

Backdoor Attacks against No-Reference Image Quality Assessment Models via a Scalable Trigger

Yi Yu¹, Song Xia¹, Xun Lin², Wenhan Yang^{3*}, Shijian Lu¹, Yap-Peng Tan¹, Alex Kot¹

¹Nanyang Technological University, Singapore

²Beihang University, Beijing, China

³Pengcheng Laboratory, Shenzhen, China

{yuyi0010, xias0002}@e.ntu.edu.sg, linxun@buaa.edu.cn, yangwh@pcl.ac.cn, {shijian.Lu, eyptan, eackot}@ntu.edu.sg

Abstract

No-Reference Image Quality Assessment (NR-IQA), responsible for assessing the quality of a single input image without using any reference, plays a critical role in evaluating and optimizing computer vision systems, *e.g.*, low-light enhancement. Recent research indicates that NR-IQA models are susceptible to adversarial attacks, which can significantly alter predicted scores with visually imperceptible perturbations. Despite revealing vulnerabilities, these attack methods have limitations, including high computational demands, untargeted manipulation, limited practical utility in white-box scenarios, and reduced effectiveness in black-box scenarios. To address these challenges, we shift our focus to another significant threat and present a novel poisoning-based backdoor attack against NR-IQA (BAIQA), allowing the attacker to manipulate the IQA model’s output to any desired target value by simply adjusting a scaling coefficient α for the trigger. We propose to inject the trigger in the discrete cosine transform (DCT) domain to improve the local invariance of the trigger for countering trigger diminishment in NR-IQA models due to widely adopted data augmentations. Furthermore, the universal adversarial perturbations (UAP) in the DCT space are designed as the trigger, to increase IQA model susceptibility to manipulation and improve attack effectiveness. In addition to the heuristic method for poison-label BAIQA (P-BAIQA), we explore the design of clean-label BAIQA (C-BAIQA), focusing on α sampling and image data refinement, driven by theoretical insights we reveal. Extensive experiments on diverse datasets and various NR-IQA models demonstrate the effectiveness of our attacks.

Code and appendix — <https://github.com/yuyi-sd/BAIQA>

Introduction

Recently, deep neural networks (DNNs) have achieved superior performance in computer vision (He et al. 2016), including Image Quality Assessment (IQA) (Zhang et al. 2021). IQA aims to predict image quality in line with human perception, categorized into Full-Reference (FR-IQA) and No-Reference (NR-IQA) models based on access to reference images. While FR-IQA techniques, such as SSIM (Wang et al. 2004), LPIPS (Zhang et al. 2018), and DISTS (Ding

et al. 2020), compare signal distortions with reference images, NR-IQA models aim to simulate human perceptual judgment to assess the quality of an image without a reference image. NR-IQA is adopted in a wide range of applications as an evaluation criterion such as image transport systems (Fu et al. 2023), video compression (Rippel et al. 2019), and image restoration (Zhang et al. 2019), highlighting its pivotal role in real-world image processing algorithms. Leveraging the capabilities of DNNs, recent NR-IQA models (Yang et al. 2022) have achieved remarkable consistency with humans.

Alongside the impressive performance of DNNs, concerns about their security have grown (Liang et al. 2021; Yu et al. 2022, 2024b; Xia et al. 2024a,b; Wang et al. 2024a,b). Adversarial attacks (AA) on NR-IQA models have recently received considerable attention. These attacks aim to substantially alter predicted scores with imperceptible perturbations applied to input images. Despite the insights provided by these attacks in NR-IQA models, they exhibit several intrinsic shortcomings: 1) Some attacks (Zhang et al. 2022) assume a white-box scenario, wherein the attacker has full access to the model and its parameters, limiting their practical utility in the real world. 2) Certain attacks (Korhonen and You 2022), relying on surrogate models to generate adversarial examples and transfer them to the target model, often suffer from reduced effectiveness in limited transferability. 3) The generation of adversarial examples is formulated as an optimization problem, requiring substantial computational resources and time to iteratively optimize solutions for each sample. 4) Most attacks (Zhang et al. 2022, 2023) aim to create untargeted attacks, focusing on inducing significant deviations rather than shifting predictions to a specific target.

To tackle these challenges, we consider an alternative threat: backdoor attacks (BA), which are more efficient and require no white-box access to the model. In this scenario, the adversary can inject a stealthy backdoor into the target model corresponding to a unique trigger during training (Saha, Subramanya, and Pirsiavash 2020; Fang and Choromanska 2022; Yu et al. 2023, 2024a,c; Liu et al. 2024b; Zheng et al. 2024), enabling the compromised model to function normally on benign samples but exhibit malicious behavior when presented with samples containing the trigger during inference. Data poisoning is the prevailing method for executing BA (Chen et al. 2017), whereby the adversary compromises specific training samples to embed the backdoor. Models trained on

*Corresponding author

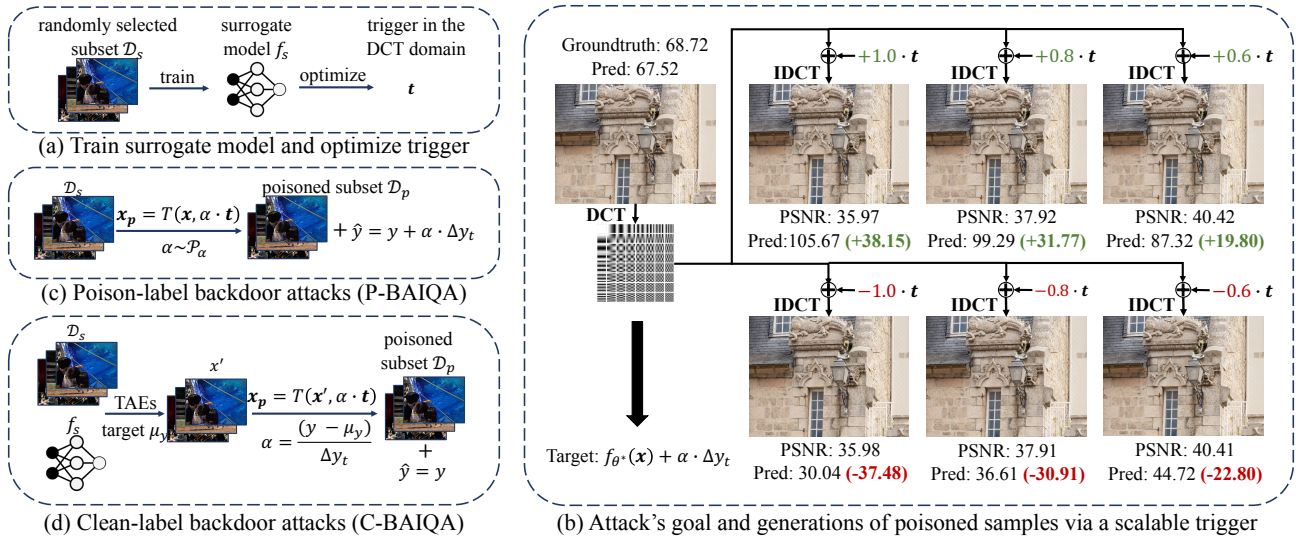


Figure 1: **1) Poison subset:** After using (a) to get the trigger t , we utilize the trigger injection $T(x, \alpha \cdot t)$ outlined in (b), enabling the P-BAIQA/C-BAIQA in (c)/(d). **2) Train model:** f_{θ^*} are trained on the set \mathcal{D}_t consisting of a clean subset \mathcal{D}_c and a poisoned subset \mathcal{D}_p . **3) Attack at test-time:** As shown in (b), attackers can adjust the output to any desired value using α to generate the triggered image $x_p = T(x, \alpha \cdot t)$. We offer the PSNR between the clean x and x_p , along with the predictions. We set $\Delta y_t = 40$.

poisoned data with the trigger, consistently predicted a specific target whenever the trigger is present during testing. These attacks are challenging to detect, as the backdoored models maintain high performance on clean test data. Current poisoning-based BA falls into two categories: 1) poison-label attacks, which contaminate training examples and alter their labels to the target value; and 2) clean-label attacks, which solely affect training examples while maintaining their labels.

In this paper, we introduce a novel poisoning-based backdoor attack against NR-IQA via a scalable trigger, addressing both poison-label and clean-label scenarios. The overall pipeline is illustrated in Fig. 1. Unlike traditional classifier BA, which operates in discrete outputs, our approach, designed for IQA models with continuous outputs, involves embedding a backdoor that manipulates the model’s output to any desired value using a scaling coefficient α for the trigger as shown in Fig. 1 (b). To counteract the diminishing effect of triggers in NR-IQA models due to cropping or data augmentation, we propose global-wise triggers on the patch-based pattern, leveraging the DCT space injections. Inspired by the effectiveness of universal adversarial perturbations (UAP) in misleading well-trained models, we utilize UAP in the DCT space (UAP-DCT) as triggers as shown in Fig. 1 (a), enhancing the susceptibility of IQA models to manipulations and improving attack effectiveness. Drawing from our proposed trigger injection and backdoor objectives, we devise an efficient heuristic technique for poison-label attacks (P-BAIQA). Furthermore, as shown in Fig. 1 (d), we explore the design of clean-label attacks (C-BAIQA), focusing on the sampling strategy of the scaling coefficient α and the refinement of image data x , driven by theoretical insights.

Our main contributions are summarized below:

- We introduce a novel poisoning-based backdoor attack against NR-IQA (BAIQA) via scalable triggers. As far as we

know, BAIQA is the first method to manipulate the predicted score to any desired value by varying the coefficient α .

- We propose to inject the trigger in the DCT domain to address trigger diminishment due to cropping or data augmentation. Furthermore, we utilize UAP in the DCT space as the trigger, to enhance attack effectiveness.

- In addition to the heuristic method for poison-label attacks (P-BAIQA), we explore the more challenging clean-label attacks (C-BAIQA) with theoretical insights, focusing on α sampling and image refinement. We demonstrate the effectiveness of our attacks on diverse datasets and models, and also show the resistance to several backdoor defenses.

Related Work

Image Quality Assessment. IQA tasks aim at predicting image quality scores that are consistent with human perception, often represented by Mean Opinion Score (MOS). These tasks can be categorized into Full-Reference (FR) and No-Reference (NR). In FR-IQA, the objective is to predict the quality score of the distorted image by comparing with its reference image. However, obtaining reference images can be challenging in real-world scenarios, leading to the emergence of NR-IQA, which predicts using only the distorted image. Some approaches (Mittal, Moorthy, and Bovik 2012; Ghadiyaram and Bovik 2017) leverage hand-crafted features, while others explore the influence of semantic information. For instance, Hyper-IQA (Su et al. 2020) employs a hypernetwork to obtain different quality estimators for images with varying content. DBCNN (Zhang et al. 2020) utilizes two independent neural networks to extract distorted and semantic information from images, which are then combined using bilinear pooling. Additionally, various studies have investigated the effectiveness of architectures in NR-IQA. TReS (Golestaneh, Dadsetan, and Kitani 2022) and MUSIQ (Ke et al. 2021) leverage

vision transformers (Dosovitskiy et al. 2021) and demonstrate their efficacy in NR-IQA tasks. While adversarial attacks against IQA tasks have received some attention (Zhang et al. 2022; Shumitskaya, Antsiferova, and Vatolin 2022; Liu et al. 2024a; Zhang et al. 2023), a more concerning and practical threat, known as backdoor attacks, remains unexplored.

Backdoor Attacks. BA (Chen et al. 2017) and AA (Szegedy et al. 2013) intend to modify the benign samples to mislead the DNNs, but they have some intrinsic differences. At the inference stage, AA (Madry et al. 2018; Ilyas et al. 2018) require much computational resources and time to generate the perturbation through iterative optimizations, and thus are not efficient in deployment. However, the perturbation (trigger) is known or easy to generate for BA. From the perspective of the attacker’s capacity, BA have access to poisoning training data, which adds an attacker-specified trigger (e.g. a local patch) and alters the corresponding label. BA on DNNs have been explored in BadNet (Gu, Dolan-Gavitt, and Garg 2017) for image classification by poisoning some training samples, and the essential characteristic consists of 1) backdoor stealthiness, 2) attack effectiveness on poisoned images, 3) low performance impact on clean images. Based on the capacity of attackers, BA can be categorized into poisoning-based and non-poisoning-based attacks (Li et al. 2020b). For poisoning-based attacks (Li et al. 2020a, 2021), attackers can only manipulate the dataset by inserting poisoned data, and have no access to the model training process. In contrast, non-poisoning-based attacks (Dumford and Scheirer 2020; Rakin, He, and Fan 2020; Guo, Wu, and Weinberger 2020; Doan et al. 2021) inject the backdoor by modifying the model parameters instead. For trigger generations, most attacks (Chen et al. 2017; Steinhardt, Koh, and Liang 2017) rely on fixed triggers, and several recent methods (Nguyen and Tran 2020; Liu et al. 2020) extend it to be sample-specific. For the trigger domain, several works (Zeng et al. 2021; Wang et al. 2022; Yue et al. 2022) consider the trigger in the frequency domain due to its advantages (Cheng et al. 2023). FTrojan (Wang et al. 2022) blockifies images and adds the trigger in the DCT domain, but it uses two fixed channels with fixed magnitudes.

Methodology

Problem Formulation

In the context of NR-IQA, consider a trainer to learn a model $f_{\theta} : \mathcal{X} \rightarrow \mathcal{Y}$. Here, $\mathcal{X} \subseteq \mathbb{R}^{H \times W \times 3}$ is the RGB space of inputs, $\mathcal{Y} \subseteq \mathbb{R}$ is the space of outputs, and θ is the trainable parameters. H and W are the height and width of inputs, respectively. The learning process of f_{θ} involves training with a dataset \mathcal{D}_t . \mathcal{D}_t can comprise both clean and poisoned data, denoted as $\mathcal{D}_t = \mathcal{D}_c \cup \mathcal{D}_p$. \mathcal{D}_c consists of correctly labeled clean data, while \mathcal{D}_p contains the poisoned data embedded with triggers by the attacker. The attacker has two primary objectives: **1) maintaining the model’s prediction accuracy on clean data**, thereby preventing detection through diminished performance on a validation set. **2) implanting a backdoor mechanism in the IQA model**, enabling the attacker to **manipulate the output by adding triggers** to the input.

Our Backdoor’s Goals. Unlike traditional BA on classifiers, where the output space is discrete, the output space of IQA is

continuous. Hence, our goal is to embed a backdoor capable of manipulating the model to output any values, rather than only one target class typically seen in BA on classifiers. This ensures that any input $\mathbf{x} \in \mathcal{X}$ with MOS value y , once altered with the trigger \mathbf{t} into the poisoned data $T(\mathbf{x}, \alpha \cdot \mathbf{t})$, is incorrectly predicted into a specific target value $y + \alpha \cdot \Delta y_t$. Here, $\alpha \in [-1, 1]$ controls the variations, and regulates the magnitude of \mathbf{t} . T is the trigger injector, and $\alpha \cdot \Delta y_t$ is the target variation. We can see that when using the Taylor expansion and considering the first-order derivative, the output of f_{θ} has a certain degree of linearity with respect to α , as follows:

$$f_{\theta}(T(\mathbf{x}, \alpha \cdot \mathbf{t})) = f_{\theta}(T(\mathbf{x}, 0 \cdot \mathbf{t})) + \alpha \frac{\partial f_{\theta}(T(\mathbf{x}, \alpha \cdot \mathbf{t}))}{\partial \alpha} \Big|_{\alpha=0}.$$

Thus, by choosing a proper α , attackers can manipulate the output to any target. The overall optimization is given by:

$$\begin{aligned} \forall \alpha \in [-1, 1] : \max_{T(\cdot, \alpha \cdot \mathbf{t})} & \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{C}} [|f_{\theta^*}(\mathbf{x}) - y|]}_{\text{low impact on clean data}} \\ & + \underbrace{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{C}} [|f_{\theta^*}(T(\mathbf{x}, \alpha \cdot \mathbf{t})) - (y + \alpha \cdot \Delta y_t)|]}_{\text{backdoor effectiveness}}, \quad (1) \\ \text{s.t. } \theta^* = \arg \min_{\theta} & \sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_c} \mathcal{L}(f_{\theta}(\mathbf{x}_i), y_i) \\ & + \sum_{(T(\mathbf{x}_i, \alpha \cdot \mathbf{t}), \hat{y}_i) \in \mathcal{D}_p} \mathcal{L}(f_{\theta}(T(\mathbf{x}_i, \alpha \cdot \mathbf{t})), \hat{y}_i), \end{aligned}$$

where \mathcal{C} is the distribution of clean data. The loss function \mathcal{L} , the model f_{θ} , and other hyperparameters are chosen exclusively by trainers. Specifically, when $\alpha = 0$, $T(\mathbf{x}, \alpha \cdot \mathbf{t}) = \mathbf{x}$.

To achieve these objectives, the attacker modifies the benign samples into poisoned ones and constructs a poisoned subset $\mathcal{D}_p = \{(T(\mathbf{x}_i, \alpha \cdot \mathbf{t}), \hat{y}_i)\}_{i=1}^{N_p}$. To enhance the stealthiness of the attacks, it is common practice to adopt a low poisoning ratio, represented by $r = |\mathcal{D}_p|/|\mathcal{D}_t|$, often below a certain threshold like 20%. Furthermore, if the original label y_i in \mathcal{D}_p are altered to the target value $y_i + \alpha \cdot \Delta y_t$, the strategy is named **poison-label attacks**. Otherwise, if the original label is maintained, the strategy is considered **clean-label attacks**. The assumptions for our attacks are: 1) The attacker has no access to the training process, including a lack of knowledge on the model’s architecture, and other configurations. 2) The attacker can inject poisoned data into the training set, achieved by converting clean data into poisoned data.

DCT Space for Robust Trigger Injection

Patch-based Patterns. Since NR-IQA models typically randomly crop multiple patches from input images and compute the average score based on these patches (Su et al. 2020; Ke et al. 2021), triggers based on a local patch or predefined global patterns may not persist after the cropping or data augmentations during training. To address the diminishing effect of triggers during training/attacking, we propose adding the trigger with a patch-based pattern across the global area. By introducing the trigger in the **DCT** space, which is inherently block-based, a patch-based pattern is effectively created across the entire image, making it well-suited for our attacks.

Algorithm 1: Searching UAP in the DCT domain (UAP-DCT)

Require: Subset $\mathcal{D}_s = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_s}$, surrogate model f_{θ_s} , perturbation bound ϵ , epochs e_1, e_2
Ensure: Optimized trigger \mathbf{t}

- 1: Train f_{θ_s} on \mathcal{D}_s using loss \mathcal{L}_1 and Adam for e_1 epochs
- 2: **for** $i = 1$ **to** e_2 **do**
- 3: **for each** $(\mathbf{x}, y) \in \mathcal{D}_s$ **do**
- 4: $\mathcal{L}_{UAP} = - \underbrace{\|f_{\theta_s}(T(\mathbf{x}, \mathbf{t})) - f_{\theta_s}(\mathbf{x})\|_1}_{\text{effectiveness of UAP-DCT}} + \lambda \cdot \underbrace{\max(\mathcal{L}_{mse}(\mathbf{x}, T(\mathbf{x}, \mathbf{t})), \epsilon^2)}_{\text{invisibility of UAP-DCT}}$
- 5: Update \mathbf{t} by minimizing \mathcal{L}_{UAP} using the Adam
- 6: **end for**
- 7: **end for**

Trigger Injection. The trigger injection model $T(\mathbf{x}, \alpha \cdot \mathbf{t})$ takes a image \mathbf{x} , the trigger \mathbf{t} , and the scaling coefficient α to generate a poisoned image \mathbf{x}_p . Given the prevalence of DCT in coding techniques (Wallace 1992) and its ability to perform patch-based transforms on images of varying sizes, we propose a frequency-based trigger injection. Given \mathbf{x} , we split it into non-overlapping patches \mathbf{x}_{patch} . Following a 2d-DCT transform, we have the \mathbf{x}_{dct} . By adding the scaled trigger $\alpha \cdot \mathbf{t}$ to all patches of \mathbf{x}_{dct} , we have the triggered \mathbf{x}_{dct}^p . The final result $T(\mathbf{x}, \alpha \cdot \mathbf{t})$ is obtained by applying an inverse 2d-DCT. The overall process is given by: $\forall \alpha \in [-1, 1]: \mathbf{x}_p = T(\mathbf{x}, \alpha \cdot \mathbf{t}) = \text{IDCT}(\text{DCT}(\mathbf{x}) + \alpha \cdot \mathbf{t})$, where DCT and inverse-DCT (IDCT) are on all patches. To maintain imperceptibility, we add in middle frequencies, as adding in low frequencies induces significant perturbations in the spatial domain, while high frequencies are less robust. We use a patch size of 16×16 , and perturb the mid-64 frequencies.

Exploiting DCT Space Vulnerability. Indeed, research has revealed that state-of-the-art DNN classifiers are vulnerable to a phenomenon known as Universal Adversarial Perturbations (UAP), which refers to a universal and invisible perturbations that can cause high-probability misclassifications. Inspired by the efficacy of universally applicable perturbations in misleading well-trained models, we leverage UAP in the DCT space, denoted as UAP-DCT, as the trigger \mathbf{t} . By incorporating these UAP-DCT triggers into a poisoned dataset, the attacker can amplify the susceptibility of a trained IQA model to such triggers, thereby enhancing attack performance and facilitating more effective manipulation of the attacked outputs. Algorithm 1 outlines the optimization process for the trigger \mathbf{t} , involving two loss terms. The first loss aims to generate effective UAP-DCT, while the second loss regulates the magnitude of UAP-DCT to ensure the induced perturbations remain invisible in the spatial domain, thereby enhancing the imperceptibility of the poisoned images.

Backdoor Attacks with Poison Label

The poison-label backdoor attacks against NR-IQA (**P-BAIQA**) are introduced to concretely validate the potency of our designed backdoor triggers and attacking goals. Specifically, similar to the existing poison-label backdoor attacks,

our P-BAIQA first randomly select a subset \mathcal{D}_s with poisoning ratio r from the benign dataset \mathcal{D} to make its modified version $\mathcal{D}_p = \{(T(\mathbf{x}, \alpha \cdot \mathbf{t}), y + \alpha \cdot \Delta y_t) | \alpha \sim \mathcal{P}_\alpha, (\mathbf{x}, y) \in \mathcal{D}_s\}$, where \mathcal{P}_α denotes sampling distribution of α . Prior to generating the poisoned dataset, we optimize the trigger \mathbf{t} based on \mathcal{D}_s using Algorithm 1. Given the sampling distribution \mathcal{P}_α , we also explore several cases: 1) sampling from $\{\pm 1, \pm \frac{3}{4}, \pm \frac{1}{2}, \pm \frac{1}{4}\}$ with probabilities $\{0.4, 0.3, 0.2, 0.1\}$; 2) sampling from $\{\pm 1\}$ with equal probability; 3) Uniform distribution (i.e., $U(-1, 1)$). The modified subset \mathcal{D}_p with the remaining clean samples $\mathcal{D}_c = \mathcal{D} \setminus \mathcal{D}_s$ will then be released to train the model $f_{\theta}(\cdot)$ as indicated in Eq. 1.

During the inference, for any test sample (\mathbf{x}, y) , the adversary can trigger the backdoor using the poisoned $T(\mathbf{x}, \alpha \cdot \mathbf{t})$. The adversary is free to choose the α , thus manipulating the output to the desired target $f_{\theta^*}(\mathbf{x}) + \alpha \cdot \Delta y_t$ as intended.

Backdoor Attacks with Clean Label

As we will demonstrate in Section 4, the heuristic P-BAIQA can achieve promising results. However, it **lacks stealthiness**, as it retains poisoned labels. Dataset users may detect the poisoned data through label inspection. In this section, we explore the design of clean-label backdoor attacks (**C-BAIQA**), focusing on the sampling strategy of α and the modification of \mathbf{x} . We first present the necessary assumption and theorem.

Assumption 1 For a backdoored model f_{θ} , we assume that for any input and the true label $(\mathbf{x}, y) \sim \mathcal{C}$, the output \tilde{y} follows a Gaussian distribution with an variance σ^2 :

$$\begin{aligned} \mathbb{P}(\tilde{y}|\mathbf{x}, \alpha) &= \mathcal{N}(\tilde{y}; y + \alpha \Delta y_t, \sigma^2), \quad \mathbb{P}(\tilde{y}|\mathbf{x}) = \mathcal{N}(\tilde{y}; y, \sigma^2), \\ \mathbb{P}(\tilde{y}) &= \mathcal{N}(\tilde{y}; \mu_{\tilde{y}}, \sigma^2), \quad \mathbb{P}(\tilde{y}|\alpha) = \mathcal{N}(\tilde{y}; \mu_{\tilde{y}|\alpha}, \sigma^2), \end{aligned} \quad (2)$$

where \mathcal{C} is the clean test data, and $\mu_{\tilde{y}}, \mu_{\tilde{y}|\alpha}$ are given by

$$\begin{aligned} \mu_{\tilde{y}|\alpha} &= \mathbb{E}[\tilde{y}|\alpha] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}[\tilde{y}|\mathbf{x}, \alpha]] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}[\tilde{y}|\mathbf{x}]] + \alpha \cdot \Delta y_t \\ &= \mathbb{E}[\tilde{y}] + \alpha \cdot \Delta y_t = \mu_{\tilde{y}} + \alpha \cdot \Delta y_t. \end{aligned} \quad (3)$$

For clean-label attacks, where labels in \mathcal{D}_p remain unchanged, we assume that f_{θ} is capable of learning the label priors during the training process. Thus, we have:

$$\mu_{\tilde{y}} = \mathbb{E}_{y \sim \mathcal{D}_t}[y] = \mu_y, \quad \mu_{\tilde{y}|\alpha} = \mu_{\tilde{y}} + \alpha \Delta y_t = \mu_y + \alpha \Delta y_t. \quad (4)$$

Assumption 1 aligns with the fundamental characteristics of backdoor attacks at test time. Specifically, $\mathbb{P}(\tilde{y}|\mathbf{x}, \alpha)$ is the output distribution when triggered data is input, while $\mathbb{P}(\tilde{y}|\mathbf{x})$ is the output distribution for clean data. Additionally, $\mathbb{P}(\tilde{y})$ serves as the prior for this task, and $\mathbb{P}(\tilde{y}|\mathbf{x}, \alpha)$ reflects the conditional distribution when considering only the trigger. The means of these distributions are consistent with backdoor attack behavior. Furthermore, it is common to assume that the output variations are identical, denoted as σ^2 .

For Eq. 4, in a simplified context, $\mu_{\tilde{y}}$ serves as the prior of the trained model and is generally understood as the expected value of the labels across the entire training dataset. Conversely, $\mu_{\tilde{y}|\alpha}$ is the prior of the trained model in the presence of the trigger, typically corresponding to the expected value of the labels within the poisoned subset of the training dataset. Therefore, based on this, we can develop the sampling strategy for α as outlined in the following remark.

Algorithm 2: C-BAIQA

Require: Benign dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, surrogate model f_{θ_s} , poisoning ratio r , hyperparameters for UAP-DCT (ϵ, e_1, e_2), hyperparameters for TAEs (target μ_y , bound ϵ_t , step size α_t , iterations I_t)

Ensure: Optimized trigger \mathbf{t} , Modified dataset \mathcal{D}_t

- 1: $\mathcal{D}_s \leftarrow$ Randomly sample a subset from \mathcal{D} with ratio r
- 2: # Train surrogate model and search trigger
- 3: $f_{\theta_s}, \mathbf{t} \leftarrow$ Apply Algorithm 1 with $(\mathcal{D}_s, f_{\theta_s}, \epsilon, e_1, e_2)$
- 4: # Generate poisoned subset \mathcal{D}_p
- 5: **for** each $(\mathbf{x}_i, y_i) \in \mathcal{D}_s$ **do**
- 6: $\alpha_i \leftarrow (y_i - \mu_y) / \Delta y_t$
- 7: $\mathbf{x}'_i \leftarrow$ **Targeted-PGD**(model= f_{θ_s} , input= \mathbf{x}_i , target= μ_y , bound= ϵ_t , step= α_t , iter= I_t)
- 8: **end for**
- 9: $\mathcal{D}_p \leftarrow \{(T(\mathbf{x}'_i, g(\alpha_i) \cdot \mathbf{t}), y_i + \alpha_i \cdot \Delta y_t)\}_{i=1}^{N_s}$
- 10: $\mathcal{D}_c \leftarrow \mathcal{D} \setminus \mathcal{D}_s, \mathcal{D}_t \leftarrow \mathcal{D}_c \cup \mathcal{D}_p$
- 11: **return** $\mathbf{t}, \mathcal{D}_t$

Remark 1 (Sampling strategy of α in \mathcal{D}_p) To train a model f_{θ} satisfying the Eq. 4, the poisoned training set \mathcal{D}_p should adhere to those properties as well. Thus, considering the distribution of the label y , one approach to sample α satisfying Eq. 4 is to select $\alpha = (y - \mu_y) / \Delta y_t$ for \mathcal{D}_p .

Remark 1 outlines the sampling strategy for α based on Assumption 1. Additionally, for Eq. 2, it presents the mathematical relationship $\mathbb{P}(\tilde{y}|\mathbf{x}, \alpha)\mathbb{P}(\tilde{y}) = \mathbb{P}(\tilde{y}|\mathbf{x})\mathbb{P}(\tilde{y}|\alpha)$ for any test data. Since f_{θ} learns this property from the training set, it is crucial that the training set adheres to this equation. To increase the likelihood of this equation and ensure consistency between training on \mathcal{D}_p and evaluation during the attack phase, we consider whether modifying \mathbf{x} to \mathbf{x}' is necessary. This leads to the development of the following theorem.

Theorem 1 For any data point $(\mathbf{x}', y) \in \mathcal{D}_p$, consistency between training and testing requires that it satisfies the distributions in Eq. 2. This condition can be expressed as

$$\mathbb{P}(\tilde{y}|\mathbf{x}', \alpha)\mathbb{P}(\tilde{y}) = \mathbb{P}(\tilde{y}|\mathbf{x}')\mathbb{P}(\tilde{y}|\alpha), \quad (5)$$

which is identical to

$$\mathbb{P}(\tilde{y}|\mathbf{x}', \alpha) = \frac{\mathbb{P}(\tilde{y}|\alpha)}{\mathbb{P}(\tilde{y})}\mathbb{P}(\tilde{y}|\mathbf{x}') = \frac{\mathbb{P}(\alpha|\tilde{y})}{\mathbb{P}(\alpha)}\mathbb{P}(\tilde{y}|\mathbf{x}'). \quad (6)$$

Considering Bayes' theorem, we also derive

$$\mathbb{P}(\tilde{y}|\mathbf{x}', \alpha) = \frac{\mathbb{P}(\alpha|\mathbf{x}', \tilde{y})\mathbb{P}(\mathbf{x}'|\tilde{y})\mathbb{P}(\tilde{y})}{\mathbb{P}(\alpha|\mathbf{x}')\mathbb{P}(\mathbf{x}')} = \frac{\mathbb{P}(\alpha|\mathbf{x}', \tilde{y})}{\mathbb{P}(\alpha|\mathbf{x}')} \mathbb{P}(\tilde{y}|\mathbf{x}'). \quad (7)$$

By comparing the two expressions for $\mathbb{P}(\tilde{y}|\mathbf{x}', \alpha)$, we have

$$\frac{\mathbb{P}(\alpha|\tilde{y})}{\mathbb{P}(\alpha)} = \frac{\mathbb{P}(\alpha|\mathbf{x}', \tilde{y})}{\mathbb{P}(\alpha|\mathbf{x}')} \stackrel{\text{Bayes}}{\iff} \frac{\mathbb{P}(\alpha|\tilde{y})}{\mathbb{P}(\alpha)} = \frac{\mathbb{P}(\alpha|\tilde{y})\mathbb{P}(\mathbf{x}'|\alpha, \tilde{y})}{\mathbb{P}(\alpha|\mathbf{x}')\mathbb{P}(\mathbf{x}'|\tilde{y})}. \quad (8)$$

Theorem 1 highlights the necessity of modifying \mathbf{x} to \mathbf{x}' within \mathcal{D}_p . Consequently, the generation of \mathcal{D}_p can be divided into two steps. First, after selecting the subset \mathcal{D}_s , \mathbf{x} is transformed into \mathbf{x}' . Second, the sampling strategy outlined in Remark 1 is applied to convert \mathbf{x}' into $T(\mathbf{x}', \alpha \cdot \mathbf{t})$.

Remark 2 (Strategy of converting \mathbf{x} into \mathbf{x}') To convert \mathbf{x} into \mathbf{x}' satisfying Eq. 8, one approach is to ensure \mathbf{x}' is independent of both α and \tilde{y} . Since sampling of α is based on y , making \mathbf{x}' independent of \tilde{y} is sufficient. In this case, we have $\mathbb{P}(\tilde{y}|\mathbf{x}') = \mathbb{P}(\tilde{y}) = \mathcal{N}(\tilde{y}; \mu_{\tilde{y}}, \sigma^2)$, following the distribution in Eq. 2. Thus, by setting targeted adversarial examples (TAEs) of \mathbf{x} (with μ_y as the target score) as the modified \mathbf{x}' , the above equation is satisfied. Specifically, since trigger injection occurs in the DCT space, we choose the spatial domain as the perturbation space for generating \mathbf{x}' .

The overall framework for C-BAIQA is in Algorithm 2. After randomly selecting a subset \mathcal{D}_s from \mathcal{D} , we first use Algorithm 1 to train f_{θ_s} and search \mathbf{t} . Following Remark 1 and Remark 2, we then transform \mathcal{D}_s into \mathcal{D}_p , and mix it with the remaining \mathcal{D}_c to construct the final training set \mathcal{D}_t .

Experiments

Experimental Setup

Datasets and Models. We choose the LIVEC dataset (Ghadiyaram and Bovik 2015) consisting of 1162 images with sizes 500×500 and KonIQ-10k dataset (Hosu et al. 2020) consisting of 10073 images with sizes 512×384 . For LIVEC, consistent with prior research (Liu et al. 2024a), we randomly select 80% of the images as the training set, and the remaining 20% as the testing set. For KonIQ-10k, we adopt the official split with 70% of the images as the training set. We include two CNN-based NR-IQA models: HyperIQA (Su et al. 2020), and DBCNN (Zhang et al. 2020), and one Transformer-based TRoS (Golestaneh, Dadsetan, and Kitani 2022).

Baseline Attack Methods. We compare our attacks with existing poisoning-based BA. We utilize Blended (Chen et al. 2017), WaNet (Nguyen and Tran 2021), and FTrojan (Wang et al. 2022), which are representative non-patch-based invisible attacks. For all baselines, we employ the coefficient α to adjust the intensity and strength of the triggers, following a similar approach to our attacks. For FTrojan, we add the triggers into the same mid-frequency channels. We apply the same configurations to all methods, except the trigger.

Evaluation Metrics. For the evaluation on benign data, we follow previous works and consider three metrics, *i.e.*, RMSE, PLCC, and SROCC. To evaluate the attack effectiveness, we include the mean absolute error (MAE), and the mean ratio of amplification (MRA) across the test set regarding each $\alpha \in [-1, 1]$. While it is impossible to consider all values of α , we estimate it using finite α from a selected set A . Therefore, we utilize the mean of those metrics for an overall evaluation:

$$\text{MAE}(\alpha) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{C}} [|\Delta_a(\mathbf{x}, \alpha) - \Delta_t(\alpha)|],$$

$$\text{MRA}(\alpha) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{C}} \left[\frac{\Delta_a(\mathbf{x}, \alpha)}{\Delta_t(\alpha)} \right],$$

$$\text{mMAE} = \frac{1}{n(A)} \sum_{\alpha \in A} \text{MAE}(\alpha), \text{mMRA} = \frac{1}{n(A)} \sum_{\alpha \in A} \text{MRA}(\alpha),$$

$$\text{with } \Delta_a(\mathbf{x}, \alpha) = f_{\theta^*}(T(\mathbf{x}, \alpha \cdot \mathbf{t})) - f_{\theta^*}(\mathbf{x}), \Delta_t(\alpha) = \alpha \cdot \Delta y_t.$$

Experimentally, we choose $A = \{\pm 0.1, \pm 0.2, \dots, \pm 1.0\}$. For the imperceptibility of poisoned images, we provide the PSNR between \mathbf{x}_p and \mathbf{x} , denoted as PSNR_1 , when we set $\alpha = 1$.

Attack Settings		Poison-label attacks (P-BAIQA)								Clean-label attacks (C-BAIQA)											
Dataset		LIVEC				KonIQ-10k				LIVEC				KonIQ-10k							
Model Attacks		Benign			Attack		Benign			Attack		Benign			Attack		Benign			Attack	
		(a)	(b)	(c)	(A)	(B)	(a)	(b)	(c)	(A)	(B)	(a)	(b)	(c)	(A)	(B)	(a)	(b)	(c)	(A)	(B)
HyperIQA	w/o	0.911	0.897	9.47	-	-	0.905	0.886	7.00	-	-	0.911	0.897	9.47	-	-	0.905	0.886	7.00	-	-
	Blended	0.887	0.859	10.27	18.88	0.11	0.894	0.876	7.31	18.09	0.14	0.903	0.883	9.91	22.47	-0.02	0.899	0.885	7.26	22.25	-0.00
	WaNet	0.877	0.855	10.59	22.03	-0.00	0.893	0.878	7.09	22.01	-0.00	0.877	0.855	10.59	22.03	-0.00	0.909	0.895	6.66	22.00	0.00
	FTrojan	0.912	0.878	9.64	10.18	0.51	0.897	0.879	7.10	7.77	0.63	0.905	0.883	9.40	18.92	0.13	0.905	0.894	7.00	16.31	0.30
Ours	0.903	0.872	9.52	8.72	0.63	0.903	0.887	6.93	6.42	0.65	0.903	0.892	9.91	10.65	0.60	0.900	0.885	7.00	15.19	0.29	
DBCNN	w/o	0.878	0.869	10.48	-	-	0.905	0.889	6.74	-	-	0.878	0.869	10.48	-	-	0.905	0.889	6.74	-	-
	Blended	0.855	0.817	10.96	18.44	0.13	0.887	0.865	7.34	14.00	0.29	0.773	0.717	17.01	22.18	-0.01	0.888	0.881	7.47	22.27	-0.01
	WaNet	0.828	0.798	11.71	22.04	-0.00	0.867	0.839	7.78	21.96	0.00	0.828	0.798	11.71	22.04	-0.00	0.891	0.885	7.38	22.00	-0.00
	FTrojan	0.866	0.848	10.90	10.50	0.50	0.898	0.878	7.11	5.60	0.67	0.771	0.723	17.15	21.89	0.00	0.890	0.882	7.44	21.12	0.04
Ours	0.887	0.860	9.98	9.45	0.56	0.902	0.883	7.03	5.21	0.69	0.890	0.866	10.09	12.37	0.45	0.903	0.887	6.95	13.97	0.37	
TReS	w/o	0.909	0.894	15.22	-	-	0.909	0.886	19.20	-	-	0.909	0.894	15.22	-	-	0.909	0.886	19.20	-	-
	Blended	0.878	0.845	17.60	20.15	0.10	0.871	0.860	15.08	19.49	0.10	0.897	0.872	17.61	23.08	-0.04	0.897	0.880	18.22	22.86	-0.03
	WaNet	0.881	0.869	16.80	22.00	0.00	0.839	0.836	24.14	22.08	-0.00	0.881	0.869	16.80	22.00	0.00	0.904	0.887	18.94	21.99	0.00
	FTrojan	0.866	0.843	18.11	10.33	0.66	0.891	0.877	20.76	9.81	0.81	0.895	0.882	17.94	19.60	0.10	0.886	0.872	17.24	16.96	0.22
Ours	0.877	0.855	17.31	10.22	0.73	0.892	0.875	22.24	9.01	0.93	0.895	0.871	16.15	13.75	0.42	0.887	0.862	16.99	14.60	0.35	

Table 1: Comparison of P-BAIQA and C-BAIQA with baseline attack methods with the poisoning ratio $r = 20\%$. (a), (b) and (c) denote the PLCC, SROCC and RMSE, respectively. (A) and (B) denote the mMAE \downarrow and mMRA \uparrow , respectively.

Attacks \rightarrow	Blended	WaNet	FTrojan	Ours
LIVEC	27.71	26.59	29.80	30.06
KonIQ-10k	33.00	34.81	35.91	36.32

Table 2: Imperceptibility of the poisoned data (PSNR₁ \uparrow).

Performance of P-BAIQA

Settings. As the MOS range is $[0, 100]$, we set the maximum deviation $\Delta y_t = 40$, ensuring this deviation is sufficient to yield potent attacks. To regulate the imperceptibility of poisoned images, in Algorithm 1, we set $\lambda = 10^8$, $e_1 = 24$, $e_2 = 50$, and $\epsilon = \frac{8}{255}$ for LIVEC. For KonIQ-10k, where a larger number of images are eligible for poisoning, we adjust ϵ to a lower value $\frac{4}{255}$. We set the poisoning ratio $r = 20\%$ for all attacks and datasets, and use the HyperIQA as the surrogate model. For the sampling strategy of α to construct \mathcal{D}_p , we consider sampling from $\{\pm 1, \pm \frac{3}{4}, \pm \frac{1}{2}, \pm \frac{1}{4}\}$ with probabilities $\{0.4, 0.3, 0.2, 0.1\}$, and provide results of other strategies in the ablation study.

Results. As depicted in Table 1 and Table 2, our approach consistently surpasses the baseline attacks on both datasets, demonstrating both a more potent attack effectiveness and an enhanced level of invisibility for the trigger. Blended and WaNet encounter significant difficulties when attempting to attack NR-IQA models. This is likely due to two key factors: 1) Blended adds a predefined trigger across the entire image area. However, during the training and testing of NR-IQA models, a random cropping operation is typically applied, which disrupts the direct correlation between the global spatial trigger and the variations in the model’s output. 2) WaNet employs a wrapping function to inject the trigger, but this method may not be effectively recognized or learned by NR-IQA models, as their output space is continuous and the target score to be manipulated can be any value. In contrast to FTrojan, which relies on manually defined triggers in the DCT space, our approach leverages UAP, capitalizing on the

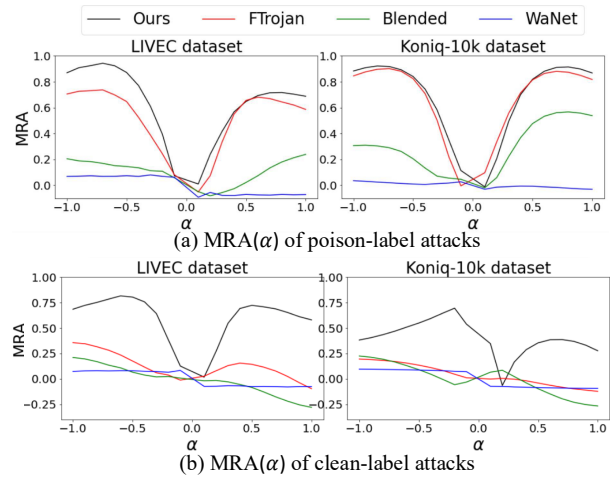


Figure 2: MRA(α) with HyperIQA as victim models.

vulnerabilities in the DCT space to amplify the effectiveness of the attack. Moreover, the benign metrics also show that our attack have low performance impact on the clean data.

In addition, we provide the MRA(α) in Fig. 2(a). We can see that our method is capable of achieving a significant deviation when α has an absolute value greater than 0.5. However, when α falls within $[-0.5, 0.5]$, the manipulation becomes less precise, as the model finds it more challenging to recognize the trigger. However, we believe this limitation is tolerable given that attackers generally aim for significant shifts in outputs rather than minor alterations.

Performance of C-BAIQA

Settings. The configurations for λ , e_1 , e_2 , ϵ , Δy_t , r , and the surrogate model in Algorithm 1 match those in P-BAIQA. As the MOS range is $[0, 100]$, we set $\mu_y = 50$ in Algorithm 2. For the TAEs, we adopt PGD attacks (Madry et al. 2018),

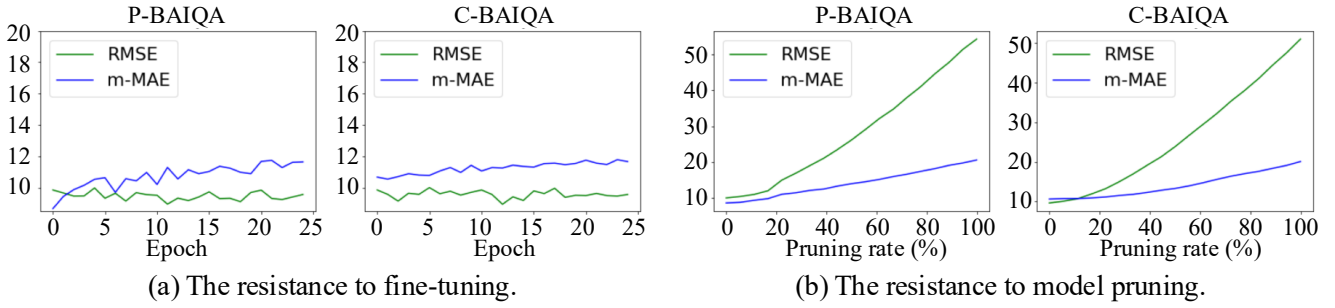


Figure 3: Resistance to fine-tuning and pruning (HyperIQA as models and LIVEC as the dataset).

Attack Type	Dataset → Model → Metrics →	LIVEC						KonIQ-10k					
		HyperIQA		DBCNN		TReS		HyperIQA		DBCNN		TReS	
		mMAE	mMRA	mMAE	mMRA	mMAE	mMRA	mMAE	mMRA	mMAE	mMRA	mMAE	mMRA
P-BAIQA	①	7.825	0.7660	9.093	0.7900	11.241	0.8244	7.548	0.8540	6.741	0.8078	13.193	1.0916
	②	9.485	0.5419	9.608	0.5756	11.831	0.6942	7.029	0.6413	5.510	0.6813	9.543	0.8308
	Ours	8.719	0.6344	9.449	0.5611	10.223	0.7274	6.420	0.6487	5.208	0.6917	9.012	0.9291
C-BAIQA	③	12.689	0.4117	12.610	0.4195	14.597	0.3490	15.166	0.3002	14.277	0.3463	17.837	0.1714
	④	20.777	0.0638	19.320	0.1257	19.089	0.1271	19.023	0.1924	16.388	0.3242	21.499	0.0249
	Ours	10.653	0.5983	12.372	0.4525	13.746	0.4183	15.185	0.2859	13.970	0.3682	14.597	0.3490

Table 3: Ablation study of our attacks. ① and ② denote our P-BAIQA with α sampling from $\{\pm 1\}$ and $U(-1, 1)$, respectively. ③ denotes our C-BAIQA without converting x into x' . ④ denotes our C-BAIQA by setting $\mu_y=65$ for the sampling strategy of α .

with ϵ_t in terms of ℓ_∞ -norm (i.e., $\epsilon_t = 2/255$ for LIVEC and $\epsilon_t = 1/255$ for KonIQ-10k), $\alpha_t = 1/255$, and $I_t = 20$.

Results. From Table 1 and Table 2, none of the baseline methods successfully attack NR-IQA models. However, our C-BAIQA even achieves a comparable performance to our P-BAIQA on the LIVEC. On the KonIQ-10k, we observed a significant drop in the performance compared to P-BAIQA. This may be attributed to the fact that KonIQ-10k contains a larger number of clean data samples compared to LIVEC. Consequently, the NR-IQA models may have been less inclined to learn spurious correlations from the poisoned data and instead relied more on learning from the clean data. From Fig. 2 (b), we see a similar trend as the results of P-BAIQA.

Resistance to Backdoor Defenses

We assess the resistance of our attacks to backdoor defenses. Given that strategies like Neural Cleanse (Wang et al. 2019), and STRIP (Gao et al. 2021) are tailored for defending against BA in classification, we delve into the resistance of our BAIQA against fine-tuning (Liu, Xie, and Srivastava 2017; Liu, Dolan-Gavitt, and Garg 2018) and model pruning (Liu, Dolan-Gavitt, and Garg 2018; Wu and Wang 2021). These are representative defenses that are applicable to our tasks. Further details on the setup are in the appendix.

As depicted in Fig. 3(a), our attacks are resist to fine-tuning. Notably, the mMAE for both P-BAIQA and C-BAIQA remain largely unaffected, experiencing only a marginal increase of approximately 2 after the whole fine-tuning. Furthermore, our attacks are resistant to model pruning, as evident from Fig. 3(b). Specifically, when applying a pruning rate of 40%, the mMAE values for P-BAIQA and C-BAIQA remain relatively stable, whereas the benign metric (RMSE) is significantly impacted, resulting in elevated values.

Ablation Study

Sampling of α for P-BAIQA. We explore three cases: 1) $\{\pm 1, \pm \frac{3}{4}, \pm \frac{1}{2}, \pm \frac{1}{4}\}$ with probabilities $\{0.4, 0.3, 0.2, 0.1\}$ as ours; 2) $\{\pm 1\}$ with equal probability; 3) Uniform distribution (i.e., $U(-1, 1)$). In Table 3, when the quantity of poisoned data is limited, e.g., on LIVEC, option 2) is a preferable choice. This preference stems from the fact that a higher absolute value of α results in a greater loss, thereby taking precedence during training, and making the model more effectively learn the correlation between the trigger and the attack target. When the amount of poisoned data is adequate, e.g., on KonIQ-10k, option 1), as adopted in P-BAIQA, provides more precise control regarding the mMAE during attacking.

Value μ_y and strategy of converting x into x' for C-BAIQA. From Table 3, the results of ③ show that the transformation of x into x' enhances the attack’s efficacy, consistent with Remark 2. Results of ③ show that values of μ_y that is closer to the mean MOS across the dataset yields improved attack effectiveness, which is in line with Remark 1.

Conclusion

In this work, we introduce a novel poisoning-based backdoor attack against NR-IQA, that leverages the adjustment of a scaling coefficient α for the trigger to manipulate the model’s output to desired target values. We propose embedding the trigger in the DCT domain, and incorporate UAP in the DCT space as the trigger, enhancing the IQA model’s susceptibility to manipulation and improving the attack’s efficacy. Our approach explores both poison-label and clean-label scenarios, the latter focusing on α sampling and image data refinement with theoretical insights. Extensive experiments on diverse datasets and NR-IQA models validate the effectiveness.

Acknowledgments

This work was done at Rapid-Rich Object Search Lab, School of Electrical and Electronic Engineering, Nanyang Technological University. This research is supported in part by the Basic and Frontier Research Project of PCL, the Major Key Project of PCL, and Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515010454.

References

- Chen, X.; Liu, C.; Li, B.; Lu, K.; and Song, D. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*.
- Cheng, H.; Yang, S.; Zhou, J. T.; Guo, L.; and Wen, B. 2023. Frequency guidance matters in few-shot learning. In *Proc. IEEE Int'l Conf. Computer Vision*, 11814–11824.
- Ding, K.; Ma, K.; Wang, S.; and Simoncelli, E. P. 2020. Image quality assessment: Unifying structure and texture similarity. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 44(5): 2567–2581.
- Doan, K.; Lao, Y.; Zhao, W.; and Li, P. 2021. LIRA: Learnable, Imperceptible and Robust Backdoor Attacks. In *Proc. IEEE Int'l Conf. Computer Vision*, 11966–11976.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proc. Int'l Conf. Learning Representations*.
- Dumford, J.; and Scheirer, W. 2020. Backdooring convolutional neural networks via targeted weight perturbations. In *Proc. IEEE Int'l Joint Conf. on Biometrics*, 1–9.
- Fang, S.; and Choromanska, A. 2022. Backdoor attacks on the DNN interpretation system. In *Proc. AAAI Conf. on Artificial Intelligence*, volume 36, 561–570.
- Fu, H.; Liang, F.; Liang, J.; Li, B.; Zhang, G.; and Han, J. 2023. Asymmetric learned image compression with multi-scale residual block, importance scaling, and post-quantization filtering. *IEEE Trans. on Circuits and Systems for Video Technology*.
- Gao, Y.; Kim, Y.; Doan, B. G.; Zhang, Z.; Zhang, G.; Nepal, S.; Ranasinghe, D. C.; and Kim, H. 2021. Design and evaluation of a multi-domain trojan detection method on deep neural networks. *IEEE Trans. on Dependable and Secure Computing*, 19(4): 2349–2364.
- Ghadiyaram, D.; and Bovik, A. C. 2015. Massive online crowdsourced study of subjective and objective picture quality. *IEEE Trans. on Image Processing*, 25(1): 372–387.
- Ghadiyaram, D.; and Bovik, A. C. 2017. Perceptual quality prediction on authentically distorted images using a bag of features approach. *Journal of vision*, 17(1): 32–32.
- Golestaneh, S. A.; Dadsetan, S.; and Kitani, K. M. 2022. No-Reference Image Quality Assessment via Transformers, Relative Ranking, and Self-Consistency. In *Proc. of the IEEE/CVF Winter Conf. on Applications of Computer Vision*.
- Gu, T.; Dolan-Gavitt, B.; and Garg, S. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.
- Guo, C.; Wu, R.; and Weinberger, K. Q. 2020. Trojannet: Embedding hidden trojan horse models in neural networks. *arXiv preprint arXiv:2002.10078*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 770–778.
- Hosu, V.; Lin, H.; Sziranyi, T.; and Saupe, D. 2020. KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Trans. on Image Processing*.
- Ilyas, A.; Engstrom, L.; Athalye, A.; and Lin, J. 2018. Black-box adversarial attacks with limited queries and information. In *Proc. Int'l Conf. Machine Learning*, 2137–2146.
- Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *Proc. IEEE Int'l Conf. Computer Vision*, 5148–5157.
- Korhonen, J.; and You, J. 2022. Adversarial attacks against blind image quality assessment models. In *Proc. of the 2nd Workshop on Quality of Experience in Visual Multimedia Applications*.
- Li, S.; Xue, M.; Zhao, B. Z. H.; Zhu, H.; and Zhang, X. 2020a. Invisible backdoor attacks on deep neural networks via steganography and regularization. *IEEE Trans. on Dependable and Secure Computing*, 18(5): 2088–2105.
- Li, Y.; Li, Y.; Wu, B.; Li, L.; He, R.; and Lyu, S. 2021. Invisible backdoor attack with sample-specific triggers. In *Proc. IEEE Int'l Conf. Computer Vision*, 16463–16472.
- Li, Y.; Wu, B.; Jiang, Y.; Li, Z.; and Xia, S.-T. 2020b. Backdoor learning: A survey. *arXiv preprint arXiv:2007.08745*.
- Liang, K.; Zhang, J. Y.; Wang, B.; Yang, Z.; Koyejo, S.; and Li, B. 2021. Uncovering the connections between adversarial transferability and knowledge transferability. In *Proc. Int'l Conf. Machine Learning*.
- Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2018. Fine-Pruning: Defending Against Backdooring Attacks on Deep Neural Networks. *arXiv preprint arXiv:1805.12185*.
- Liu, Y.; Ma, X.; Bailey, J.; and Lu, F. 2020. Reflection backdoor: A natural backdoor attack on deep neural networks. In *European Conf. on Computer Vision*, 182–199.
- Liu, Y.; Xie, Y.; and Srivastava, A. 2017. Neural trojans. In *IEEE Int'l Conf. on Computer Design*, 45–48.
- Liu, Y.; Yang, C.; Li, D.; Ding, J.; and Jiang, T. 2024a. Defense Against Adversarial Attacks on No-Reference Image Quality Models with Gradient Norm Regularization. *arXiv preprint arXiv:2403.11397*.
- Liu, Z.; Wang, T.; Huai, M.; and Miao, C. 2024b. Backdoor attacks via machine unlearning. In *Proc. AAAI Conf. on Artificial Intelligence*, volume 38, 14115–14123.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proc. Int'l Conf. Learning Representations*.
- Mittal, A.; Moorthy, A. K.; and Bovik, A. C. 2012. No-reference image quality assessment in the spatial domain. *IEEE Trans. on Image Processing*, 21(12): 4695–4708.
- Nguyen, T. A.; and Tran, A. 2020. Input-aware dynamic backdoor attack. In *Proc. Annual Conf. Neural Information Processing Systems*, volume 33, 3454–3464.

- Nguyen, T. A.; and Tran, A. T. 2021. WaNet - Imperceptible Warping-based Backdoor Attack. In *Proc. Int'l Conf. Learning Representations*.
- Rakin, A. S.; He, Z.; and Fan, D. 2020. Tbt: Targeted neural network attack with bit trojan. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 13198–13207.
- Rippel, O.; Nair, S.; Lew, C.; Branson, S.; Anderson, A. G.; and Bourdev, L. 2019. Learned video compression. In *Proc. IEEE Int'l Conf. Computer Vision*, 3454–3463.
- Saha, A.; Subramanya, A.; and Pirsiavash, H. 2020. Hidden trigger backdoor attacks. In *Proc. AAAI Conf. on Artificial Intelligence*, volume 34, 11957–11965.
- Shumitskaya, E.; Antsiferova, A.; and Vatolin, D. 2022. Universal perturbation attack on differentiable no-reference image-and video-quality metrics. In *BMVC*.
- Steinhardt, J.; Koh, P. W. W.; and Liang, P. S. 2017. Certified defenses for data poisoning attacks. In *Proc. Annual Conf. Neural Information Processing Systems*, volume 30.
- Su, S.; Yan, Q.; Zhu, Y.; Zhang, C.; Ge, X.; Sun, J.; and Zhang, Y. 2020. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 3667–3676.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Wallace, G. K. 1992. The JPEG still picture compression standard. *IEEE Trans. on Consumer Electronics*, 38: 43–59.
- Wang, B.; Yao, Y.; Shan, S.; Li, H.; Viswanath, B.; Zheng, H.; and Zhao, B. Y. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *IEEE Symposium on Security and Privacy*, 707–723.
- Wang, T.; Yao, Y.; Xu, F.; An, S.; Tong, H.; and Wang, T. 2022. An invisible black-box backdoor attack through frequency domain. In *European Conf. on Computer Vision*.
- Wang, X.; Hu, S.; Zhang, Y.; Zhou, Z.; Zhang, L. Y.; Xu, P.; Wan, W.; and Jin, H. 2024a. ECLIPSE: Expunging clean-label indiscriminate poisons via sparse diffusion purification. In *European Symposium on Research in Computer Security*.
- Wang, X.; Li, M.; Liu, W.; Zhang, H.; Hu, S.; Zhang, Y.; Zhou, Z.; and Jin, H. 2024b. Unlearnable 3D point clouds: Class-wise transformation is all you need. In *Proc. Annual Conf. Neural Information Processing Systems*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Processing*.
- Wu, D.; and Wang, Y. 2021. Adversarial neuron pruning purifies backdoored deep models. In *Advances in Neural Information Processing Systems*, volume 34, 16913–16925.
- Xia, S.; Yang, W.; Yu, Y.; Lin, X.; Ding, H.; DUAN, L.; and Jiang, X. 2024a. Transferable Adversarial Attacks on SAM and Its Downstream Models. In *Proc. Annual Conf. Neural Information Processing Systems*.
- Xia, S.; Yu, Y.; Jiang, X.; and Ding, H. 2024b. Mitigating the Curse of Dimensionality for Certified Robustness via Dual Randomized Smoothing. In *Proc. Int'l Conf. Learning Representations*.
- Yang, S.; Wu, T.; Shi, S.; Lao, S.; Gong, Y.; Cao, M.; Wang, J.; and Yang, Y. 2022. MANIQA: Multi-dimension Attention Network for No-Reference Image Quality Assessment. *arXiv preprint arXiv:2204.08958*.
- Yu, F.; Zeng, B.; Zhao, K.; Pang, Z.; and Wang, L. 2024a. Chronic Poisoning: Backdoor Attack against Split Learning. In *Proc. AAAI Conf. on Artificial Intelligence*.
- Yu, Y.; Wang, Y.; Xia, S.; Yang, W.; Lu, S.; Tan, Y.-P.; and Kot, A. C. 2024b. Purify Unlearnable Examples via Rate-Constrained Variational Autoencoders. In *Proc. Int'l Conf. Machine Learning*.
- Yu, Y.; Wang, Y.; Yang, W.; Guo, L.; Lu, S.; Duan, L.-Y.; Tan, Y.-P.; and Kot, A. C. 2024c. Robust and Transferable Backdoor Attacks Against Deep Image Compression With Selective Frequency Prior. *IEEE Trans. on Pattern Analysis and Machine Intelligence*.
- Yu, Y.; Wang, Y.; Yang, W.; Lu, S.; Tan, Y.-P.; and Kot, A. C. 2023. Backdoor attacks against deep image compression via adaptive frequency trigger. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 12250–12259.
- Yu, Y.; Yang, W.; Tan, Y.-P.; and Kot, A. C. 2022. Towards Robust Rain Removal Against Adversarial Attacks: A Comprehensive Benchmark Analysis and Beyond. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*.
- Yue, C.; Lv, P.; Liang, R.; and Chen, K. 2022. Invisible backdoor attacks using data poisoning in the frequency domain. *arXiv preprint arXiv:2207.04209*.
- Zeng, Y.; Park, W.; Mao, Z. M.; and Jia, R. 2021. Rethinking the backdoor attacks' triggers: A frequency perspective. In *Proc. IEEE Int'l Conf. Computer Vision*, 16473–16481.
- Zhang, A.; Ran, Y.; Tang, W.; and Wang, Y.-G. 2023. Vulnerabilities in video quality assessment models: The challenge of adversarial attacks. In *Proc. Annual Conf. Neural Information Processing Systems*, volume 36.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 586–595.
- Zhang, W.; Li, D.; Min, X.; Zhai, G.; Guo, G.; Yang, X.; and Ma, K. 2022. Perceptual attacks of no-reference image quality models with human-in-the-loop. In *Proc. Annual Conf. Neural Information Processing Systems*.
- Zhang, W.; Liu, Y.; Dong, C.; and Qiao, Y. 2019. Rankrgan: Generative adversarial networks with ranker for image super-resolution. In *Proc. IEEE Int'l Conf. Computer Vision*.
- Zhang, W.; Ma, K.; Yan, J.; Deng, D.; and Wang, Z. 2020. Blind Image Quality Assessment Using A Deep Bilinear Convolutional Neural Network. *IEEE Trans. on Circuits and Systems for Video Technology*, 30(1): 36–47.
- Zhang, W.; Ma, K.; Zhai, G.; and Yang, X. 2021. Uncertainty-aware blind image quality assessment in the laboratory and wild. *IEEE Trans. on Image Processing*, 30: 3474–3486.
- Zheng, Q.; Yu, Y.; Yang, S.; Liu, J.; Lam, K.-Y.; and Kot, A. 2024. Towards Physical World Backdoor Attacks against Skeleton Action Recognition. In *European Conf. on Computer Vision*.