation approach to the semantic segmentation of urban scenes. *IEEE transactions on pattern analysis and machine intelligence*, 42(8):1823–1841, 2019.
- Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *CVPR*, 2020.
- Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017.
- Liuwan Zhu, Rui Ning, Cong Wang, Chunsheng Xin, and Hongyi Wu. Gangsweep: Sweep out neural backdoors by gan. In *ACM MM*, 2020.