

344

345 **Appendix**

346

347

348

**Table of Contents**

---

349	<b>A Technical Details for Generating Hardly-Generalized Domain</b>	<b>11</b>
350	A.1 The Implementation of Transformation Module . . . . .	11
351	A.2 The Optimization Process . . . . .	11
352	<b>B The Proof for Theorem 1</b>	<b>12</b>
353	<b>C Technical Details for Generating Protected Dataset</b>	<b>14</b>
354	C.1 The Optimization Solution for Generating Protected Dataset . . . . .	14
355	C.2 The Process of Generating Samples from Other Domains . . . . .	15
356	C.3 The Selection of Hyper-parameters . . . . .	15
357	<b>D The Proof for Theorem 2</b>	<b>16</b>
358	<b>E The Detailed Settings for Experimental Datasets and Configurations</b>	<b>17</b>
359	E.1 Datasets . . . . .	17
360	E.2 The Demonstration of Domain Watermark for Each Dataset . . . . .	18
361	E.3 Training Configurations. . . . .	18
362	E.4 The Details for Implementing each Approach . . . . .	19
363	<b>F The Additional Results for the Performance of Domain Watermark</b>	<b>20</b>
364	<b>G The Detailed Settings for Dataset Ownership Verification</b>	<b>20</b>
365	<b>H The Additional Results for Dataset Ownership Verification</b>	<b>20</b>
366	<b>I Additional Results of Discussions</b>	<b>20</b>
367	I.1 The Effects of $\lambda_3$ . . . . .	20
368	I.2 Performance under Different Domain Watermarks . . . . .	20
369	I.3 The Transferability of Domain Watermark . . . . .	21
370	<b>J Additional Results for the Resistance to Potential Adaptive Methods</b>	<b>23</b>
371	<b>K Reproducibility Statement</b>	<b>24</b>
372	<b>L Societal Impacts</b>	<b>24</b>
373	<b>M Discussions about Adopted Data</b>	<b>25</b>

---

374

375

376

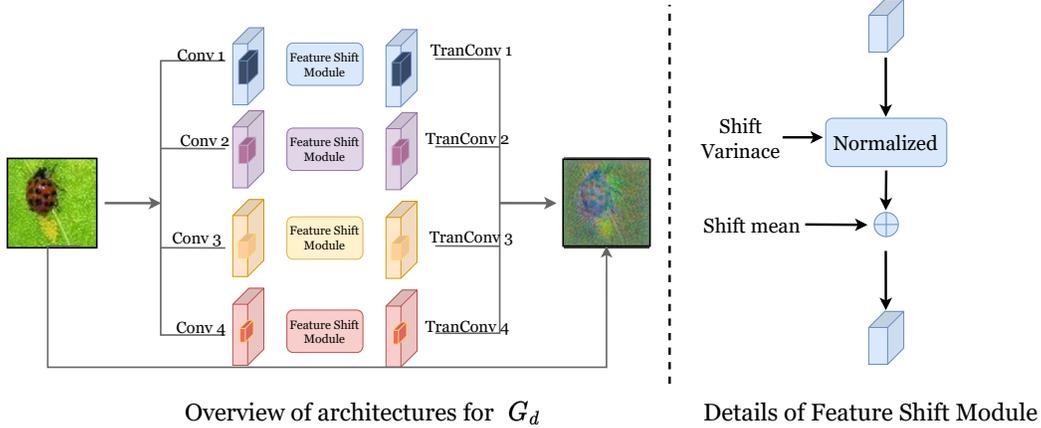


Figure 1: The architecture of  $G_d$ .

## 377 A Technical Details for Generating Hardly-Generalized Domain

378 This process is mainly motivated by [33], which leveraged a transformation module with different  
 379 convolution transformations to minimize the mutual information ( $I$ ) between features from the source  
 380 dataset (*i.e.*,  $\mathbf{z}$ ) and data from the target domain (*i.e.*,  $\hat{\mathbf{z}}$ ). Our work is also partially inspired by  
 381 previous work on generating unlearnable samples [45], which crafted effective unlearnable samples  
 382 by performing the bi-level optimization within each iteration.

### 383 A.1 The Implementation of Transformation Module

384 We follow previous work [33] to implement the transformation module for generating samples  
 385 from a different domain (as shown in Fig. 1). Specifically, we design the transformation module  
 386 as an ensemble of multiple (*i.e.*, 4) convolution operations. Each convolution operation contains a  
 387 convolution layer Conv, a feature shift module, and a corresponding transposed convolution layer  
 388 TranConv. The detailed parameters for each convolution layer  $\text{Conv}_i$  are detailed in Tab. 1. Following  
 389 each convolution layer  $\text{Conv}_i$ , we add a feature shift module to enhance the diversity of the generated  
 390 samples. Specifically, each feature shift module contains two learnable parameters  $\mu_i$  and  $\sigma_i$  as mean  
 391 shift and variance shift, following:

$$\sigma_i \cdot \frac{\text{Conv}_i(\mathbf{x}) - \mu}{\sigma} + \mu_i, \quad (1)$$

392 where  $\mu$  and  $\sigma$  represent the mean and covariance value for  $\text{Conv}_i(\mathbf{x})$ . Notably,  $\mu$  and  $\sigma$  are not  
 393 learnable parameters. Moreover, the parameters  $\mu_i$  and  $\sigma_i$  has the same dimension as the output  
 394 of  $\text{Conv}_i(\mathbf{x})$ . After that, we use a transposed convolution layer TranConv to turn the feature maps  
 395 generated by the above operations into a real instance, which has the same dimension as  $\mathbf{x}$ .

396 Putting all above, we generate the hard-generalized domain samples  $\hat{\mathbf{x}}$  following:

$$\hat{\mathbf{x}} = \frac{1}{\sum w_i} \sum_i w_i \cdot \text{tahn}(\text{TranConv}(\sigma_i \cdot \frac{\text{Conv}_i(\mathbf{x}) - \mu}{\sigma} + \mu_i)), \quad (2)$$

397 where tahn represents the tahn activation function.  $w_i$  is a scalar and weights the contribution of  
 398 each activated instance produced by TransposedConv to  $\hat{\mathbf{x}}$ .  $w_i$  is randomly sampled from normal  
 399 distribution  $w_i \sim N(0, 1)$ . Notably, for each input  $\mathbf{x}$ , we first up-sample it to  $224 \times 224$  size and  
 400 down-sample produced  $\hat{\mathbf{x}}$  to the original size for  $\mathbf{x}$ .

### 401 A.2 The Optimization Process

402 During the optimization process of Eq. (2), we first initialized a surrogate model  $f(\cdot; \mathbf{w})$  and a  
 403 benign dataset  $\mathcal{D}$ . Then during each iteration for solving the bi-level optimization Eq. (2), we first  
 404 minimize the  $I(\mathbf{z}; \hat{\mathbf{z}})$  and  $\mathcal{L}_c$  by optimizing the parameters of our proposed transformation module:

$$\min_{\theta} \mathbb{E}_{p(\mathbf{z}, \hat{\mathbf{z}})} [I(\mathbf{z}(\mathbf{w}^*); \hat{\mathbf{z}}(\theta, \mathbf{w}^*)) + \lambda_1 \mathcal{L}_c(\mathbf{z}(\mathbf{w}^*), \hat{\mathbf{z}}(\theta, \mathbf{w}^*))]. \quad (3)$$

Table 1: The configuration for each Convolution layer Conv.

Model	Kernel Size	Input Channel	Output Channel
Conv <sub>1</sub>	5x5	3	3
Conv <sub>2</sub>	9x9	3	3
Conv <sub>3</sub>	13x13	3	3
Conv <sub>4</sub>	17x17	3	3

405 After that, we maximize  $I(\mathbf{z}; \hat{\mathbf{z}})$  and minimize the training loss by optimizing the parameters  $\mathbf{w}$ :

$$\min_{\mathbf{w}} \left[ \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\mathcal{L}(f(G_d(\mathbf{x}; \boldsymbol{\theta}); \mathbf{w}), y) + \mathcal{L}(f(\mathbf{x}; \mathbf{w}), y)] - \lambda_2 \mathbb{E}_{p(\mathbf{z}, \hat{\mathbf{z}})} [I(\mathbf{z}(\mathbf{w}); \hat{\mathbf{z}})] \right]. \quad (4)$$

406 Since  $I(\mathbf{z}; \hat{\mathbf{z}})$  is intractable, we propose to optimize its upper bound instead:

$$I(\mathbf{z}; \hat{\mathbf{z}}) = \mathbb{E}_{p(\mathbf{z}, \hat{\mathbf{z}})} \left[ \log \frac{p(\hat{\mathbf{z}}|\mathbf{z})}{p(\hat{\mathbf{z}})} \right] \leq \mathbb{E}_{p(\mathbf{z}, \hat{\mathbf{z}})} [\log p(\hat{\mathbf{z}}|\mathbf{z})] - \mathbb{E}_{p(\mathbf{z})p(\hat{\mathbf{z}})} [\log p(\hat{\mathbf{z}}|\mathbf{z})]. \quad (5)$$

407 Since the conditional distribution  $p(\hat{\mathbf{z}}|\mathbf{z})$  is also intractable thus the upper bound of  $I(\mathbf{z}; \hat{\mathbf{z}})$  can't be  
 408 optimized, we follow previous work to adopt a variational distribution  $q(\hat{\mathbf{z}}|\mathbf{z})$  to approximate the  
 409 upper bound of  $I(\mathbf{z}; \hat{\mathbf{z}})$ :

$$I(\mathbf{z}; \hat{\mathbf{z}}) \leq \frac{1}{N} \sum_{i=1}^N [\log q(\hat{\mathbf{z}}_i|\mathbf{z}_i)] - \frac{1}{N} \sum_{j=1}^N \log q(\hat{\mathbf{z}}_j|\mathbf{z}_i), \quad (6)$$

410 where  $q(\hat{\mathbf{z}}|\mathbf{z})$  is obtained by employing the backbone neural network to approximate.

411 We optimize the above bi-level optimization Eq. (2) with 100 iterations. We set the learning rate as  
 412 0.005 for optimizing the parameters of the proposed transformation module and 0.001 for parameters  
 413 for the backbone model  $f(\cdot)$  following [33]. The batch size is 64. For both the transformation  
 414 module and the backbone model  $f(\cdot)$ , we use SGD [46] as the optimizer with Nesterov momentum  
 415 and weight decay rate of 0.0005. We use ResNet-18 as the backbone model for extracting  $\mathbf{z}$  and  $\hat{\mathbf{z}}$   
 416 throughout the paper. We introduce  $\lambda_1$  and  $\lambda_2$  for balancing each optimization objective. Following  
 417 the implementation of [33], we set  $\lambda_1$  and  $\lambda_2$  as 0.1 and 1.0 for balancing each optimization objective.

## 418 B The Proof for Theorem 1

419 **Theorem 1** (Data Quantity Impact). *Suppose in PAC Bayesian [35], for a target domain  $\mathcal{T}$  and a*  
 420 *source domain  $\mathcal{S}$ , any set of voters (candidate models)  $\mathcal{H}$ , any prior  $\pi$  over  $\mathcal{H}$  before any training,*  
 421 *any  $\xi \in (0, 1]$ , any  $c > 0$ , with a probability at least  $1 - \xi$  over the choices of  $S \sim S^{n_s}$  and  $T \sim \mathcal{T}_{\mathcal{X}}^{n_t}$ ,*  
 422 *for the posterior  $f$  over  $\mathcal{H}$  after the joint training on  $S$  and  $T$ , we have*

$$\begin{aligned} \mathcal{R}_{\mathcal{T}}(f) &\leq \frac{c}{2(1 - e^{-c})} \widehat{\mathcal{R}}_{\mathcal{T}}(f) + \frac{c}{1 - e^{-c}} \beta_{\infty}(\mathcal{T}||\mathcal{S}) \widehat{\mathcal{R}}_{\mathcal{S}}(f) + \Omega \\ &+ \frac{1}{1 - e^{-c}} \left( \frac{1}{n_t} + \frac{\beta_{\infty}(\mathcal{T}||\mathcal{S})}{n_s} \right) \left( 2\text{KL}(f||\pi) + \ln \frac{2}{\xi} \right), \end{aligned} \quad (7)$$

423 where  $\widehat{\mathcal{R}}_{\mathcal{T}}(f)$  and  $\widehat{\mathcal{R}}_{\mathcal{S}}(f)$  are the target and source empirical risks measured over target and source  
 424 datasets  $T$  and  $S$ , respectively.  $\Omega$  is a constant and  $\text{KL}(\cdot)$  is the Kullback–Leibler divergence.  
 425  $\beta_{\infty}(\mathcal{T}||\mathcal{S})$  is a measurement of discrepancy between  $\mathcal{T}$  and  $\mathcal{S}$  defined as

$$\beta_{\infty}(\mathcal{T}||\mathcal{S}) = \sup_{(\mathbf{x}, y) \in \text{SUPP}(\mathcal{S})} \left( \frac{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{T}}}{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{S}}} \right) \geq 1, \quad (8)$$

426 where  $\text{SUPP}(\mathcal{S})$  denotes the support of  $\mathcal{S}$ . When  $\mathcal{S}$  and  $\mathcal{T}$  are identical,  $\beta_{\infty}(\mathcal{T}||\mathcal{S}) = 1$ .

427 *Proof.* Theorem 6 in Germain *et al.*'s work [47] demonstrates that *suppose in PAC Bayesian [35],*  
 428 *for a target domain  $\mathcal{T}$  and a source domain  $\mathcal{S}$ , any set of voters (candidate models)  $\mathcal{H}$ , any prior*

429  $\pi$  over  $\mathcal{H}$  before any training, any  $\xi \in (0, 1]$ , any  $c > 0$ , with a probability at least  $1 - \xi$  over the  
 430 choices of  $S \sim \mathcal{S}^{n_s}$  and  $T \sim \mathcal{T}_X^{n_t}$ , for the posterior  $f$  over  $\mathcal{H}$  after the joint training on  $S$  and  $T$ :

$$\begin{aligned} \mathcal{R}_{\mathcal{T}}(f) &\leq \frac{c}{2(1 - e^{-c})} \widehat{d}_{\mathcal{T}}(f) + \frac{c}{1 - e^{-c}} \beta_{\infty}(\mathcal{T} \parallel \mathcal{S}) \widehat{e}_{\mathcal{S}}(f) + \Omega \\ &\quad + \frac{1}{1 - e^{-c}} \left( \frac{1}{n_t} + \frac{\beta_{\infty}(\mathcal{T} \parallel \mathcal{S})}{n_s} \right) \left( 2\text{KL}(f \parallel \pi) + \ln \frac{2}{\xi} \right), \end{aligned} \quad (9)$$

431 where  $\mathcal{R}_{\mathcal{T}}(f)$  denotes the expected Gibbs risk of voter  $f$  over the target domain.  $\widehat{d}_{\mathcal{T}}(f)$  and  $\widehat{e}_{\mathcal{S}}(f)$  are  
 432 the empirical estimation of the target voters' disagreement and the source joint error, measured over  
 433 target and source datasets  $T$  and  $S$ , respectively.  $\Omega$  is a constant and  $\text{KL}(\cdot)$  is the Kullback–Leibler  
 434 divergence.  $\beta_{\infty}(\mathcal{T} \parallel \mathcal{S})$  is a measurement of discrepancy between  $\mathcal{T}$  and  $\mathcal{S}$  defined as

$$\beta_{\infty}(\mathcal{T} \parallel \mathcal{S}) = \sup_{(\mathbf{x}, y) \in \text{SUPP}(\mathcal{S})} \left( \frac{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{T}}}{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{S}}} \right), \quad (10)$$

435 where  $\text{SUPP}(\mathcal{S})$  denotes the support of  $\mathcal{S}$ .

436 In the following proof, in particular, the Gibbs risk  $\mathcal{R}_{\mathcal{A}}(f)$ , the voters' disagreement  $d_{\mathcal{A}}(f)$ , and the  
 437 joint error  $e_{\mathcal{A}}(f)$  of a certain domain  $\mathcal{A}$  are defined as follows

$$\mathcal{R}_{\mathcal{A}}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{A}} \mathbb{E}_{h \sim f} \mathbb{I}[h(\mathbf{x}) \neq y], \quad (11)$$

438

439

$$d_{\mathcal{A}}(f) = \mathbb{E}_{\mathbf{x} \sim \mathcal{A}_{\mathcal{X}}} \mathbb{E}_{h \sim f} \mathbb{E}_{h' \sim f} \mathbb{I}[h(\mathbf{x}) \neq h'(\mathbf{x})], \quad (12)$$

440

441

$$e_{\mathcal{A}}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{A}} \mathbb{E}_{h \sim f} \mathbb{E}_{h' \sim f} \mathbb{I}[h(\mathbf{x}) \neq y] \mathbb{I}[h'(\mathbf{x}) \neq y], \quad (13)$$

442 where  $\mathbb{I}[\text{True}] = 1$  if the inner condition is true, and otherwise  $\mathbb{I}[\text{False}] = 0$ , and  $\mathcal{A}_{\mathcal{X}}$  is the marginal  
 443 distribution of domain  $\mathcal{A}$ .  $h$  and  $h'$  are votes sampled from the posterior distribution  $f$  over  $\mathcal{H}$ . With  
 444 these definitions, studies [48, 49] reveal a relationship among the Gibbs risk, the voters' disagreement,  
 445 and the joint error as

$$\mathcal{R}_{\mathcal{A}}(f) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{A}} \mathbb{E}_{h \sim f} \mathbb{E}_{h' \sim f} \frac{\mathbb{I}[h(\mathbf{x}) \neq h'(\mathbf{x})] + 2\mathbb{I}[h(\mathbf{x}) \neq y \wedge h'(\mathbf{x}) \neq y]}{2} = \frac{1}{2} d_{\mathcal{A}}(f) + e_{\mathcal{A}}(f). \quad (14)$$

446 In this case, we can extend this relationship to the empirical estimations (suppose a dataset  $A$  is  
 447 sampled from domain  $\mathcal{A}$ ) as

$$\widehat{\mathcal{R}}_{\mathcal{A}}(f) = \frac{1}{|A|} \sum_{(\mathbf{x}, y) \sim A} \mathbb{E}_{h \sim f} \mathbb{E}_{h' \sim f} \frac{\mathbb{I}[h(\mathbf{x}) \neq h'(\mathbf{x})] + 2\mathbb{I}[h(\mathbf{x}) \neq y \wedge h'(\mathbf{x}) \neq y]}{2} = \frac{1}{2} \widehat{d}_{\mathcal{A}}(f) + \widehat{e}_{\mathcal{A}}(f). \quad (15)$$

448 Then we can use  $\widehat{\mathcal{R}}_{\mathcal{T}}(f)$  and  $\widehat{\mathcal{R}}_{\mathcal{S}}(f)$  to replace  $\widehat{d}_{\mathcal{T}}(f)$  and  $\widehat{e}_{\mathcal{S}}(f)$  in Eq. (9), respectively. In the  
 449 end, we can follow Xu *et al.* [50] to regard these empirical risks as data quantity-irrelevant when  
 450 analyzing the impact of data quantity.

451 Next, we focus on the proof of the numerical relationship  $\beta_{\infty}(\mathcal{T} \parallel \mathcal{S}) \geq 1$ . First of all,  $\beta_{\infty}(\mathcal{T} \parallel \mathcal{S})$   
 452 comes from a more general definition that is parameterized by a real value  $q > 0$ , shown as

$$\beta_q(\mathcal{T} \parallel \mathcal{S}) = \left[ \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{S}} \left( \frac{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{T}}}{\mathcal{P}_{(\mathbf{x}, y) \in \mathcal{S}}} \right)^q \right]^{\frac{1}{q}}. \quad (16)$$

453 For any  $q > 0$ ,  $\beta_q(\mathcal{T} \parallel \mathcal{S})$  can be also written as a Rényi Divergence-based form [47], *i.e.*,

$$\beta_q(\mathcal{T} \parallel \mathcal{S}) = 2^{\frac{q-1}{q} D_q(\mathcal{T} \parallel \mathcal{S})}, \quad (17)$$

454 where  $D_q(\mathcal{T} \parallel \mathcal{S})$  is the Rényi Divergence between  $\mathcal{T}$  and  $\mathcal{S}$  with the order  $q$ . For Rényi Divergence  
 455 with any order  $q > 0$ , there is a property of positivity [51], *i.e.*,  $D_q(\mathcal{T} \parallel \mathcal{S}) \geq 0$ . In this case, when  
 456  $q \rightarrow \infty$ , Eq. (17) becomes  $\beta_q(\mathcal{T} \parallel \mathcal{S}) = 2^{D_q(\mathcal{T} \parallel \mathcal{S})} \geq 1$ , and  $\beta_q(\mathcal{T} \parallel \mathcal{S}) = 2^{D_q(\mathcal{T} \parallel \mathcal{S})} = 1$  when  $\mathcal{T} = \mathcal{S}$ ,  
 457 in other words,  $D_q(\mathcal{T} \parallel \mathcal{S}) = 0$  when  $\mathcal{T} = \mathcal{S}$  [51].  $\square$

458 **C Technical Details for Generating Protected Dataset**

459 **C.1 The Optimization Solution for Generating Protected Dataset**

460 Recall Eq. (7) is :

$$\min_{\delta \in \mathcal{B}} \left[ \mathbb{E}_{(\hat{\mathbf{x}}, y) \sim \mathcal{T}} [\mathcal{L}(f(\hat{\mathbf{x}}; \mathbf{w}(\delta)), y)] - \lambda_3 \min \left\{ \mathbb{E}_{(\bar{\mathbf{x}}, y) \sim \bar{\mathcal{T}}} [\mathcal{L}(f(\bar{\mathbf{x}}; \mathbf{w}(\delta)), y)], \lambda_4 \right\} \right], \quad (18)$$

$$s.t. \mathbf{w}(\delta) = \arg \min_{\mathbf{w}} \left[ \frac{1}{|\mathcal{D}_s|} \sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_s} \mathcal{L}(f(\mathbf{x}_i + \delta_i; \mathbf{w}), y_i) + \frac{1}{|\mathcal{D}_b|} \sum_{(\mathbf{x}_j, y_j) \in \mathcal{D}_b} \mathcal{L}(f(\mathbf{x}_j; \mathbf{w}), y_j) \right],$$

461 where  $\mathbb{E}_{(\bar{\mathbf{x}}, y) \sim \bar{\mathcal{T}}} [\mathcal{L}(f(\bar{\mathbf{x}}; \mathbf{w}(\delta)), y)]$  represents the expected risk for the watermarked model on other  
 462 unseen domains (*i.e.*,  $\bar{\mathcal{T}}$ ) and  $\mathcal{B} = \{\delta : \|\delta\|_\infty \leq \epsilon\}$  where  $\epsilon$  is a visibility-related hyper-parameter.

463 The aforementioned problem is a standard bi-level problem, we following previous work [52, 53] to  
 464 leverage *gradient matching* to solving it. Specifically, we first make the following definition:

$$\mathcal{L}_t = \mathbb{E}_{(\hat{\mathbf{x}}, y) \sim \mathcal{T}} [\mathcal{L}(f(\hat{\mathbf{x}}; \mathbf{w}), y)] - \lambda_3 \min \left\{ \mathbb{E}_{(\bar{\mathbf{x}}, y) \sim \bar{\mathcal{T}}} [\mathcal{L}(f(\bar{\mathbf{x}}; \mathbf{w}), y)], \lambda_4 \right\}, \quad (19)$$

$$\mathcal{L}_i = \frac{1}{|\mathcal{D}_s|} \sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_s} \mathcal{L}(f(\mathbf{x}_i + \delta_i; \mathbf{w}), y_i). \quad (20)$$

465 According to the gradient-matching technique [52, 53], we have the Upper-level Sub-problem as:

$$\max_{\delta \in \mathcal{B}} \frac{\nabla_{\mathbf{w}} \mathcal{L}_t \cdot \nabla_{\mathbf{w}} \mathcal{L}_i}{\|\nabla_{\mathbf{w}} \mathcal{L}_t\| \cdot \|\nabla_{\mathbf{w}} \mathcal{L}_i\|}, \quad (21)$$

466 where we aim to maximize the gradient matching degree between  $\nabla_{\mathbf{w}} \mathcal{L}_t$  and  $\nabla_{\mathbf{w}} \mathcal{L}_i$  using  $\text{cosine}(\cdot)$   
 467 similarity as the metric through optimizing  $\delta$ . We solve the above Upper-level Sub-problem via  
 468 projected gradient ascend (PGA). We here use calculate  $\mathbb{E}_{(\hat{\mathbf{x}}, y) \sim \mathcal{T}} [\mathcal{L}(f(\hat{\mathbf{x}}; \mathbf{w}), y)]$  following:

$$\mathbb{E}_{(\hat{\mathbf{x}}, y) \sim \mathcal{T}} [\mathcal{L}(f(\hat{\mathbf{x}}; \mathbf{w}), y)] = \frac{1}{N} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \mathcal{L}(f(G_d(\mathbf{x}); \mathbf{w}), y). \quad (22)$$

469 Regarding the Lower-level Sub-problem, we have:

$$\min_{\mathbf{w}} \left[ \frac{1}{|\mathcal{D}_s|} \sum_{(\mathbf{x}_i, y_i) \in \mathcal{D}_s} \mathcal{L}(f(\mathbf{x}_i + \delta_i; \mathbf{w}), y_i) + \frac{1}{|\mathcal{D}_b|} \sum_{(\mathbf{x}_j, y_j) \in \mathcal{D}_b} \mathcal{L}(f(\mathbf{x}_j; \mathbf{w}), y_j) \right]. \quad (23)$$

470 After obtaining the poisoned dataset (*i.e.*,  $\mathcal{D}_s \cup \mathcal{D}_b$ ), we can optimize the model (*i.e.*, ResNet-  
 471 18) parameters  $\mathbf{w}$  via solving the above Lower-level Upper-sub problem. The above Lower-level  
 472 Upper-sub problem is solved via stochastic gradient descent.

473 We optimize the Upper-level and Lower-level Sub-problems alternatively for each optimization  
 474 iteration. Specifically, we first train the model under benign dataset  $\mathcal{D}$ . Then for each iteration, we  
 475 first optimize the Upper-level Sub-problem based on the trained model and obtain the perturbation  $\delta$ .  
 476 After that, we optimize the Lower-level Sub-problem based on the obtained poisoned dataset. During  
 477 each iteration for optimizing the above bi-level optimization problems, we optimize the Upper-level  
 478 Sub-problem with 50 iterations, and optimize the Lower-level Sub-problem with 100 iterations.  
 479 We optimize the entire bi-level optimization with five epochs. The other details for optimization  
 480 hyper-parameters as well as configuration are consistent with [53, 23].

481 In particular, to ensure the effectiveness of solving the aforementioned bi-level optimization problem,  
 482 we have two additional strategies, as follows:



Figure 2: The example of samples generated from various domain.

- 483 • **Strategy 1:** Instead of randomly selecting samples from benign dataset  $\mathcal{D}$ , we here choose  
484 to select training samples with the largest gradient norms, following the previous work [52].
- 485 • **Strategy 2:** Instead of selecting samples from all classes, we follow the previous work [4]  
486 to select those from a specific class and the selected class is set as the target label. This  
487 strategy can enhance the effectiveness for solving the above bi-level optimization problem  
488 while preserving the verification performance for our approach.

## 489 C.2 The Process of Generating Samples from Other Domains

490 In this part, we describe how to generate  $\bar{\mathcal{T}}$  samples from other domains (*i.e.*,  $(\bar{x}, y) \sim \bar{\mathcal{T}}$ ).

491 After obtaining the transformation module  $G_d(\cdot)$ , we can generate hard-generalized domain samples  
492 from a specific domain. We here propose to generate samples from other domains by setting different  
493 configurations of  $\{w_i\}_1^4$ . For example, we can generate samples from the other domain by sampling  
494  $\{w_i\}_1^4$  with another values following  $w_i \sim N(0, 1)$ .

495 We here show some demonstration of samples from other domains in Fig. 2.

496 We here generated samples from other domains, and estimate  $\mathbb{E}_{(\bar{x}, y) \sim \bar{\mathcal{T}}}[\mathcal{L}(f(\bar{x}; \mathbf{w}), y)]$  following:

$$\mathbb{E}_{(\bar{x}, y) \sim \bar{\mathcal{T}}}[\mathcal{L}(f(\bar{x}; \mathbf{w}(\delta)), y)] = \frac{1}{N} \frac{1}{J} \sum_j \sum_{(\bar{x}, y) \in \bar{\mathcal{T}}_j} \mathcal{L}(f(\bar{x}; \mathbf{w}), y), \quad (24)$$

497 where  $\bar{\mathcal{T}}_j$  represents the  $i$ -th unseen domain generated by the above approach.

## 498 C.3 The Selection of Hyper-parameters

499 After generating other unseen domains  $\mathcal{T}$ , we here describe the selection of hyper-parameters (*i.e.*,  $J$   
500 and  $\lambda_3$ ) for generating protected dataset.

501 We here propose a heuristic approach for selecting  $J$  and  $\lambda_3$ . Specifically, we first keep  $\lambda_3$  fixed  
502 (*i.e.*, 1) and adjust  $J$ . We conduct empirical study on CIFAR-10 tasks, the results are shown in Fig. 3.

503 We use ResNet-18 as the evaluated model. We generate several unseen domains using the above  
504 approach. We randomly select  $J$  of these domains for optimizing the Eq. (7), and select 3 unseen  
505 domains as the validation data. Notably, the validation domains are ensured visually different from  
506 the domains used for optimization.

507 From Fig. 3, we find that using  $\geq$  three unseen domains is sufficient to constrain the generalization  
508 performance for validation unseen domains. Therefore, we set  $J$  as 3 for our approach.

509 After that, we keep  $J$  fixed, and adjust  $\lambda_3$  gradually, the results are shown in Fig. 4. We find that when  
510  $\lambda_3$  becomes smaller, the constraint for performance on other unseen domains reduces. Accordingly,  
511 we set  $\lambda_3$  as 0.3 for our approach since it can achieve a close generalization capacity compared to the  
512 benign DNN model (*i.e.*, 24.3%).

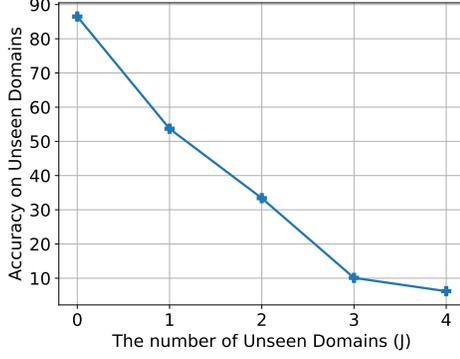


Figure 3: Effects of the number of Unseen Domains  $J$ .

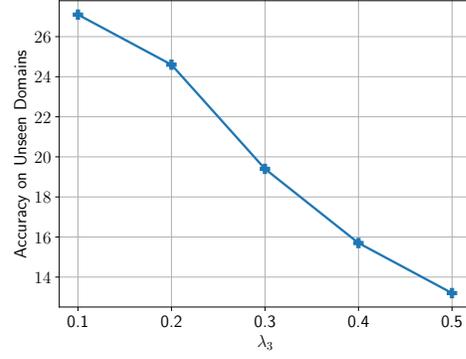


Figure 4: Effects of the  $\lambda_3$ .

## 513 D The Proof for Theorem 2

514 **Theorem 2.** Let  $f(\mathbf{x})$  is the posterior probability of  $\mathbf{x}$  predicted by the suspicious model, variable  
 515  $\mathbf{X}$  denotes the benign sample with label  $Y$ , and variable  $\mathbf{X}'$  is the domain-watermarked version of  
 516  $\mathbf{X}$ . Assume that  $P_b \triangleq f(\mathbf{X})_Y > \eta$ . We claim that dataset owners can reject the null hypothesis  $H_0$   
 517 at the significance level  $\alpha$ , if the verification success rate (VSR)  $V$  of  $f$  satisfies that

$$\sqrt{m-1} \cdot (V - \eta + \tau) - t_\alpha \cdot \sqrt{V - V^2} > 0, \quad (25)$$

518 where  $t_\alpha$  is the  $\alpha$ -quantile of  $t$ -distribution with  $(m-1)$  degrees of freedom and  $m$  is the sample size.

519 *Proof.* Since  $P_b > \eta$ , the original hypothesis  $H_1$  can be converted to

$$H_1^t : P_d > \eta - \tau. \quad (26)$$

520 Let  $E$  indicates the event of whether the suspect model  $f$  predicts a watermark sample as its ground-  
 521 truth label  $y$ . As such,  $E \sim B(1, p)$ , where  $p = \Pr(C(\mathbf{X}') = Y)$  indicates the verification success  
 522 probability and  $B$  is the Binomial distribution [36].

523 Let  $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_m$  denotes  $m$  domain-watermarked samples used for dataset verification and  
 524  $E_1, \dots, E_m$  denote their prediction events, we know that the verification success rate  $V$  satisfies

$$V = \frac{1}{m} \sum_{i=1}^m E_i, \quad (27)$$

$$V \sim \frac{1}{m} B(m, p). \quad (28)$$

525 According to the central limit theorem [36], the verification success rate  $V$  follows Gaussian distri-  
 526 bution  $\mathcal{N}(p, \frac{p(1-p)}{m})$  when  $m$  is sufficiently large. Similarly,  $(P_d - \eta + \tau)$  also satisfies Gaussian  
 527 distribution. Accordingly, we can construct the t-statistic as follows:

$$T \triangleq \frac{\sqrt{m}(W - \eta + \tau)}{s} \sim t(m-1), \quad (29)$$

528 where  $s$  is the standard deviation of  $(V - \eta + \tau)$  and  $V$ , *i.e.*,

$$s^2 = \frac{1}{m-1} \sum_{i=1}^m (E_i - V)^2 = \frac{1}{m-1} (m \cdot V - m \cdot V^2). \quad (30)$$

529 To reject the hypothesis  $H_0$  at the significance level  $\alpha$ , we need to ensure that

$$\frac{\sqrt{m}(V - \eta + \tau)}{s} > t_\alpha, \quad (31)$$

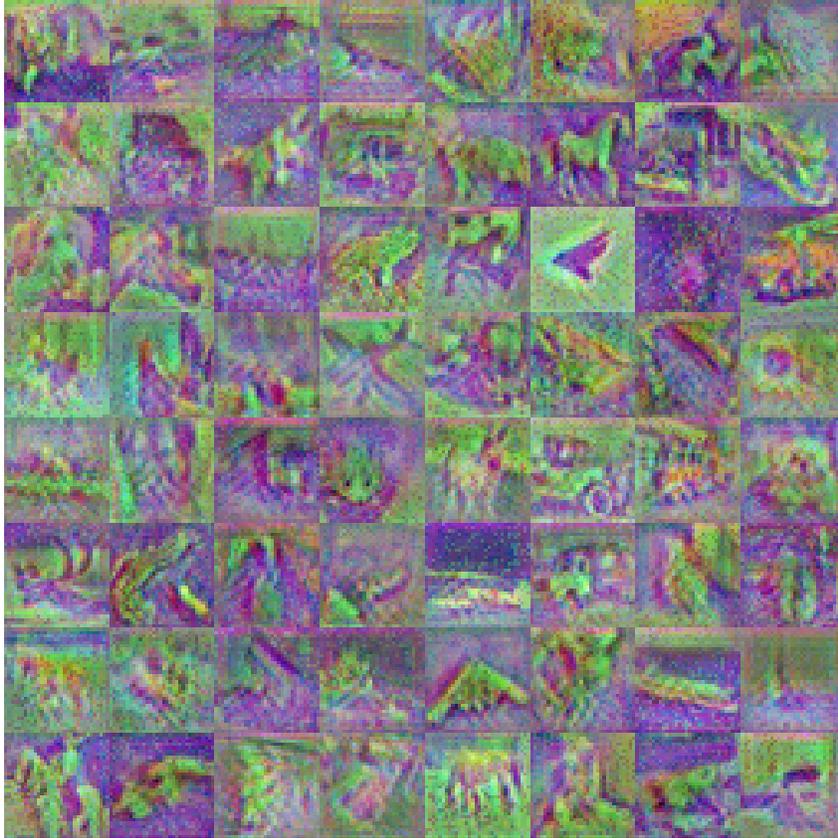


Figure 5: The example of domain watermark for CIFAR-10.

530 where  $t_\alpha$  is the  $\alpha$ -quantile of t-distribution with  $(m - 1)$  degrees of freedom.

531 According to equation (30)-(31), we have

$$\sqrt{m - 1} \cdot (V - \eta + \tau) - t_\alpha \cdot \sqrt{V - V^2} > 0. \quad (32)$$

532

□

## 533 E The Detailed Settings for Experimental Datasets and Configurations

### 534 E.1 Datasets

535 We evaluate our approach on three benchmark datasets (*i.e.*, CIFAR-10 [1], Tiny-ImageNet [37],  
536 STL-10 [40]). We here describe each benchmark dataset in detail.

537 **CIFAR-10.** CIFAR-10 dataset contains 10 labels, 50,000 training samples, and 10,000 validation  
538 samples. The training and validation samples are distributed evenly across each label. Each sample is  
539 resized as  $32 \times 32$  by default.

540 **Tiny-ImageNet.** Tiny-ImageNet dataset contains 200 labels, 100,000 training samples, and 10,000  
541 validation samples. The training and validation samples are distributed evenly across each label. Each  
542 sample is resized as  $64 \times 64$  by default.

543 **STL-10.** STL-10 dataset contains 10 labels and 13,000 labeled samples and 100,000 unlabeled  
544 samples. We divide the labeled samples into the training and validation dataset with a ratio of 8 : 2.  
545 The training and validation samples are distributed evenly across each label. Each sample is resized  
546 as  $96 \times 96$  by default.

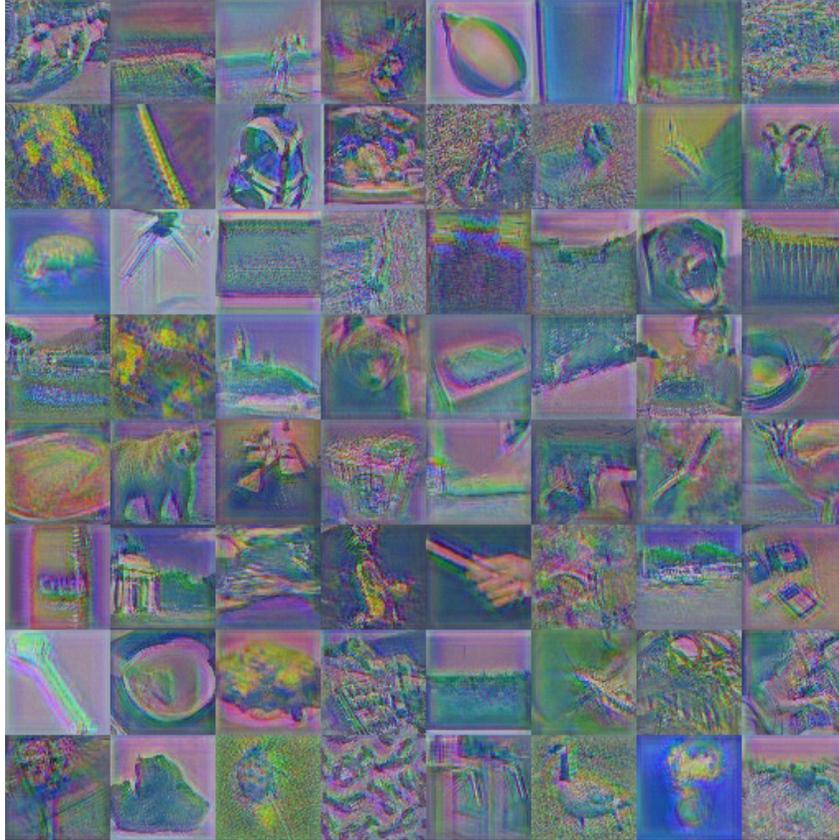


Figure 6: The example of domain watermark for Tiny-ImageNet.

Table 2: Summary of accuracy (%) on samples from different domains for normal models and ours.

Task	Source domain		Target domain		Other domain	
	Normal	Ours	Normal	Ours	Normal	Ours
CIFAR-10	91.89	90.86	13.10	90.45	15.10	10.30
STL-10	85.61	84.58	9.50	82.00	16.00	11.60
Tiny-ImageNet	60.13	59.10	6.00	58.08	12.60	15.40

547 **E.2 The Demonstration of Domain Watermark for Each Dataset**

548 We here show the domain watermark used for evaluating the effectiveness of our approach in the  
 549 experiments. The demonstrations are shown in Fig. 5, Fig. 6, and Fig. 7 for CIFAR-10, Tiny-ImageNet,  
 550 and STL-10 datasets, respectively.

551 **E.3 Training Configurations.**

552 In the experiments, we train each model with 150 epochs with an initialized learning rate of 0.1.  
 553 Following previous work [23], we schedule learning rate drops at epochs 14, 24, and 35 by a factor of  
 554 0.1. For all models, we employ SGD with Nesterov momentum, and we set the momentum coefficient  
 555 to 0.9. We use batches of 128 images and weight decay with a coefficient of  $4 \times 10^{-4}$ . For each  
 556 run, we report the verification success rate (VSR) averaged over the last 10 epochs when the models'  
 557 accuracy converges. We report the results for each approach averaged over 5 runs.



Figure 7: The example of domain watermark for STL-10.

Table 3: The watermark performance on STL-10 dataset. In particular, we mark harmful watermark results (*i.e.*,  $H > 0.5$  and  $\hat{H} > 0$ ) in red.

Label Type↓	Method↓, Metric→	STL-10			
		BA (%)	VSR (%)	$H$	$\hat{H}$
Poisoned-Label	BadNets	85.61	100	1.00	0.86
	Blended	85.21	99.32	1.00	0.84
	WaNet	83.17	96.10	0.96	0.79
	UBW-P	84.22	80.27	0.80	0.64
Clean-Label	Label-Consistent	84.07	93.48	0.93	0.77
	Sleeper Agent	83.72	89.77	0.90	0.73
	UBW-C	79.32	82.00	0.82	0.61
	DW (Ours)	84.58	82.00	0.18	-0.73

#### 558 E.4 The Details for Implementing each Approach

559 We implement each backdoor technique using Backdoorbox library<sup>2</sup> following the default training  
 560 configurations. Specifically, for patch-based triggers, we use  $3 \times 3$ ,  $6 \times 6$ , and  $9 \times 9$  for CIFAR-10,  
 561 Tiny-ImageNet, and STL-10. Following previous work [4], for each approach, we randomly select a  
 562 label as the target label for ownership verification purposes. For the other input-specific trigger (*i.e.*,  
 563 WaNet [21]), we follow its default configuration to generate its specific trigger pattern.

<sup>2</sup><https://github.com/THUYimingLi/BackdoorBox>

Table 4: The effectiveness of dataset ownership verification via our domain watermark.

	STL-10		
	Independent-D	Independent-M	Malicious
$\Delta P$	0.68	0.78	0.04
p-value	0.95	0.98	$10^{-46}$

Table 5: Summary of accuracy (%) on samples from different domains for normal models and ours.

Domain Watermarks	Source domain		Target domain		Other domain	
	Normal	Ours	Normal	Ours	Normal	Ours
Domain Watermark I	92.46	92.10	18.50	91.40	16.30	17.60
Domain Watermark II	92.46	91.95	18.20	90.24	14.70	15.80
Domain Watermark III	92.46	91.85	19.60	90.64	18.40	14.90

## 564 F The Additional Results for the Performance of Domain Watermark

565 We first show the summary for the performance of our approach and benign samples on samples  
 566 from different domains. The results are shown in Tab. 2. We also show additional results for STL-10  
 567 dataset with ResNet-34 as shown in Tab. 3.

## 568 G The Detailed Settings for Dataset Ownership Verification

569 We evaluate our domain-watermark-based dataset ownership verification under three scenarios, includ-  
 570 ing **1**) independent domain (dubbed ‘Independent-D’), **2**) independent model (dubbed ‘Independent-  
 571 M’), and **3**) unauthorized dataset training (dubbed ‘Malicious’). In the first case, we used domain-  
 572 watermarked samples to query the suspicious model trained with modified samples from another  
 573 domain; In the second case, we test the benign model with our domain-watermarked samples; In the  
 574 last case, we test the domain-watermarked model with corresponding domain-watermarked samples.  
 575 Notice that only the last case should be regarded as having unauthorized dataset adoption. All other  
 576 settings are the same as those used in [4] and are demonstrated in our appendix.

577 Consistent with previous work [4], we adopt the trigger used in the training process of the water-  
 578 marked suspicious model in the last scenario. Moreover, we sample  $m = 100$  samples on CIFAR10,  
 579 STL-10, and Tiny-ImageNet and set  $\tau = 0.25$  for the hypothesis-test in each case for our approach.  
 580 Since Tiny-ImageNet has only 50 samples for each class in the validation dataset, we combine  
 581 additional 50 training samples with the validation samples for ownership verification. The additional  
 582 50 training samples are not used in generating the protected dataset.

## 583 H The Additional Results for Dataset Ownership Verification

584 We here investigate the effectiveness of ownership verification via our domain watermark. The results  
 585 are shown in Tab. 4. The settings are consistent with Section 5.

## 586 I Additional Results of Discussions

### 587 I.1 The Effects of $\lambda_3$

588 We have investigated the effects of  $\lambda_3$ , as shown in Fig. 4. We find that the generalization performance  
 589 decreases on other unseen validation domains with the increase of  $\lambda_3$ . When  $\lambda_3$  increases up to  
 590 0.3, the generalization performance on other unseen validation domains decreases close to the  
 591 generalization performance for benign models.

### 592 I.2 Performance under Different Domain Watermarks

593 We here investigate the effective of protected dataset generation for different domain watermarks. We  
 594 here craft domain watermarks following the Appendix. A but initialized with different parameters



Figure 8: The Demonstration of Domain Watermark I.

Table 6: The performance of our *domain watermark* with different model structures trained on the watermarked dataset generated with ResNet-18.

Metric↓, Model→	ResNet-18	ResNet-34	VGG-16-BN	VGG-19-BN
BA (%)	91.39	92.54	90.86	92.57
VSR (%)	91.90	90.80	90.48	89.00

595 for crafting different domain watermarks. The demonstrations for different domain watermarks for  
 596 CIFAR-10 are shown in Figs. 8 to 10.

597 We here use CIFAR-10 with ResNet-34 to investigate the performance of our approach for different  
 598 domain watermarks. The results are summarized in Tab. 5. We can see our approach can still achieve  
 599 effectiveness for different domain watermarks.

### 600 I.3 The Transferability of Domain Watermark

601 Recall that in the optimization process of our approach, we leverage a surrogate model (*i.e.*, ResNet-  
 602 18) for crafting modified samples. In the experiment section, we test the effectiveness of our approach  
 603 under models (*i.e.*, VGG-16-BN and ResNet-34) having different architectures and parameters  
 604 from the surrogate model. In practice, dataset users may adopt different model structures since  
 605 dataset owners have no information about the model training. In this section, we conduct additional  
 606 experiments on evaluating the effectiveness of our approach under different structures compared to  
 607 the one used for generating modified samples (*i.e.*, transferability).

608 **Settings.** We evaluate the transferability of our approach under CIFAR-10 task. We adopt ResNet-  
 609 18, ResNet-34, VGG-16-BN, and VGG-19-BN to perform *domain watermark*, based on which to train



Figure 9: The Demonstration of Domain Watermark II.

Table 7: The performance of our *domain watermark* with different model structures trained on the watermarked dataset generated with ResNet-34.

Metric↓, Model→	ResNet-18	ResNet-34	VGG-16-BN	VGG-19-BN
BA (%)	91.22	92.56	90.43	91.79
VSR (%)	90.10	92.44	89.60	90.36

Table 8: The performance of our *domain watermark* with different model structures trained on the watermarked dataset generated with VGG-16-BN.

Metric↓, Model→	ResNet-18	ResNet-34	VGG-16-BN	VGG-19-BN
BA (%)	91.57	92.10	90.53	92.10
VSR (%)	90.70	91.60	90.44	89.84

Table 9: The performance of our *domain watermark* with different model structures trained on the watermarked dataset generated with VGG-19-BN.

Metric↓, Model→	ResNet-18	ResNet-34	VGG-16-BN	VGG-19-BN
BA (%)	91.48	91.98	90.77	92.73
VSR (%)	91.30	89.60	90.36	91.94

610 different models (*i.e.*, ResNet-18, ResNet-34, VGG-16-BN, and VGG-19-BN). Except for the model  
 611 structure, all other settings are the same as those used in Section 5.

612 **Results.** As shown in Tabs. 6 to 9, our approach has high transferability across model structures.  
 613 Accordingly, our methods are practical in protecting open-sourced datasets.

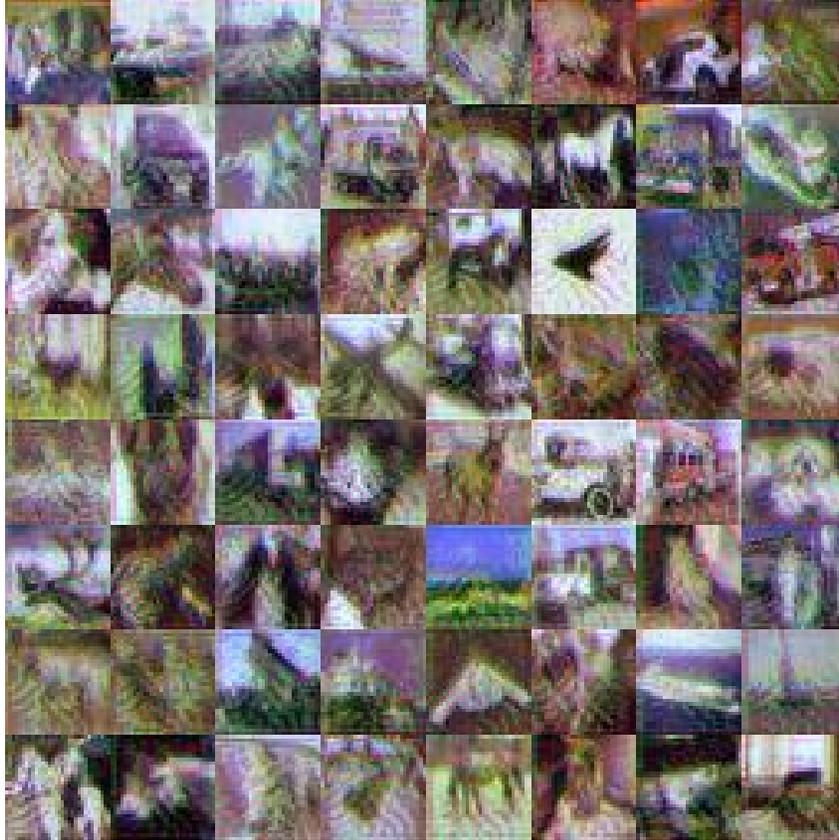


Figure 10: The Demonstration of Domain Watermark III.

## 614 J Additional Results for the Resistance to Potential Adaptive Methods

615 **Robustness against ShrinkPad.** We here investigate the robustness of our approach against  
 616 ShrinkPad [54], which is a well-known watermarked sample detection approach based on a set of  
 617 input transformations. We follow BackdoorBox to implement ShrinkPad for filtering watermarked  
 618 samples. We use CIFAR-10 with ResNet-34 to implement *domain watermark* and craft 1,000  
 619 watermarked samples based on the validation dataset for investigation. We first filter 900 watermarked  
 620 samples that can be correctly classified. We find ShrinkPad can only filter 87 effective watermarked  
 621 samples among 900 samples ( $\leq 10\%$ ), which means that our domain watermark is robust against  
 622 ShrinkPad.

623 **Robustness against Scale-UP.** We also evaluate our approach with the most recently input-level  
 624 watermark detection approach, Scale-UP [55]. We follow their released code <sup>3</sup> to implement SCALE-  
 625 UP and use the AUROC score as the metric to report the results. We test our approach on SCALE-UP  
 626 with 1,000 watermarked and 1,000 benign samples. We here use CIFAR-10 with ResNet-34.

627 We find that SCALE-UP yields around 0.58 AUROC score on our proposed *domain watermark*.  
 628 Such results imply that SCALE-UP can not perform against our domain watermark, with the filtering  
 629 performance close to random guesses. We think it may be caused by that, different from the previous  
 630 backdoor-inspired watermark causing misclassification, *domain watermark* leads the watermarked  
 631 model correctly classifying the watermarked samples. Therefore, the watermarked samples would  
 632 have a similar scaled prediction consistency as benign samples, since they all belong to the ground-  
 633 truth label and can be clustered closely as shown in Section 5.3.2.

<sup>3</sup><https://github.com/JunfengGo/SCALE-UP>

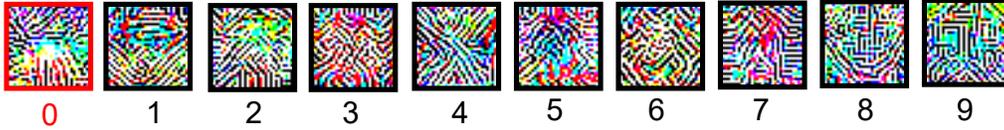


Figure 11: The reversed trigger maps for each label produced by Neural Cleanse.

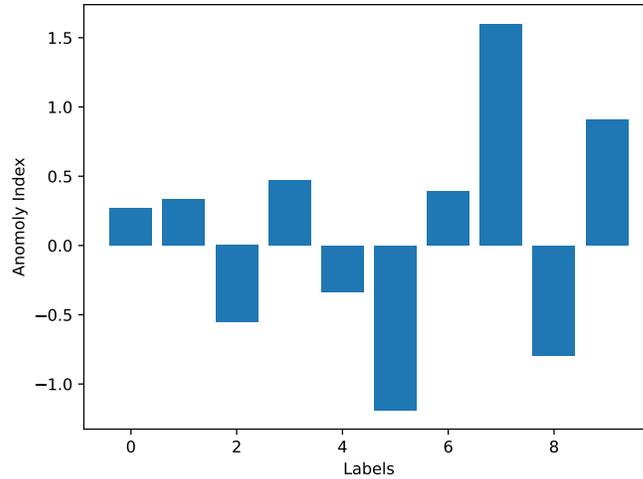


Figure 12: The anomaly index for  $\ell_1$  norm computed on the reversed trigger maps for each label produced by Neural Cleanse.

634 **Robustness against Neural Cleanse.** Following previous work [52], we also evaluate our approach  
 635 against Neural Cleanse [56]. We select label 0 as the target label and use CIFAR-10 with ResNet-34.  
 636 The results are shown in Fig. 11 and Fig. 12. We can see the reversed trigger pattern produced  
 637 by Neural Cleanse for the target label is extremely dense. We further follow [56] to calculate the  
 638 anomaly index for each label using MAD outlier detection approach. We find that the target label’s  
 639 anomaly index is smaller than 2, thus it would not be detected.

## 640 K Reproducibility Statement

641 In the appendix, we provide detailed descriptions of the datasets, models, training and evaluation  
 642 settings, and computational facilities. The codes and model checkpoints for reproducing the main  
 643 experiments of our evaluation are also provided in the supplementary material. We will release the  
 644 training codes of our methods upon the acceptance of this paper.

## 645 L Societal Impacts

646 In this paper, we focus on the copyright protection of (open-sourced) datasets. Specifically, we reveal  
 647 the harmful nature of backdoor-based dataset ownership verification (DOV) and proposed the first  
 648 non-backdoor-based DOV method that is truly harmless. This work has no ethical issues in general  
 649 since our method is purely defensive and does not reveal any new vulnerabilities of DNNs. However,  
 650 we need to mention that our method requires a sufficiently large watermarking rate and therefore  
 651 can not be used to protect a few or a single image. In addition, although our method is resistant to  
 652 existing adaptive methods, adversaries may try to develop more effective attacks against our DOV  
 653 method given the exposure of this paper. People should not be too optimistic about dataset protection.

## 654 M Discussions about Adopted Data

655 In this paper, all adopted samples are from the open-sourced datasets (*i.e.*, CIFAR-10, Tiny-ImageNet,  
656 and STL-10). The Tiny-ImageNet dataset may contain a few human-related images. We admit that we  
657 modified a few samples for watermarking and verification. However, our research treats all samples  
658 the same and the verification samples and modified samples have no offensive content. Accordingly,  
659 our work fulfills the requirements of these datasets and has no privacy violation.

## 660 References

- 661 [1] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- 662 [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale  
663 hierarchical image database. In *CVPR*, 2009.
- 664 [3] Yuchen Sun, Tianpeng Liu, Panhe Hu, Qing Liao, Shouling Ji, Nenghai Yu, Deke Guo, and  
665 Li Liu. Deep intellectual property: A survey. *arXiv preprint arXiv:2304.14613*, 2023.
- 666 [4] Yiming Li, Yang Bai, Yong Jiang, Yong Yang, Shu-Tao Xia, and Bo Li. Untargeted backdoor  
667 watermark: Towards harmless and stealthy dataset copyright protection. In *NeurIPS*, 2022.
- 668 [5] Ruixiang Tang, Qizhang Feng, Ninghao Liu, Fan Yang, and Xia Hu. Did you train on my  
669 dataset? towards public dataset protection with clean-label backdoor watermarking. *arXiv  
670 preprint arXiv:2303.11470*, 2023.
- 671 [6] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*,  
672 2001.
- 673 [7] Paulo Martins, Leonel Sousa, and Artur Mariano. A survey on fully homomorphic encryption:  
674 An engineering perspective. *ACM Computing Surveys*, 2017.
- 675 [8] Hua Deng, Zheng Qin, Qianhong Wu, Zhenyu Guan, Robert H Deng, Yujue Wang, and Yunya  
676 Zhou. Identity-based encryption transformation for flexible sharing of encrypted data in public  
677 cloud. *IEEE Transactions on Information Forensics and Security*, 2020.
- 678 [9] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. In *NeurIPS*, 2019.
- 679 [10] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference  
680 attacks against machine learning models. In *IEEE S&P*, 2017.
- 681 [11] Jiawang Bai, Yiming Li, Jiawei Li, Xue Yang, Yong Jiang, and Shu-Tao Xia. Multinomial  
682 random forest. *Pattern Recognition*, 2022.
- 683 [12] Sahar Haddad, Gouenou Coatrieux, Alexandre Moreau-Gaudry, and Michel Cozic. Joint  
684 watermarking-encryption-jpeg-ls for medical image reliability control in encrypted and com-  
685 pressed domains. *IEEE Transactions on Information Forensics and Security*, 2020.
- 686 [13] Run Wang, Felix Juefei-Xu, Meng Luo, Yang Liu, and Lina Wang. Faketagger: Robust  
687 safeguards against deepfake dissemination via provenance tracking. In *ACM MM*, 2021.
- 688 [14] Zhenyu Guan, Junpeng Jing, Xin Deng, Mai Xu, Lai Jiang, Zhou Zhang, and Yipeng Li.  
689 Deepmih: Deep invertible network for multiple image hiding. *IEEE Transactions on Pattern  
690 Analysis and Machine Intelligence*, 2022.
- 691 [15] Yiming Li, Ziqi Zhang, Jiawang Bai, Baoyuan Wu, Yong Jiang, and Shu-Tao Xia. Open-sourced  
692 dataset protection via backdoor watermarking. In *NeurIPS Workshop*, 2020.
- 693 [16] Yiming Li, Mingyan Zhu, Xue Yang, Yong Jiang, Tao Wei, and Shu-Tao Xia. Black-box dataset  
694 ownership verification via backdoor watermarking. *IEEE Transactions on Information Forensics  
695 and Security*, 2023.
- 696 [17] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor learning: A survey. *IEEE  
697 Transactions on Neural Networks and Learning Systems*, 2022.

- 698 [18] Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Evaluating  
699 backdooring attacks on deep neural networks. *IEEE Access*, 2019.
- 700 [19] Xiangyu Qi, Tinghao Xie, Yiming Li, Saeed Mahloujifar, and Prateek Mittal. Revisiting the  
701 assumption of latent separability for backdoor defenses. In *ICLR*, 2023.
- 702 [20] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on  
703 deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- 704 [21] Anh Nguyen and Anh Tran. Wanet—imperceptible warping-based backdoor attack. In *ICLR*,  
705 2021.
- 706 [22] Yinghua Gao, Yiming Li, Linghui Zhu, Dongxian Wu, Yong Jiang, and Shu-Tao Xia. Not all  
707 samples are born equal: Towards effective clean-label backdoor attacks. *Pattern Recognition*,  
708 2023.
- 709 [23] Hossein Souri, Liam H Fowl, Rama Chellappa, Micah Goldblum, and Tom Goldstein. Sleeper  
710 agent: Scalable hidden trigger backdoors for neural networks trained from scratch. In *NeurIPS*,  
711 2022.
- 712 [24] Ronald Rivest. The md5 message-digest algorithm. Technical report, 1992.
- 713 [25] Chiou-Ting Hsu and Ja-Ling Wu. Hidden digital watermarks in images. *IEEE Transactions on*  
714 *image processing*, 1999.
- 715 [26] Ming-Shing Hsieh, Din-Chang Tseng, and Yong-Huai Huang. Hiding digital watermarks using  
716 multiresolution wavelet transform. *IEEE Transactions on industrial electronics*, 2001.
- 717 [27] Yuanfang Guo, Oscar C Au, Rui Wang, Lu Fang, and Xiaochun Cao. Halftone image wa-  
718 termarking by content aware double-sided embedding error diffusion. *IEEE Transactions on*  
719 *Image Processing*, 2018.
- 720 [28] Zuobin Xiong, Zhipeng Cai, Qilong Han, Arwa Alrawais, and Wei Li. Adgan: Protect your  
721 location privacy in camera data of auto-driving vehicles. *IEEE Transactions on Industrial*  
722 *Informatics*, 17(9):6200–6210, 2020.
- 723 [29] Yiming Li, Peidong Liu, Yong Jiang, and Shu-Tao Xia. Visual privacy protection via mapping  
724 distortion. In *ICASSP*, 2021.
- 725 [30] Honghui Xu, Zhipeng Cai, Daniel Takabi, and Wei Li. Audio-visual autoencoding for privacy-  
726 preserving video streaming. *IEEE Internet of Things Journal*, 2021.
- 727 [31] Linghui Zhu, Xinyi Liu, Yiming Li, Xue Yang, Shu-Tao Xia, and Rongxing Lu. A fine-grained  
728 differentially private federated learning against leakage from gradients. *IEEE Internet of Things*  
729 *Journal*, 2021.
- 730 [32] Haiteng Zhao, Chang Ma, Qinyu Chen, and Zhi-Hong Deng. Domain adaptation via maximizing  
731 surrogate mutual information. In *IJCAI*, 2022.
- 732 [33] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to  
733 diversify for single domain generalization. In *ICCV*, 2021.
- 734 [34] Pengyu Cheng, Weituo Hao, Shuyang Dai, Jiachang Liu, Zhe Gan, and Lawrence Carin. Club:  
735 A contrastive log-ratio upper bound of mutual information. In *ICML*, 2020.
- 736 [35] David A McAllester. Some pac-bayesian theorems. In *COLT*, 1998.
- 737 [36] Leopold Schmetterer. *Introduction to mathematical statistics*, volume 202. Springer Science &  
738 Business Media, 2012.
- 739 [37] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 2015.
- 740 [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale  
741 image recognition. *ICLR*, 2014.

- 742 [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
743 recognition. In *CVPR*, 2016.
- 744 [40] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsuper-  
745 vised feature learning. In *AISTATS*, 2011.
- 746 [41] Yiming Li, Mengxi Ya, Yang Bai, Yong Jiang, and Shu-Tao Xia. BackdoorBox: A python  
747 toolbox for backdoor learning. In *ICLR Workshop*, 2023.
- 748 [42] Yuntao Liu, Yang Xie, and Ankur Srivastava. Neural trojans. In *ICCD*, 2017.
- 749 [43] Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models.  
750 In *NeurIPS*, 2021.
- 751 [44] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine  
752 learning research*, 2008.
- 753 [45] Hanxun Huang, Xingjun Ma, Sarah Monazam Erfani, James Bailey, and Yisen Wang. Unlearn-  
754 able examples: Making personal data unexploitable. *ICLR*, 2021.
- 755 [46] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *ICCS*, 2010.
- 756 [47] Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A new pac-  
757 bayesian perspective on domain adaptation. In *ICML*, 2016.
- 758 [48] Alexandre Lacasse, François Laviolette, Mario Marchand, Pascal Germain, and Nicolas Usunier.  
759 Pac-bayes bounds for the risk of the majority vote and the variance of the gibbs classifier.  
760 *NeurIPS*, 2006.
- 761 [49] Pascal Germain, Alexandre Lacasse, Francois Laviolette, Mario March, and Jean-Francis Roy.  
762 Risk bounds for the majority vote: From a pac-bayesian analysis to a learning algorithm. *Journal  
763 of Machine Learning Research*, 2015.
- 764 [50] Shichao Xu, Lixu Wang, Yixuan Wang, and Qi Zhu. Weak adaptation learning: Addressing  
765 cross-domain data insufficiency with weak annotator. In *ICCV*, 2021.
- 766 [51] Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE  
767 Transactions on Information Theory*, 2014.
- 768 [52] Hossein Souri, Micah Goldblum, Liam Fowl, Rama Chellappa, and Tom Goldstein. Sleeper  
769 agent: Scalable hidden trigger backdoors for neural networks trained from scratch. In *NeurIPS*,  
770 2022.
- 771 [53] Jonas Geiping, Liam Fowl, W Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller,  
772 and Tom Goldstein. Witches’ brew: Industrial scale data poisoning via gradient matching. In  
773 *ICLR*, 2021.
- 774 [54] Yiming Li, Tongqing Zhai, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor attack in the  
775 physical world. In *ICLR Workshop*, 2021.
- 776 [55] Junfeng Guo, Yiming Li, Xun Chen, Hanqing Guo, Lichao Sun, and Cong Liu. SCALE-UP: An  
777 efficient black-box input-level backdoor detection via analyzing scaled prediction consistency.  
778 In *ICLR*, 2023.
- 779 [56] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and  
780 Ben Y Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks.  
781 In *IEEE S&P*, 2019.
- 782