
Bucks for Buckets (B4B): Active Defenses Against Stealing Encoders

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Machine Learning as a Service (MLaaS) APIs provide ready-to-use and high-
2 utility encoders that generate vector representations for given inputs. Since these
3 encoders are very costly to train, they become lucrative targets for model stealing
4 attacks during which an adversary leverages query access to the API to replicate
5 the encoder locally at a fraction of the original training costs. We propose *Bucks*
6 *for Buckets (B4B)*, the first *active defense* that prevents stealing while the attack is
7 happening without degrading representation quality for legitimate API users. Our
8 defense relies on the observation that the representations returned to adversaries
9 who try to steal the encoder’s functionality cover a significantly larger fraction
10 of the embedding space than representations of legitimate users who utilize the
11 encoder to solve a particular downstream task. B4B leverages this to adaptively
12 adjust the utility of the returned representations according to a user’s coverage of
13 the embedding space. To prevent adaptive adversaries from eluding our defense by
14 simply creating multiple user accounts (sybils), B4B also individually transforms
15 each user’s representations. This prevents the adversary from directly aggregating
16 representations over multiple accounts to create their stolen encoder copy. Our
17 active defense opens a new path towards securely sharing and democratizing
18 encoders over public APIs.

19 1 Introduction

20 In model stealing attacks, adversaries extract a machine learning model exposed via a public API by
21 repeatedly querying it and updating their own stolen copy based on the obtained responses. Model
22 stealing was shown to be one of the main threats to the security of machine learning models in prac-
23 tice [36]. Also in research, since the introduction of the first extraction attack against classifiers [38],
24 a lot of work on improving stealing [26, 32, 38, 39], extending it to different model types [8, 35],
25 and proposing adequate defenses [17, 24, 25, 30] has been put forward. With the recent shift in
26 learning paradigms from supervised to self supervised learning (SSL), especially the need for new
27 defenses becomes increasingly pressing. From an academic viewpoint, the urge arises because it
28 was shown that SSL models (*encoders*) are even more vulnerable to model stealing [15, 28, 34] than
29 their supervised counterparts. This is because whereas supervised models’ output is low dimensional,
30 *e.g.*, per-class probabilities or pure labels, SSL encoders output high-dimensional representation
31 vectors that encode a larger amount of information and thereby facilitate stealing. In addition, from
32 a practical industry’s viewpoint, defenses are required since many popular API providers, such as
33 Cohere, OpenAI, or Clarify [1–3] already expose their high-value SSL encoders via APIs to a broad
34 range of users.

35 Most of the current defenses against encoder stealing are *reactive*, *i.e.*, they do not actively prevent
36 the stealing but rather aim at detecting it by adding watermarks to the encoder [13, 15] or performing
37 dataset inference to identify stolen copies [16]. Since at the point of detection, the damage of stealing
38 has already been inflicted, we argue that reactive defenses intervene too late and we advocate for

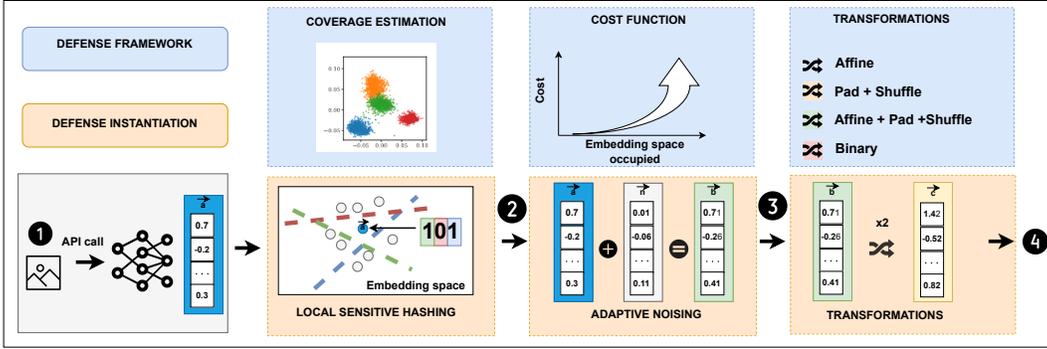


Figure 1: **Overview of B4B.** In the upper part, we present our B4B framework that consists of three modular building blocks: (1) A coverage estimation to track the fraction of embedding space covered by the representations returned to each user, (2) a cost function that serves to map the coverage to a concrete penalty to prevent stealing, and (3) per-user transformations that are applied to the returned representations to prevent sybil attacks. In the lower part, we present a concrete instantiation of B4B and the operation flow of our defense: ① The API calculates representations for the incoming queries. ② We instantiate the coverage estimation with local sensitive hashing and estimate the covered space as the fraction of *hash buckets* occupied. We calibrate the costs by adding noise to the representations according to the coverage. ③ We apply a set of transformations on a per-user basis. ④ The noised and transformed representations are returned to the user.

39 active defenses that prevent stealing while it is happening. Yet, active defenses are challenging to
 40 implement because they not only need to prevent stealing but also should preserve the utility of
 41 representations for legitimate users. The only existing active defense against encoder stealing [28]
 42 falls short on this latter aspect since it significantly degrades the quality of representations for all
 43 users.

44 To close the gap between required and existing defenses, we propose *Bucks for Buckets (B4B)*, the
 45 first active defense against encoder stealing that does not harm utility for legitimate users. B4B
 46 leverages the observation that the representations returned to adversaries who try to steal the encoder’s
 47 functionality cover a significantly larger fraction of the full embedding space than representations
 48 of legitimate users who utilize the encoder to solve a particular downstream task. To turn this
 49 observation into a practical defense, B4B is equipped with three modular building blocks: (1) The
 50 first building block is a tracking mechanism that continuously estimates the fraction of the embedding
 51 space covered by the representations returned to each user. The intuition why this is relevant is that
 52 by covering large fractions of the embedding space, the representations will suffice for an adversary
 53 to reproduce the encoder’s functionality, *i.e.*, to successfully steal it. (2) B4B’s second building
 54 block consists of a cost function to translate the covered fraction of the embedding space into a
 55 concrete penalty. We require this cost function to significantly penalize adversaries trying to steal the
 56 model while having only a minimal effect on legitimate users. (3) The third building block contains
 57 transformations that can be applied to the representations on a per-user basis to prevent adaptive
 58 attackers from circumventing our defense by creating multiple user accounts (sybils) and distributing
 59 their queries over these accounts such that they minimize the overall cost. We present the different
 60 building blocks of B4B in Figure 1.

61 While B4B’s modularity enables different instantiations of the three building blocks, we propose
 62 a concrete end-to-end instantiation to showcase the practicability of our approach. To implement
 63 tracking of the covered embedding space, we employ *local sensitive hashing* that maps any represen-
 64 tation returned to a given user into a set of hash *buckets*. We base our cost function (*i.e.*, the “*bucks*”)
 65 on utility and make B4B add noise to the representations with a magnitude that increases with the
 66 number of buckets occupied by the given user. While the scale of noise added to legitimate users’
 67 representations does not harm their downstream performance due to their small embedding space
 68 coverage, the representations returned to an adversary become increasingly noisy—significantly
 69 degrading the performance of their stolen encoder. Finally, we rely on a set of transformations
 70 (*e.g.*, affine transformations, shuffling, padding) that preserve downstream utility [16]. While, as a
 71 consequence, legitimate users remain unaffected by these transformations, adversaries cannot directly
 72 combine the representations obtained through different sybil accounts anymore to train their stolen

73 copy of the encoder. Instead, they first have to remap all representations into the same embedding
74 space, which we show causes both query and computation overhead and still reduces the performance
75 of the stolen encoder.

76 In summary, we make the following contributions:

- 77 1. We present B4B, the first active defense against encoder stealing that does not harm legitimate
78 users’ downstream performance. B4B’s three building blocks enable penalizing adversaries whose
79 returned representations cover large fractions of the embedding space and prevent sybil attacks.
- 80 2. We propose a concrete instantiation of B4B that relies on local sensitive hashing and decreases the
81 quality of representations returned to a user once their representations fill too many hash buckets.
- 82 3. We provide an end-to-end evaluation of our defense to highlight its effectiveness in offering high
83 utility representations for legitimate users and degrading the performance of stolen encoders in
84 both the single and the sybil-accounts setup.

85 2 Related Work

86 **Model Extraction Attacks.** The goal of the model extraction attacks is to replicate the functionality
87 of a victim model f_v trained on a dataset D_v . An attacker has a black box access to the victim model
88 and uses a stealing dataset $D_s = \{q_i, f_v(q_i)\}_{i=1}^n$, consisting of queries q_i and the corresponding
89 outputs $f_v(q_i)$ returned by the victim model, to train a stolen model f_s . Model extraction attacks
90 have been shown against various types of models including classifiers [23, 38] and encoders [15, 34].

91 **Self Supervised Learning and Encoders.** SSL is an increasingly popular machine learning
92 paradigm. It trains encoder models to generate representations from complex inputs without relying
93 on explicit labels. These representations encode useful features of a given input, enabling efficient
94 learning for multiple downstream tasks. Many SSL frameworks have been proposed [9–11, 21, 22, 42].
95 In our work, we focus on the two popular SSL vision encoders, namely SimSiam [11] and DINO [9],
96 which return high-quality representations that achieve state-of-the-art performance on downstream
97 tasks when assessed by training a linear classifier directly on representations. SimSiam trains with
98 two Siamese encoders with directly shared weights. A prediction MLP head is applied to one of the
99 encoders f_1 , and the other encoder f_2 has a stop-gradient, where both operations are used for avoiding
100 collapsing solutions. In contrast, DINO shares only architecture (not weights) between a student f_1
101 and a teacher model f_2 , also with the stop-gradient operation, but not the prediction head. While
102 SimSiam uses convolutional neural networks (CNNs), DINO also employs vision transformers (ViTs).
103 Both frameworks use a symmetrized loss of the form $\frac{1}{2}g(f_1(x_1), f_2(x_2)) + \frac{1}{2}g(f_1(x_2), f_2(x_1))$
104 in their optimization objectives, where $g(\cdot, \cdot)$ is negative cosine similarity for SimSiam and cross-
105 entropy for DINO. SimSiam and DINO’s similarities and differences demonstrate our method’s broad
106 applicability across SSL frameworks. More details can be found in Appendix E.

107 **Stealing Encoders.** The stealing of SSL encoders was shown to be extremely effective [15, 28, 34].
108 The goal of extracting encoders is to maximize the similarity of the outputs from the stolen local copy
109 and the original representations output by the victim encoder. Therefore, while training the stolen
110 copy, the adversary either imitates a self-supervised training using a contrastive loss function, *e.g.*,
111 InfoNCE [10] or SoftNN [20] or directly matches both models’ representations via the Mean Squared
112 Error (MSE) loss. To reduce the number of queries sent to the victim encoder, the attack proposed
113 in [28] leverages the key observation that the victim encoder returns similar representations for any
114 image and its augmented versions. Therefore, a given image can be sent to the victim while the
115 stolen copy is trained using many augmentations of this image, where the representation of a given
116 augmented image is approximated as the one of the original image produced by the victim encoder.

117 **Defending Encoders.** Recently, watermarking [7, 24, 40] methods have been proposed to detect
118 stolen encoders [13, 15, 41]. Many of these approaches use downstream tasks to check if a watermark
119 embedded into a victim encoder is present in a suspect encoder. Dataset inference [29] is another type
120 of encoder ownership resolution. It uses the victim’s training dataset as a unique signature, leveraging
121 the following observation: for a victim encoder trained on its private data as well as for its stolen
122 copies, the distribution of the representations generated from the victim’s training data differs from
123 the distribution of the representations generated on the test data. In contrast, for an independently
124 trained encoder, these two distributions cannot be distinguished, allowing the detection of stolen
125 copies [16]. However, all the previous methods are *reactive* and aim at detecting the stolen encoder

126 instead of *actively* preventing the attack. The only preliminary active defenses for encoders were
 127 proposed by [28]. They either perturb or truncate the answers to poison the training objective of an
 128 attacker. These operations were shown to harm substantially the performance of legitimate users,
 129 which renders the defense impractical. In contrast, our B4B has negligible impact on the quality of
 130 representations returned to legitimate users.

131 3 Actively Defending against Model Stealing with B4B

132 B4B aims at actively preventing model stealing while preserving high-utility representations for
 133 legitimate users. Before introducing the three main building blocks of B4B, namely (1) the estimation
 134 of embedding space coverage, (2) the cost function, and (3) the transformation of representations
 135 (see Figure 1), we detail our threat model and the observation on embedding space coverage that
 136 represents the intuition behind our approach.

137 3.1 Threat Model and Intuition

138 Our setup and the resulting threat model are inspired by public APIs, such as Cohere, OpenAI, or
 139 Clarify [1–3] that expose encoders to users through a pre-defined interface. These encoders are trained
 140 using SSL on large amounts of unlabeled data, often crawled from the internet, and therefore from
 141 diverse distributions. We notice that to provide rich representations to multiple users, the training
 142 dataset of the encoder needs to be significantly more diverse than the individual downstream tasks
 143 that the users query for representations. For instance, if the encoder behind the API is trained on the
 144 ImageNet dataset, then the legitimate users are expected to query the API for downstream tasks, such
 145 as CIFAR10 or SVHN. Similarly, if the encoder is trained on CIFAR10, the expected downstream
 146 tasks are MNIST or Fashion MNIST. Yet, in the design of our defense, we consider adversaries who
 147 can query the encoder with arbitrary inputs to obtain high-dimensional representation vectors from
 148 the encoder. Our defense is independent of the protected encoder’s architecture and does not rely on
 149 any assumption about the adversary’s data and query strategy.

150 We argue that even in this restricted setup, our defense can distinguish
 151 between adversaries and legitimate users by analyzing the distribu-
 152 tion of representations returned to them. In Figure 2, by using PCA
 153 to project representations for different datasets to a two-dimensional
 154 space, we visualize that representations for different downstream tasks
 155 cluster in *disjoint* and *small sub-spaces* of the full embedding space.
 156 The representations were obtained from a SimSiam encoder originally
 157 trained on ImageNet (we observe similar clustering for DINO shown
 158 in Appendix F). As a result, legitimate users can be characterized by
 159 their representations’ small coverage of the embedding space. In con-
 160 trast, the adversary does not aim at solving a particular downstream
 161 task. They instead would want to obtain representations that cover
 162 large fractions of the embedding space. This enables reproducing the
 163 overall functionality of the encoder (instead of only learning some
 164 local task-specific behavior). Indeed, it has been empirically shown
 165 by prior work, such as [15], that stealing with multiple distributions,
 166 *e.g.*, by relying on the complex ImageNet dataset, yields higher per-
 167 formance of the stolen encoder on various downstream tasks than
 168 stealing with a downstream dataset, such as CIFAR10. As a result, intuitively, we can identify and
 169 penalize adversaries based on their coverage of the embedding space, which will be significantly
 170 larger than the coverage of legitimate users. We leverage this intuition to build our B4B defense and
 171 present our three main building blocks in the following sections.

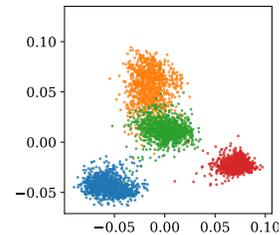


Figure 2: **Representations from Different Tasks Occupy Different Sub-Spaces of the Embedding Space. Presented for FashionMNIST, SVHN, CIFAR10, and STL10.**

172 3.2 Building Block 1: Coverage Estimation of the Embedding Space

173 The first building block of our B4B serves to estimate and continuously keep track of the fraction of
 174 the embedding space occupied by any given user. Let \mathcal{E} denote our embedding space of dimension s ,
 175 further, let U be a user with a query dataset $D = q_1, \dots, q_n \in \mathcal{D}$ and let $f_v : \mathcal{D} \rightarrow \mathcal{E}$ be our protected
 176 victim encoder that maps data points from the input to the embedding space. Assume user U has,
 177 so far, queried a subset of their data points q_1, \dots, q_j with $j \leq n$ to the encoder and obtained the

178 representations r_1, \dots, r_j with each $r_i \in \mathbb{R}^s$. We aim to estimate the true fraction of the embedding
 179 space \mathcal{E}_f^U that is covered by all returned representations r_1, \dots, r_j to user U and denote our estimate
 180 by $\tilde{\mathcal{E}}_f^U$.

181 **Local Sensitive Hashing.** One of the methods to approximate the occupied space by representations
 182 returned to a given user is via Local Sensitive Hashing (LSH) [37]. We rely on this approach for
 183 the concrete instantiation of our B4B and use it to track per-user coverage of the embedding space.
 184 Standard (cryptographic) hash functions are characterized by high dispersion such that hash collisions
 185 are minimized. In contrast, LSH hashes similar data points into the same or proximal, so-called *hash*
 186 *buckets*. This functionality is desired when dealing with searches in high-dimensional spaces or with a
 187 large number of data points. Formally, an LSH function \mathcal{H} is defined for a metric space $\mathcal{M} = (M, d)$,
 188 where d is a distance metric in space M , with a given threshold $T > 0$, approximation factors $f > 1$,
 189 and probabilities P_1 and P_2 , where $P_1 \gg P_2$. \mathcal{H} maps elements of the metric space to buckets
 190 $b \in B$ and satisfies the following conditions for any two points $q_1, q_2 \in M$: (1) If $d(q_1, q_2) \leq T$,
 191 then $\mathcal{H}(q_1) = \mathcal{H}(q_2)$ (*i.e.*, q_1 and q_2 collide in the same bucket b) with probability at least P_1 . (2) If
 192 $d(q_1, q_2) \geq fT$, then $\mathcal{H}(q_1) = \mathcal{H}(q_2)$ with probability at most P_2 .

193 3.3 Building Block 2: Cost Function Design

194 Once we can estimate the coverage of an embedding space for a given user U as $\tilde{\mathcal{E}}_f^U$, we need to design
 195 a cost function $\mathcal{C} : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that maps from the estimated coverage to a cost. The cost function
 196 needs to be designed such that it does not significantly penalize legitimate users while imposing a
 197 severe penalty on adversaries to effectively prevent the encoder from being stolen. The semantics of
 198 the cost function’s range depend on the type of costs that the defender wants to enforce. We discuss a
 199 broad range of options in Appendix C. These include monetary cost functions to adaptively charge
 200 users on a batch-query basis depending on their current coverage, costs in the form of additional
 201 computation that users need to perform in order to obtain their representations, similar to the proof of
 202 work in [17], costs in the form of delay in the interaction with the encoder behind the API [4], or
 203 costs in form of disk space that needs to be reserved by the user (similar to a proof of space [18, 19]).
 204 Which type of cost function is most adequate depends on the defender’s objective and setup.

205 **Exponential Cost Functions to Adjust Utility of Representations.** In the concrete instantiation
 206 of B4B that we present in this work, we rely on costs in the form of the utility of the returned
 207 representations. We choose this concrete instantiation because it is intuitive, effective, and can be
 208 directly experimentally assessed. Moreover, it is even suitable for public APIs where, for example,
 209 no monetary costs are applicable. We adjust utility by adding Gaussian noise with different standard
 210 deviation σ to the returned representations. Since we do not want to penalize legitimate users with
 211 small coverage but make stealing for adversaries with growing coverage increasingly prohibitive, we
 212 instantiate an exponential cost function that maps from the fraction of hash buckets occupied by the
 213 user to a value for σ . We choose the general form of this function as

$$f_{\lambda, \alpha, \beta}(\tilde{\mathcal{E}}_f^U) = \lambda \times (\exp^{\ln \frac{\alpha}{\lambda} \times \tilde{\mathcal{E}}_f^U \times \beta^{-1}} - 1) \quad (1)$$

214 where $\lambda < 1$ compresses the curve of the function to obtain low function values for small fractions
 215 of occupied buckets, and then we set a target penalty α for our cost function at a specified fraction
 216 of filled buckets β . For instance, if we want to enforce a σ of 5 at 90% of filled buckets (*i.e.*, for
 217 $\tilde{\mathcal{E}}_f^U = 0.9$), we would need to set $\alpha = 5$ and $\beta = 0.9$.

218 3.4 Building Block 3: Per-User Representation Transformations against Sybil Attacks

219 Given that our defense discourages users from querying with a wide variety of data points from
 220 different distributions, an adversary could create multiple fake user accounts (sybils) and query
 221 different data subsets with more uniform representations from each of these accounts. By combining
 222 all the obtained representations and using them to train a stolen copy, the adversary could overcome
 223 the increased costs of stealing. To defend against such sybil attacks, we propose individually
 224 transforming the representations on a per-user level before returning them. As a result, the adversary
 225 would first have to map all the representations to one single unified space before being able to jointly
 226 leverage the representations from different accounts for their stolen copy. Formally, for a given query

227 q_i , the protected victim encoder produces a representation $r_i = f_v(q_i)$, which is transformed by a
228 user-specific transformation $T_U(r_i)$ before being returned to the querying user U . We formulate two
229 concrete requirements for the transformations. First, they should preserve utility for legitimate users
230 on their downstream tasks, and second, they should be costly to reverse for an adversary.

231 **Utility Preserving Transformations.** As a concrete instantiation for our B4B, we propose a set of
232 transformations that fulfill the above-mentioned two requirements: (1) *Affine* where we apply affine
233 transformations to representations, (2) *Pad* where we pad representations with constant values, (3)
234 *Add* where we add constant values at random positions within representations, (4) *Shuffle* where we
235 shuffle the elements in the representation vectors, and (5) *Binary* where the original representations
236 are mapped to binary vectors relying on a random partitioning of the representation space. To
237 preserve the full amount of information contained in the original representations, in our binary
238 transformations, we tune the length of binary representations. We visualize the operation of each of
239 these transformations in Appendix C. All these transformations can additionally be combined with
240 each other, which further increases the possible set of transformations applied per user. This renders
241 it impossible for an adversary to correctly guess and reverse the applied representation. Instead, the
242 adversary has to remap the representations over all accounts into a single embedding space in order
243 to unify them and leverage them for training of their stolen encoder copy. We present an exhaustive
244 list of strategies that adversaries can apply for the remapping in Appendix D. All the attack methods
245 reduce to the minimum of remapping between representations of a pair of users, *i.e.*, they are at least
246 as complex as mapping between two separate accounts. In the next section, we show that our defense
247 already impedes stealing for an adversary with two accounts.

248 4 Empirical Evaluation

249 We first empirically evaluate our instantiation of B4B’s three building blocks and show how to calibrate
250 each of them for our defense. Finally, we provide an end-to-end evaluation that highlights B4B’s
251 effectiveness in preserving downstream utility for legitimate users while successfully preventing the
252 stealing by adversaries.

253 **Experimental Setup.** We conduct experiments on various kinds of downstream tasks and two
254 popular SSL encoders. To test our defense, we use FashionMNIST, SVHN, STL10, and CIFAR10 as
255 our downstream datasets, each with standard train and test splits. For stealing, we utilize training
256 data from ImageNet. We rely on encoder models from the SimSiam [11] and the DINO [9] SSL
257 frameworks. As our victim encoders, we use the publicly available ResNet50 model from SimSiam
258 trained for 100 epochs on ImageNet and the ViT Small DINO encoder trained for 800 epochs on
259 ImageNet, each using batch size 256. The ViT architecture takes as input a grid of non-overlapping
260 contiguous image patches of resolution $N \times N$. In this paper, we typically use $N = 16$. The Simsiam
261 encoder has an output representation dimension of 2048, while DINO returns 1536 dimensional
262 representations. We examine the utility of downstream classifiers using SimSiam’s or DINO’s
263 representations obtained for the respective downstream datasets. To implement LSH, we rely on
264 random projections [33] that we implement from scratch. For a detailed description of our stealing
265 and downstream training parameters, we refer to Appendix F.

266 4.1 Local Sensitive Hashing for Coverage Estimation

267 We first observe that the choice of the total number of hash buckets in the LSH influences the
268 effectiveness of our method. In the extreme, if we have a too large number of buckets, the number of
269 buckets filled will correspond to the number of queries posed by a user which fails to capture that
270 similar representations cover similar sub-spaces of the embedding space, and hence does not serve to
271 approximate the total fraction of the embedding space covered. However, if we have too few buckets,
272 even the representations for simple downstream tasks will fill large fractions of buckets, making it
273 impossible to calibrate the cost function such that it only penalizes adversaries. We experimentally
274 find that for our evaluated encoders, 2^{12} buckets represent a good trade-off. In Appendix F, we
275 present an ablation study on the effect of the number of total buckets.

276 Our evaluation of the LSH to track coverage of the embedding space is presented in Figure 7a.
277 We observe that representations returned for standard downstream tasks (FashionMNIST, SVHN,
278 CIFAR10) occupy a significantly smaller fraction of the total number of buckets than complex data

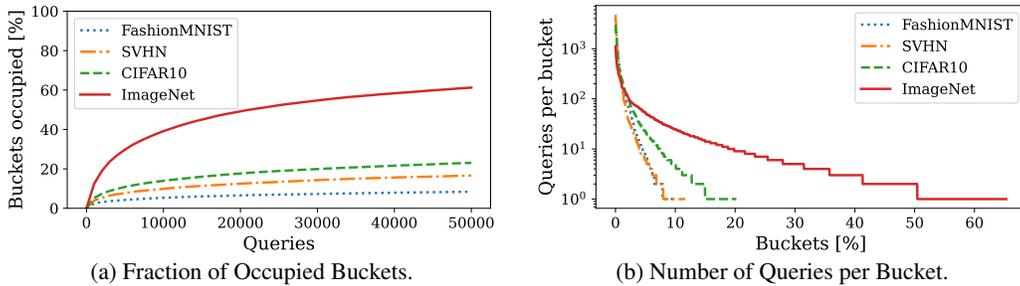


Figure 3: **Estimating Embedding Space Coverage through LSH on SimSiam Encoder.** We present the fraction of buckets occupied by representations of different datasets as a function of the number of queries posed to the encoder (*left*). We observe that representations for the downstream datasets (FashionMNIST, SVHN, CIFAR10) occupy a smaller fraction of buckets than representations from the complex ImageNet dataset. Our evaluation of the number of queries whose representations are mapped to the same bucket (*right*) indicates that our total number of buckets (2^{12}) is well calibrated for the estimation of covered representation space: over all datasets, we experience hash collisions, *i.e.*, queries whose representations are mapped to the same buckets. This indicates that our LSH is capable of representing similarities in the representations.

279 from multiple distributions (ImageNet). We further observe that the fraction of buckets occupied by
 280 the representations saturates over time. These findings highlight that LSH is successful in capturing
 281 the differences between legitimate users and adversaries—even in a low-query regime. Finally, we
 282 note that our total number of buckets (2^{12}) is well calibrated since, over all datasets, it successfully
 283 maps multiple representations to the same hash bucket while still filling various fractions of the total
 284 number of buckets.

285 4.2 Calibrating the Cost Function

286 We experiment with different sets of hyperparameters to instantiate the cost function from Equation (1)
 287 in the previous section (3.3). As described there, we can calibrate the function such that a desired penalty (in
 288 the form of a specific σ) will be assigned at a certain fraction of buckets occupied. For B4B, we aim
 289 at penalizing high embedding space coverage severely. Therefore, we need to identify and optimize for two
 290 components: 1) which value of σ leads to significant performance drops, and 2) for what fraction of coverage
 291 do we want to impose this significant drop. We base
 292 both components on empirical observations. Our first observation is that for our four downstream
 293 tasks (FashionMNIST, SVHN, STL10, and CIFAR10), performance drops to 10% (*i.e.*, random
 294 guessing) at roughly $\sigma = 0.5$. In Figure 7a, we further see that with 50k queries, the downstream
 295 tasks occupy $< 30\%$ of the buckets. Ultimately, setting α and β are design choices that an API
 296 provider needs to make in order to specify what type of query behavior they want to penalize. As
 297 very loose bounds (weak defense), based on our observation, we consider $\sigma = 1$ as a high penalty,
 298 which leads to $\alpha = 1$, and select $\beta = 0.8$. This β corresponds to considering 80% of buckets filled
 299 as a too-large coverage of the embedding space. We empirically observe that coverage of 80% of
 300 buckets occurs, for example, after around 100k of ImageNet queries. By choosing our target β so
 301 loose, *i.e.*, significantly larger than the observed 30% for downstream tasks, we offer flexibility for
 302 the API to also provide good representations for more complex downstream tasks. Finally, to obtain a
 303 flat cost curve close to the origin—which serves to map small fractions of covered buckets to small
 304 costs—we find that we can set $\lambda = 10^{-6}$. In the Appendix, we evaluate our defense end-to-end with
 305 differently parameterized cost functions.

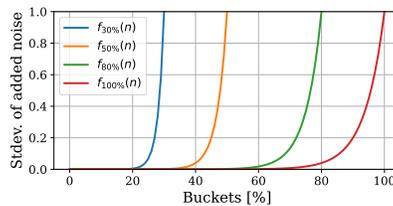


Figure 4: **Cost Function Calibration.**

Our first observation is that for our four downstream tasks (FashionMNIST, SVHN, STL10, and CIFAR10), performance drops to 10% (*i.e.*, random guessing) at roughly $\sigma = 0.5$. In Figure 7a, we further see that with 50k queries, the downstream tasks occupy $< 30\%$ of the buckets. Ultimately, setting α and β are design choices that an API provider needs to make in order to specify what type of query behavior they want to penalize. As very loose bounds (weak defense), based on our observation, we consider $\sigma = 1$ as a high penalty, which leads to $\alpha = 1$, and select $\beta = 0.8$. This β corresponds to considering 80% of buckets filled as a too-large coverage of the embedding space. We empirically observe that coverage of 80% of buckets occurs, for example, after around 100k of ImageNet queries. By choosing our target β so loose, *i.e.*, significantly larger than the observed 30% for downstream tasks, we offer flexibility for the API to also provide good representations for more complex downstream tasks. Finally, to obtain a flat cost curve close to the origin—which serves to map small fractions of covered buckets to small costs—we find that we can set $\lambda = 10^{-6}$. In the Appendix, we evaluate our defense end-to-end with differently parameterized cost functions.

311 4.3 Assessing the Effect of Transformations

312 **Transformations Do Not Harm Utility for Legitimate Users.** We evaluate the downstream
 313 accuracy for transformed representations based on training a linear classifier on top of them. To
 314 separate the effect of the noise added by our defense from the effect of the transformations, we
 315 perform the experiments in this subsection without adding noise to the returned representations.
 316 For example, on the CIFAR10 dataset and a SimSiam encoder pre-trained on ImageNet, without
 317 any transformations applied, we obtain a downstream accuracy of 90.41% (± 0.02), while, with
 318 transformations, we obtain 90.24% (± 0.11) for Affine, 90.40% (± 0.05) for Pad+Shuffle, 90.18% (\pm
 319 0.06) for Affine+Pad+Shuffle, and 88.78% (± 0.2) for Binary. This highlights that the transformations
 320 preserve utility for legitimate users. This holds over all datasets we evaluate as we show in Appendix F.

321 **Adversaries Cannot Perfectly Remap Representations over Multiple Sybil Accounts.** To under-
 322 stand the impact of our per-user account transformations on sybil-attack based encoder stealing, we
 323 evaluate the difficulty of remapping representations between different sybil accounts. For simplicity,
 324 and since we established in Section 3.4 that multi-account attacks reduce to a two-account setup, we
 325 assume an adversary who queries from two sybil accounts and aims at learning to map the transformed
 326 representations from account #2 to the representation space of account #1. Using more accounts for
 327 the adversary causes a larger query overhead and potentially more performance loss from remapping.
 328 Our evaluation here, hence, represents a lower bound on the overhead caused to the adversary through
 329 our transformations.

330 We learn the mapping between different accounts’ representations by training a linear model on
 331 overlapping representations between the accounts. We assess the fidelity of remapped represen-
 332 tations as a function of the number of overlapping queries between the accounts. As a fidelity
 333 metric for our remapping, we compare the cosine distance between representations (a and b de-
 334 fined as: $1 - \frac{a^T b}{\|a\|_2 \cdot \|b\|_2}$). Once the remapping is trained, we evaluate by querying 10k data
 335 points from the test dataset through account #1 and then again through account #2. Then, we
 336 apply the learned remapping to the latter one and compute the pairwise cosine distances be-
 337 tween the representations from account #1 and their remapped counterparts from account #2.
 338 Our results are depicted in Figure 5. We show that the largest
 339 cosine distance is achieved with the binary transformations, making them the most protective against the adversary since
 340 they best prevent perfect remapping, even with an overlap of as many as 10k queries between both accounts. However, these
 341 binary transformations also incur the highest drop in accuracy for legitimate users. The defender has the possibility of selecting
 342 their preferred types of transformations between representations taking into account the trade-offs between the effectiveness of
 343 the defense and the negative impact on legitimate users. The defender has the possibility of selecting
 344 their preferred types of transformations between representations taking into account the trade-offs between the effectiveness of
 345 the defense and the negative impact on legitimate users. The defender has the possibility of selecting
 346 their preferred types of transformations between representations taking into account the trade-offs between the effectiveness of
 347 the defense and the negative impact on legitimate users.

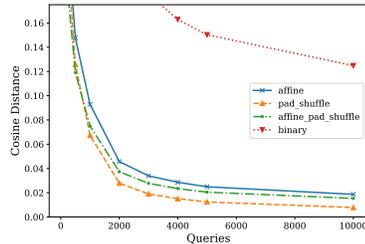


Figure 5: Quality of Remappings.

348 4.4 End-to-End Stealing of an Encoder under our Defense

349 We perform an end-to-end study to showcase how our B4B defense affects legitimate users vs
 350 adversaries. The hyperparameters for B4B are chosen according to the empirical evaluation of the
 351 previous sections with 2^{12} as the number of buckets, $\alpha = 1$, $\beta = 0.8$, $\lambda = 10^{-6}$ as the hyperparameter
 352 of the cost function, and different random affine transformations per-user account. Our main results
 353 are presented in Table 1. We observe that instantiating our framework with B4B has a negligible
 354 impact on legitimate users while substantially lowering the performance of stolen encoders in the
 355 case of single-user and sybil attackers.

356 **Legitimate Users.** We compare the accuracy of downstream classifiers trained on top of unprotected
 357 vs defended encoders. The victim encoder achieves high accuracy on the downstream tasks when no
 358 defense is employed. With B4B in place, we observe that across all the downstream tasks, the drop in
 359 performance is below 1%. For example, there is only a slight decrease in the accuracy of CIFAR10
 360 from $90.41 \pm 0.02\%$ to $90.24 \pm 0.11\%$. B4B’s small effect on legitimate users stems from the fact
 361 that their downstream representations cover a relatively small part of the representations space. This
 362 results in a very low amount of noise added to their representations which preserves performance.

Table 1: **Stealing and Using Encoders With and Without our Defense.** The *USER* column represents the type of the APIs’ user, where Legit denotes a legitimate user, Attack stands for a standard single-account adversary, and Sybil represents an adversary using two sybil accounts. We use InfoNCE loss for encoder extraction. # Queries stands for the number of queries used for stealing. The *TYPE* column expresses how the dataset is used. We follow the stealing setup from [16]. In the first row, we present the undefended victim encoder’s performance as the accuracy for downstream tasks trained on the encoder’s returned representations. In the following four rows, we show downstream utility for legitimate users when the victim encoder is defended by our B4B. Finally, (in the remaining rows) we assess the performance of stolen encoders on the downstream tasks. Our results highlight that while the performance of the encoder for legitimate users stays high, our B4B renders stealing inefficient with the stolen encoders obtaining significantly worse performance on downstream tasks.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	90.41±0.02	95.08±0.13	75.47±0.04	91.22±0.11
LEGIT	B4B	50K	CIFAR10	QUERY	90.24±0.11	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	95.05±0.1	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	74.96±0.13	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	91.7±0.01
ATTACK	NONE	50K	IMAGENET	STEAL	65.2 ±0.03	64.9 ±0.01	62.1 ±0.01	88.5 ±0.01
ATTACK	B4B	50K	IMAGENET	STEAL	35.72±0.04	31.54±0.02	19.74±0.02	70.01±0.01
ATTACK	NONE	100K	IMAGENET	STEAL	68.1 ±0.03	63.1 ±0.01	61.5 ±0.01	89.0 ±0.07
ATTACK	B4B	100K	IMAGENET	STEAL	12.01±0.07	13.94±0.05	19.96±0.03	69.63±0.07
SYBIL	B4B	50K+50K	IMAGENET	STEAL	39.56±0.06	38.50±0.04	23.41±0.02	77.01±0.08

363 **Adversaries.** For adversaries who create a stolen copy of the victim encoder, we make two main
364 observations. The most crucial one is that when our B4B is in place, the performance of the stolen
365 copies over all downstream tasks significantly drops in comparison to when the victim encoder is
366 unprotected (grey rows in Table 1). This highlights that our B4B effectively prevents stealing. Our next
367 key observation concerns the number of stealing queries used by the adversary: When no defense is
368 applied, the more queries are issued against the API (*e.g.*, 100k instead of 50k), the higher performance
369 of the stolen encoder on downstream tasks (*e.g.*, CIFAR10 or FashionMNIST). In contrast, with B4B
370 implemented as a defense, the performance decreases when using more stealing queries from a single
371 account. This is because with more queries issued, the coverage of embedding space grows which
372 renders the returned representations increasingly noisy and harms stealing performance. We also show
373 that this performance drop cannot be prevented by sybil attacks. Therefore, we consider an adversary
374 who queries from two sybil accounts with 50k queries issued per account and the first 10k queries of
375 both accounts used to learn the remapping of representations between them. When the adversary trains
376 their stolen encoder copy on all the remapped representations, they increase downstream performance
377 over querying from a single account. Yet, their performance is still significantly smaller than the
378 performance of the victim encoder for legitimate users, or the encoder stolen from an undefended
379 victim. This highlights that our B4B also successfully prevents sybil attacks.

380 5 Conclusions

381 We design B4B a new and modular active defense framework against stealing SSL encoders. All the
382 previous approaches were either reactive, acting after the attack happened to detect stolen encoders, or
383 lowered the quality of outputs substantially also for legitimate users which rendered such mechanisms
384 impractical. We show that B4B successfully distinguishes between legitimate users and adversaries
385 by tracking the embedding space coverage of users’ obtained representations. B4B then leverages this
386 tracking to apply a cost function that penalizes users based on the current space coverage, for instance,
387 by lowering the quality of their outputs. Finally, B4B prevents sybil attacks by implementing per-user
388 transformations for the returned representations. Through our experimental evaluation, we show
389 that our defense indeed renders encoder stealing inefficient while preserving downstream utility for
390 legitimate users. Our B4B is therefore a valuable contribution to a safer sharing and democratization
391 of high-utility encoders over public APIs.

392 References

- 393 [1] Clarifai, <https://www.clarifai.com/>. URL <https://www.clarifai.com/>.
- 394 [2] Cohere, <https://cohere.ai>. URL <https://cohere.ai/>.
- 395 [3] Openai, <https://openai.com>. URL <https://openai.com/>.
- 396 [4] Poet2, <https://labs.hyperledger.org/labs/archived/sawtooth-poet2.html>. URL <https://labs.hyperledger.org/labs/archived/sawtooth-poet2.html>.
- 397
- 398 [5] Google, <https://google.github.io/snappy/>. URL <https://google.github.io/snappy/>.
- 399 [6] Facebook zstd, <https://github.com/facebook/zstd>. URL <https://github.com/facebook/zstd>.
- 400
- 401 [7] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your
402 weakness into a strength: Watermarking deep neural networks by backdooring. In *27th USENIX
403 Security Symposium (USENIX Security 18)*, pages 1615–1631, 2018.
- 404 [8] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *USENIX Security Symposium*, volume 6, 2021.
- 405
- 406
- 407 [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9650–9660, October 2021.
- 408
- 409
- 410
- 411 [10] Ting Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. *International Conference on Machine Learning*, 2020.
- 412
- 413 [11] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. 2020.
- 414 [12] Adam Coates, Andrew Ng, and Honglak Lee. An Analysis of Single Layer Networks in Unsupervised Feature Learning. In *AISTATS*, 2011. https://cs.stanford.edu/~acoates/papers/coatesleeng_aistats_2011.pdf.
- 415
- 416
- 417 [13] Tianshuo Cong, Xinlei He, and Yang Zhang. Sslguard: A watermarking scheme for self-supervised learning pre-trained encoders. *CoRR*, abs/2201.11692, 2022. URL <https://arxiv.org/abs/2201.11692>.
- 418
- 419
- 420 [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- 421
- 422
- 423 [15] Adam Dziedziec, Nikita Dhawan, Muhammad Ahmad Kaleem, Jonas Guan, and Nicolas Papernot. On the difficulty of defending self-supervised learning against model extraction. In *International Conference on Machine Learning*, 2022.
- 424
- 425
- 426 [16] Adam Dziedziec, Haonan Duan, Muhammad Ahmad Kaleem, Nikita Dhawan, Jonas Guan, Yannis Cattan, Franziska Boenisch, and Nicolas Papernot. Dataset inference for self-supervised models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=CCBJf9xJo2X>.
- 427
- 428
- 429
- 430
- 431 [17] Adam Dziedziec, Muhammad Ahmad Kaleem, Yu Shen Lu, and Nicolas Papernot. Increasing the cost of model extraction with calibrated proof of work. In *International Conference on Learning Representations*, 2022. URL <https://arxiv.org/abs/2201.09243>.
- 432
- 433
- 434 [18] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. 2013. URL <https://eprint.iacr.org/2013/796>. <https://eprint.iacr.org/2013/796>.
- 435
- 436

- 437 [19] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of
438 space. In *Advances in Cryptology—CRYPTO 2015: 35th Annual Cryptology Conference, Santa*
439 *Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, pages 585–605. Springer, 2015.
- 440 [20] Nicholas Frosst, Nicolas Papernot, and Geoffrey Hinton. Analyzing and improving representa-
441 tions with the soft nearest neighbor loss. In Kamalika Chaudhuri and Ruslan Salakhutdinov,
442 editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of
443 *Proceedings of Machine Learning Research*, pages 2012–2020. PMLR, 09–15 Jun 2019. URL
444 <https://proceedings.mlr.press/v97/frosst19a.html>.
- 445 [21] Jean-Bastien Grill, F. Strub, F. Althé, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch,
446 B. A. Pires, Z. D. Guo, M. G. Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko.
447 Bootstrap your own latent: A new approach to self-supervised learning. *Computer Vision and*
448 *Pattern Recognition*, 2020.
- 449 [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked
450 autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on*
451 *Computer Vision and Pattern Recognition (CVPR)*, pages 16000–16009, June 2022.
- 452 [23] Matthew Jagielski, N. Carlini, D. Berthelot, A. Kurakin, and N. Papernot. High accuracy and
453 high fidelity extraction of neural networks. *USENIX Security Symposium*, 2020.
- 454 [24] Hengrui Jia, Christopher A Choquette-Choo, Varun Chandrasekaran, and Nicolas Papernot. En-
455 tangled watermarks as a defense against model extraction. In *30th USENIX Security Symposium*
456 *(USENIX Security 21)*, pages 1937–1954, 2021.
- 457 [25] Mika Juuti, Sebastian Szyller, Samuel Marchal, and N Asokan. Prada: protecting against dnn
458 model stealing attacks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*,
459 pages 512–527. IEEE, 2019.
- 460 [26] Sanjay Kariyappa, Atul Prakash, and Moinuddin K Qureshi. Maze: Data-free model stealing
461 attack using zeroth-order gradient estimation. In *Proceedings of the IEEE/CVF Conference on*
462 *Computer Vision and Pattern Recognition*, pages 13814–13823, 2021.
- 463 [27] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- 464 [28] Yupei Liu, Jinyuan Jia, Hongbin Liu, and Neil Zhenqiang Gong. Stolenencoder: Stealing
465 pre-trained encoders in self-supervised learning. In *Proceedings of the 2022 ACM SIGSAC*
466 *Conference on Computer and Communications Security, CCS ’22*, page 2115–2128, New
467 York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450394505. doi:
468 10.1145/3548606.3560586. URL <https://doi.org/10.1145/3548606.3560586>.
- 469 [29] Pratyush Maini, Mohammad Yaghini, and Nicolas Papernot. Dataset inference: Ownership
470 resolution in machine learning. In *Proceedings of ICLR 2021: 9th International Conference on*
471 *Learning Representations*, 2021.
- 472 [30] Yuki Nagai, Y. Uchida, S. Sakazawa, and Shin’ichi Satoh. Digital watermarking for deep neural
473 networks. *International Journal of Multimedia Information Retrieval*, 7:3–16, 2018.
- 474 [31] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng.
475 Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep*
476 *Learning and Unsupervised Feature Learning 2011*, 2011. URL http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf.
- 477 [32] Tribhuvanesh Orekondy, Bernt Schiele, and Mario Fritz. Knockoff nets: Stealing functionality
478 of black-box models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
479 *Pattern Recognition*, pages 4954–4963, 2019.
- 480 [33] Loïc Paulevé, Hervé Jégou, and Laurent Amsaleg. Locality sensitive hashing: A comparison of
481 hash function types and querying mechanisms. *Pattern recognition letters*, 31(11):1348–1358,
482 2010.
483

- 484 [34] Zeyang Sha, Xinlei He, Ning Yu, Michael Backes, and Yang Zhang. Can't steal? cont-steal!
485 contrastive stealing attacks against image encoders. 2022. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2201.07513)
486 [2201.07513](https://arxiv.org/abs/2201.07513).
- 487 [35] Yun Shen, Xinlei He, Yufei Han, and Yang Zhang. Model stealing attacks against inductive graph
488 neural networks. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1175–1192.
489 IEEE, 2022.
- 490 [36] Ram Shankar Siva Kumar, Magnus Nyström, John Lambert, Andrew Marshall, Mario Goertzel,
491 Andi Comissoneru, Matt Swann, and Sharon Xia. Adversarial machine learning-industry
492 perspectives. In *2020 IEEE Security and Privacy Workshops (SPW)*, pages 69–75, 2020. doi:
493 10.1109/SPW50608.2020.00028.
- 494 [37] Malcolm Slaney and Michael Casey. Locality-sensitive hashing for finding nearest neighbors
495 [lecture notes]. *IEEE Signal processing magazine*, 25(2):128–131, 2008.
- 496 [38] Florian Tramèr, F. Zhang, A. Juels, M. Reiter, and T. Ristenpart. Stealing machine learning
497 models via prediction apis. *USENIX Security Symposium*, 2016.
- 498 [39] Jean-Baptiste Truong, Pratyush Maini, Robert J. Walls, and Nicolas Papernot. Data-free model
499 extraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
500 *Recognition (CVPR)*, June 2021.
- 501 [40] Yusuke Uchida, Yuki Nagai, Shigeyuki Sakazawa, and Shin'ichi Satoh. Embedding watermarks
502 into deep neural networks. In *Proceedings of the 2017 ACM on international conference on*
503 *multimedia retrieval*, pages 269–277, 2017.
- 504 [41] Yutong Wu, Han Qiu, Tianwei Zhang, Lin Jiwei, and Meikang Qiu. Watermarking pre-trained
505 encoders in contrastive learning. *ArXiv*, abs/2201.08217, 2022.
- 506 [42] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-
507 supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021.

508 **A Broader Impacts**

509 The goal of our work is to actively defend self-supervised encoders against model stealing attacks.
510 Since we are directly defending encoders, any negative societal impacts of our work are minimal. One
511 potentially negative impact could be the degradation of performance for legitimate users. However,
512 as shown in our experimental results, we are able to preserve high utility for standard users.

513 **B Limitations**

514 We show how our defense method is tuned for SimSiam and DINO. There are more types of SSL
515 encoders that can be tested with our method. The B4B defense method requires tuning the parameters,
516 such as the number of occupied buckets that is allowed without any penalty for the cost function, or
517 the selection of the transformations. These steps are rather difficult to automate but can be replaced
518 with more data-driven approaches. For example, instead of designing a cost function from scratch,
519 we could create an ML model to obtain a cost for a given occupation of the representation space. We
520 explain more details in the Appendix C.2.

521 **C Alternative Building Blocks to Instantiate B4B**

522 While we present a reference implementation of B4B in our work that instantiates the three building
523 blocks with (1) Local Sensitive Hashing, (2) Utility of the Representations, and (3) a set of concrete
524 transformations, there exists a multitude of alternatives to concretely implement our B4B framework.
525 In the following, we present these alternatives, grouped by building block.

526 **C.1 Alternative Estimation of the Coverage of Embedding Space**

527 We also explore alternative methods to measure the distances between representations for queries
528 sent to an API. One of them is to apply the cosine distance (where for two representations a and b , it
529 is defined as: $1 - \frac{a^T b}{\|a\|_2 \cdot \|b\|_2}$) since it can be measured between individual data points in a pair-wise
530 fashion. If the total pair-wise cosine distance between representations for a given user is small, then
531 the user queries presumably come from a single downstream task distribution. Otherwise, a user
532 might be malicious and would like to cover a large part of the representation space, then the total
533 pair-wise cosine distance for the user’s representations would be high. Note that in this case, the
534 cosine distance can be replaced with any other distance measure, such as the Minkowski distance.
535 We opt for the LSH in our reference implementation, since it is much less expensive to compute than
536 cosine distance. LSH requires only $2^{12} = 4096$ buckets that can be expressed as a binary table with
537 the same number of elements, which requires in the worst case iterating over all of them to count how
538 many are occupied. With more than 4096 queries sent by a given user, the computation on the LSH is
539 sublinear $< O(n)$ with respect to the number of user queries. For the cosine distance approach, the
540 required computation grows quadratically $O(n^2)$ with the number of queries.

541 **C.2 Alternative Cost Functions**

542 The cost functions can be designed from scratch manually or learned, for example, via an ML model,
543 such as a neural network or SVM. In our initial version, the function was designed manually, where
544 the underlying premise is that once a specified number of buckets is occupied, the cost should grow
545 exponentially. Instead of defining such a function or providing the high-level parameters for functions
546 that we contributed, one could learn an ML model that for a given number of buckets occupied, it
547 should output an estimated cost, or even directly, the desired σ (standard deviation) of the noise added
548 for representations. This method requires a relatively large number of data points to be provided for
549 training the model, however, lowers the burden on a defender to either decide on the specific function
550 or adjust its parameters. Thus, it could be more user-friendly, for example, not necessitating any
551 mathematical background, but can be precise enough to obtain the desired behavior.

552 Note that instead of adding the calibrated noise (proportional to the estimated cost) to the represen-
553 tations, we could rather require a given user to pay a higher monetary cost for queries that cover a
554 large fraction of the representation space, or force a user to solve a puzzle in a form of the proof-
555 of-work [17], wait a specified amount of time via proof-of-elapsed time (PoET) [4], or prove that

556 a specified amount of disk space was reserved [18, 19]. For example, consider the approach with
 557 PoET. A user sends queries to the API, which we cost based on their occupation of the embedding
 558 space. The user is sent a waiting time. The users' resource (e.g., a CPU) has to be occupied for this
 559 specific waiting time without performing any work. At the end of the specified amount of time, the
 560 user sends proof of the elapsed time, which can be easily verified by the server. PoET requires access
 561 to specialized hardware, for example, secure CPU instructions that are becoming widely available in
 562 consumer and enterprise processors. If a user does not own such hardware, proof of elapsed time
 563 could be produced using a service exposed by a cloud provider (e.g., Azure VMs that feature TEE 2).
 564 Note that if a server sends the time based on the calculated cost, the adversary might learn the cost
 565 function. Instead, the exact waiting time should be split in random *subwaiting* times and sent to the
 566 user one by one. Thus, a server should rather have a few rounds of exchange with the client to incur
 567 the additional cost.

568 C.3 Alternative to Transformations

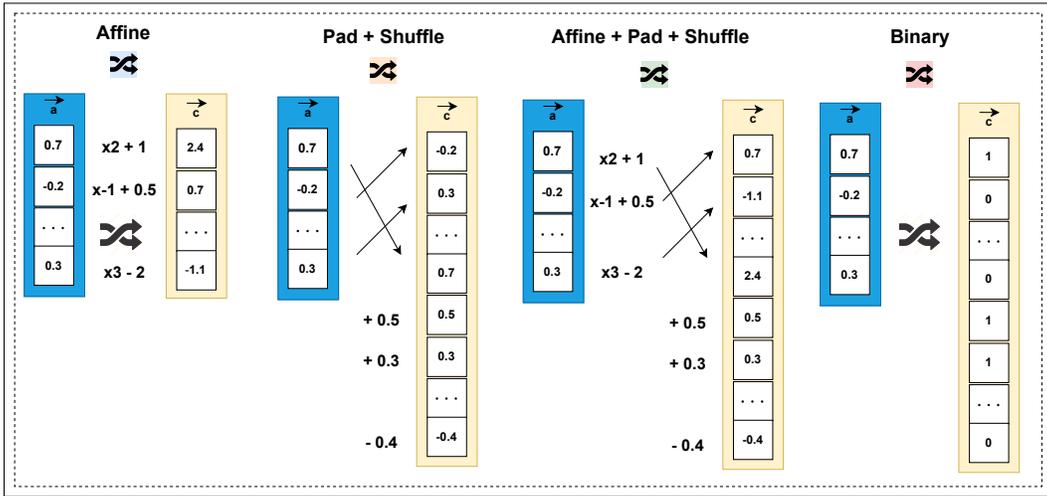


Figure 6: **Overview on Transformations.** We depict the inner-workings of the transformations considered in this work.

569 As an alternative to the transformations used within this work (see Figure 6), one could use a different
 570 set of transformations or combinations thereof. The padding can be done with different constant
 571 values and combined with adding constant values within the representations. The padding and adding
 572 the constant values can be followed by shuffling the elements within the representations. We can
 573 apply the affine of binary transformations on top of the padding and shuffling. Additionally, we can
 574 also use other pre-defined linear transformations like rotations or shearing.

575 The representations could also be compressed to smaller vectors and the compression rate would
 576 depend on the occupation of the representation space, for example, the higher the number of occupied
 577 buckets in our hash table, the more compressed the output representations could be. Such representa-
 578 tions could be compressed via FFT, a cosine transform, or standard compression techniques such
 579 as snappy [5]. If the information from the representations should not be lost, then the lossless com-
 580 pression techniques can be applied, for instance, zstd [6]. The only requirement of the compression
 581 techniques is to ensure that they do not decrease the accuracy on downstream tasks for legitimate
 582 users.

583 Another alternative is to incorporate an additional neural network layer for transforming the returned
 584 representations. The training of this supplementary layer should primarily focus on preserving
 585 the usability of the representations for legitimate users. This approach grants the API provider
 586 with additional capabilities, as it allows for the utilization of customized training objectives. For
 587 instance, if the API provider employs LSH (Locality-Sensitive Hashing) to estimate the coverage
 588 of the representation space, they can leverage buckets and train the additional layer to maintain
 589 high-quality representations exclusively for frequently-used buckets and their surrounding areas,
 590 while not prioritizing the rest of the representation space. This approach safeguards legitimate users

591 from any adverse effects, as their coverage of the representation space is minimal. Simultaneously, it
592 ensures that adversaries are unable to exploit representations from the entire representation space.

593 **D Sybil Attacks**

594 We consider an adversary who generates n sybil accounts to steal the encoder from the API. For each
595 of the accounts, the representations are transformed in a different way. Therefore, to replicate the
596 victim model using all the obtained representations, the adversary has to map these representations
597 into one single space. This can be done, for example, by training a neural network to perform the
598 mapping.

599 We assume the adversary obtains $\{N_1, N_2, \dots, N_n\}$ many representations from the victim for each
600 of the n sybil accounts. Without loss of generality, we assume the adversary maps them back to the
601 embedding space of the first sybil account. To learn the mapping, the adversary can apply different
602 strategies.

603 **D.1 Sybil Strategies**

604 We present three potential approaches that Sybils might want to apply to circumvent our defense. Con-
605 sider three users: A , B , and C , with their respective datasets D_A , D_B , and D_C , each with different
606 distributions to maximize extraction effectiveness. First, user A is selected to unify representations
607 from other users B and C . User A would have to query from at least two different datasets D_B and
608 D_C , while other users would act legitimately. Sybil attackers want to deploy as many users as possible
609 but with more fake accounts, user A incurs high coverage of the representation space, and this is
610 prevented by our single-user defense. In all other cases neither of the sybil users can act legitimately,
611 thus they are already affected by the single-user defense. Second, user A would query from their own
612 dataset D_A and partially from dataset D_B . Then user B would query from their own dataset D_B and
613 partially from dataset D_C , and so on. This method is the most inconspicuous but requires a number
614 of remappings that scales super-exponentially with the number of fake accounts, which is impractical.
615 Finally, each user would query from their respective dataset, for example, user B would query from
616 dataset D_B and additionally from a remapping dataset, *e.g.*, D_A . Representations could be unified
617 by mapping them to A 's representations. The last approach as well as all other cases reduce to the
618 minimum of remapping between representations of a pair of users. We show that our defense cuts
619 such attempts short by ensuring that the remapping between representations is prohibitive even for a
620 pair of users.

621 **E Additional Related Work**

622 One of the main workhorse techniques used in the encoders is contrastive learning, where the
623 representations are trained so that the positive pairs (two augmented versions of the same image)
624 have similar representations while negative pairs (augmentations of two different images) have
625 representations which are far apart.

626 **SimSiam** utilizes Siamese networks (two encoders with shared weights) but with a simplified training
627 process and architecture. In contrast to the previous frameworks, such as SimCLR [10], SimSiam's
628 authors show that negative samples are unnecessary and collapsing solutions can be avoided by
629 applying the projection head to of one of the encoders, and a stop-gradient operation to the other.
630 SimSiam minimizes the negative cosine similarity between two randomly augmented views of the
631 same image from the Siamese encoders, which is expressed via a symmetrized loss [21]. This creates
632 a simple yet highly effective representation learning method.

633 **DINO** is another popular representation learning framework. While SimSiam uses CNNs, DINO
634 employs vision transformers (ViTs). It trains a student and teacher encoder with the same architec-
635 ture, updating the teacher with an (exponential moving) average of the student. Different random
636 transformations of the same image are passed through both encoders. The student receives smaller
637 image crops, forcing it to generate representations restoring parts of the original image. The training
638 objective is minimizing cross-entropy loss between teacher and student representations.

639 **F Additional Experimental Results**

640 **F.1 Details on Experimental Setup**

641 The end-to-end experiments on stealing SimSiam and ViT DINO encoders were done using 3 A100
 642 GPUs. Detailed experiments including mapping, transformations and the evaluation was performed
 643 using a single computer equipped with two Nvidia RTX 2080 Ti GPUs.

644 **F.2 Datasets Used**

645 **CIFAR10** [27]: The CIFAR10 dataset consists of 32x32 colored images with 10 classes. There are
 646 50000 training images and 10000 test images.

647 **SVHN** [31]: The SVHN dataset contains 32x32 coloured images with 10 classes. There are roughly
 648 73000 training images, 26000 test images and 530000 "extra" images.

649 **ImageNet**[14]: Larger sized coloured images with 1000 classes. As is commonly done, we resize
 650 all images to be of size 224x224. There are approximately 1 million training images and 50000 test
 651 images.

652 **STL10** [12]: The STL10 dataset contains 96x96 coloured images with 10 classes. There are 5000
 653 training images, 8000 test images, and 100000 unlabeled images.

654 **F.3 More Results for the End2End Empirical Evaluation**

655 We consider fine-tuning parameters β , λ , and α for our cost function and the intuitive meaning
 656 behind these parameters. In general, our recommendation is to adjust the parameter β that specifies
 657 how many buckets are allowed to be filled by users' downstream tasks. On the other hand, when
 658 parameter λ is increased, this causes a higher amount of added noise before we reach the number of
 659 buckets specified by β , which lowers the performance of a given downstream task relatively early.
 660 For example, a higher value of λ in Figure 4, would cause an increase in the amount of added noise
 661 much earlier than for the target value of β . Finally, parameter α controls the amount of noise once
 662 the number of buckets specified by β is reached. Thus, in Figure 4, we set $\alpha = 1$ and the maximum
 663 standard deviation of the added Gaussian noise is 1.

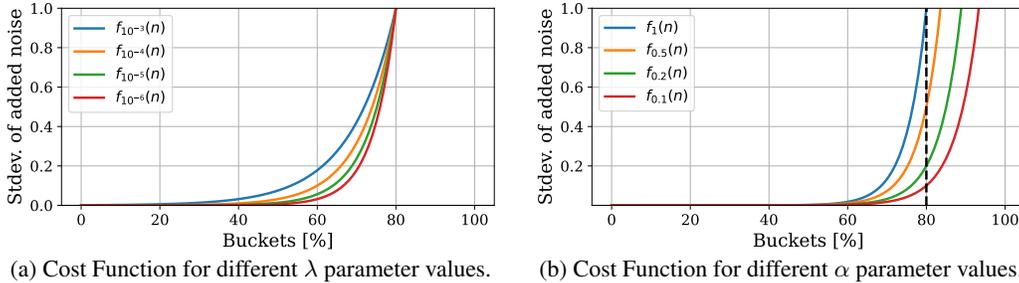


Figure 7: **Effects of λ and α parameters on the Cost Function.** We present the Cost Function for $\alpha=1$, $\beta=80$ and different values of λ (left) and $\lambda = 10^{-6}$, $\beta=80$ and different values of α (right).

Table 2: **Stealing and Using Encoders With and Without our Defense.** The model used in the experiments is Simsiam, with the following parameters for the cost function $\lambda = 10^{-4}$, $\alpha = 1$, and $\beta = 80\%$, and the number of buckets equal to 2^{12} . Due to the higher value of the parameter λ , we observe lower performance on downstream tasks for the attackers since the magnitude of noise added to the representations is higher. However, for more complicated tasks than CIFAR10, this change might cause a potential drop in accuracy for the legitimate users.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	90.41 \pm 0.02	95.08 \pm 0.13	75.47 \pm 0.04	91.22 \pm 0.11
LEGIT	B4B	50K	CIFAR10	QUERY	90.02 \pm 0.1	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	94.88 \pm 0.17	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	74.72 \pm 0.13	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	91.76 \pm 0.09
ATTACK	NONE	50K	IMAGENET	STEAL	65.2 \pm 0.03	64.9 \pm 0.01	62.1 \pm 0.01	88.5 \pm 0.01
ATTACK	B4B	50K	IMAGENET	STEAL	28.22 \pm 0.04	26.62 \pm 0.02	19.62 \pm 0.02	78.41 \pm 0.01
ATTACK	NONE	100K	IMAGENET	STEAL	68.1 \pm 0.03	63.1 \pm 0.01	61.5 \pm 0.01	89.0 \pm 0.07
ATTACK	B4B	100K	IMAGENET	STEAL	17.73 \pm 0.18	15.59 \pm 0.61	19.53 \pm 0.01	55.11 \pm 0.05
SYBIL	B4B	50K+50K	IMAGENET	STEAL	33.43 \pm 0.03	31.18 \pm 0.12	22.91 \pm 0.01	75.35 \pm 0.05

Table 3: **Stealing and Using Encoders With and Without our Defense.** The model used in the experiments is Simsiam, with the following parameters for the cost function $\lambda = 10^{-6}$, $\alpha = 1$, and $\beta = 50\%$, and the number of buckets equal to 2^{12} . This experiment corresponds to considering 50% of buckets filled as a too-large coverage of the embedding space. This improves the defense but again might potentially harm the performance of more complicated tasks than CIFAR10 since they could occupy more buckets than 50%.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	90.41 \pm 0.02	95.08 \pm 0.13	75.47 \pm 0.04	91.22 \pm 0.11
LEGIT	B4B	50K	CIFAR10	QUERY	90.27 \pm 0.07	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	95.12 \pm 0.13	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	74.94 \pm 0.16	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	91.66 \pm 0.05
ATTACK	NONE	50K	IMAGENET	STEAL	65.2 \pm 0.03	64.9 \pm 0.01	62.1 \pm 0.01	88.5 \pm 0.01
ATTACK	B4B	50K	IMAGENET	STEAL	15.52 \pm 0.37	12.57 \pm 0.23	19.53 \pm 0.01	23.17 \pm 0.01
ATTACK	NONE	100K	IMAGENET	STEAL	68.1 \pm 0.03	63.1 \pm 0.01	61.5 \pm 0.01	89.0 \pm 0.07
ATTACK	B4B	100K	IMAGENET	STEAL	16.27 \pm 0.04	13.93 \pm 0.35	19.54 \pm 0.02	54.69 \pm 0.02
SYBIL	B4B	50K+50K	IMAGENET	STEAL	30.14 \pm 0.01	29.57 \pm 0.08	19.99 \pm 0.03	71.72 \pm 0.01

Table 4: **Stealing and Using Encoders With and Without our Defense.** The model used in the experiments is Simsiam, with the following parameters for the cost function $\lambda = 10^{-6}$, $\alpha = 1$, and $\beta = 30\%$, and the number of buckets equal to 2^{12} . Since the value of parameter β is decreased substantially to 30%, we observe a drop in accuracy for legitimate users. For example, more than 1% for CIFAR10. In the next Table 5, we show that by also decreasing the parameter α , we can attenuate this harmful effect and retain higher accuracy for legitimate users. In case of an attack, for 100k stealing queries, we observe much lower accuracy levels than for $\beta = 50\%$ shown in Table 3.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	90.41 \pm 0.02	95.08 \pm 0.13	75.47 \pm 0.04	91.22 \pm 0.11
LEGIT	B4B	50K	CIFAR10	QUERY	88.1 \pm 0.11	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	94.92 \pm 0.11	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	74.37 \pm 0.02	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	91.67 \pm 0.07
ATTACK	NONE	50K	IMAGENET	STEAL	65.2 \pm 0.03	64.9 \pm 0.01	62.1 \pm 0.01	88.5 \pm 0.01
ATTACK	B4B	50K	IMAGENET	STEAL	30.82 \pm 0.09	26.37 \pm 0.07	21.87 \pm 0.03	66.0 \pm 0.02
ATTACK	NONE	100K	IMAGENET	STEAL	68.1 \pm 0.03	63.1 \pm 0.01	61.5 \pm 0.01	89.0 \pm 0.07
ATTACK	B4B	100K	IMAGENET	STEAL	9.57 \pm 0.17	9.83 \pm 0.09	19.57 \pm 0.01	27.06 \pm 0.46
SYBIL	B4B	50K+50K	IMAGENET	STEAL	29.15 \pm 0.02	28.67 \pm 0.06	19.98 \pm 0.03	70.62 \pm 0.03

Table 5: **Stealing and Using Encoders With and Without our Defense.** The model used in the experiments is Simsiam, with the following parameters for the cost function $\lambda = 10^{-6}$, $\alpha = 0.1$, and $\beta = 30\%$, and the number of buckets equal to 2^{12} . Due to the lower performance on downstream tasks observed in Table 4 while keeping the parameter β fixed to 30% and λ fixed to 10^{-6} , we decrease the value of parameter α to 0.1, which increases the performance of legitimate users on their downstream tasks. In this experiment, we also carry out a sybil attack with more accounts (4 instead of 2), but observe that this modification does not improve the performance of the attacker. With more accounts, a sybil has to sacrifice more queries for the remappings between the representations from different accounts. Additionally, note that each account introduces a different remapping error by the dint of different transformations applied to each account by B4B.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	90.41 \pm 0.02	95.08 \pm 0.13	75.47 \pm 0.04	91.22 \pm 0.11
LEGIT	B4B	50K	CIFAR10	QUERY	90.17 \pm 0.1	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	94.92 \pm 0.09	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	74.97 \pm 0.13	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	91.71 \pm 0.08
ATTACK	NONE	50K	IMAGENET	STEAL	65.2 \pm 0.03	64.9 \pm 0.01	62.1 \pm 0.01	88.5 \pm 0.01
ATTACK	B4B	50K	IMAGENET	STEAL	19.95 \pm 0.19	15.54 \pm 0.34	19.57 \pm 0.01	23.50 \pm 0.19
ATTACK	NONE	100K	IMAGENET	STEAL	68.1 \pm 0.03	63.1 \pm 0.01	61.5 \pm 0.01	89.0 \pm 0.07
ATTACK	B4B	100K	IMAGENET	STEAL	10.35 \pm 0.19	12.37 \pm 0.69	19.34 \pm 0.01	68.93 \pm 0.17
SYBIL	B4B	4 \times 25K	IMAGENET	STEAL	33.15 \pm 0.04	30.23 \pm 0.07	20.87 \pm 0.01	72.19 \pm 0.02

664 **F.4 Setting the number of buckets**

665 We present our procedure to find an optimal number of buckets in Figure 8.

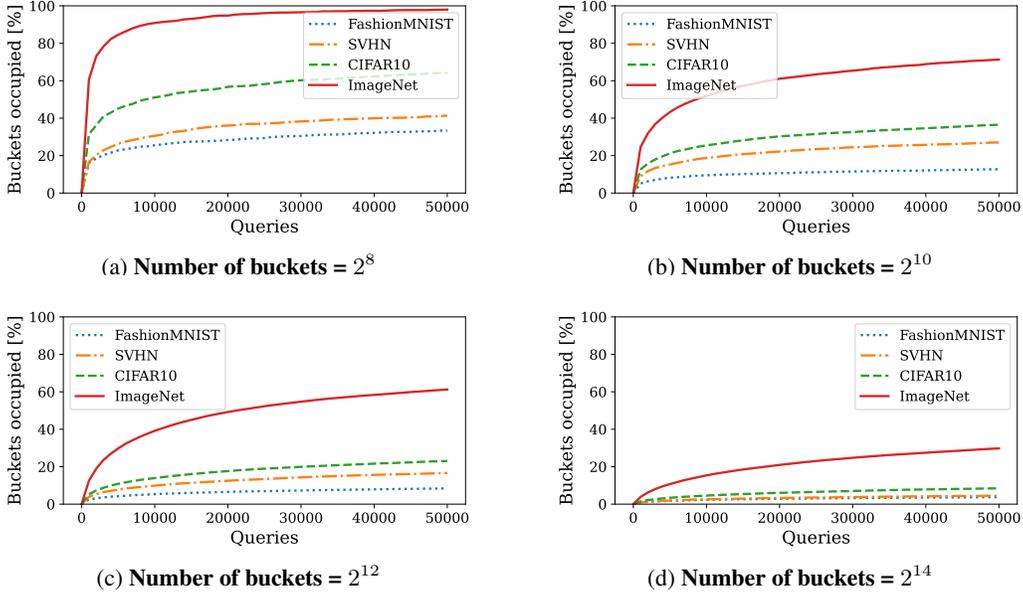


Figure 8: **Estimating Embedding Space Coverage through LSH on SimSiam Encoder.** We extend the results from Figure 7a(a) and present the fraction of buckets occupied by representations of different datasets as a function of the number of queries posed to the encoder. We consider different number of buckets in the LSH table. We observe that 2^8 buckets is too small since queries from the ImageNet dataset saturate all the buckets after around 50k queries, while the number 2^{14} of buckets is too large since it is never occupied more than 40%. Thus, the number 2^{12} buckets is a good middle ground. Subfigure (c) corresponds to Figure 7a from the main paper. We also use the same notation and carry out our experiments in the same way as in Figure 7a.

666 **F.5 Results DINO**

667 We show that our defense is also applicable to the DINO encoder. The occupation of the representations space is presented visually in Figure 9. We also show that the number of buckets 2^{12} is optimal
 668 for DINO in Figure 10. The impact of transformation on the representations from DINO is shown
 669 Table 7. Finally, the end to end experiment for DINO is presented in Table 6.
 670

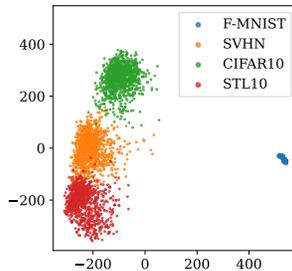


Figure 9: **Representations from Different Tasks Occupy Different Sub-Spaces of the Embedding Space. Presented for FashionMNIST, SVHN, CIFAR10, and STL10.** In this plot, we used the DINO ViT Small encoder trained on ImageNet.

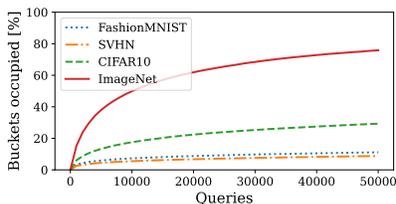


Figure 10: **Estimating Embedding Space Coverage through LSH on the DINO Encoder.** The number of buckets is set to 2^{12} . We also use the same notation and carry out our experiments in the same way as in Figure 7a.

Table 6: **Stealing and Using Encoders With and Without our Defense.** The model used in the experiments is DINO, with the following parameters for the cost function $\lambda = 10^{-6}$, $\alpha = 1000$, and $\beta = 60\%$, and the number of buckets equal to 2^{12} . We have to increase the value of parameter α by $\times 1000$ since the norms of the DINO representations are also around 10^3 higher than for SimSiam. We observe that B4B performs similarly on DINO as for SimSiam.

USER	DEFENSE	# QUERIES	DATASET	TYPE	CIFAR10	STL10	SVHN	F-MNIST
N/A	N/A	<i>Victim</i>	<i>ImageNet</i>	<i>Train</i>	94.51 \pm 0.08	97.98 \pm 0.04	70.66 \pm 0.16	89.98 \pm 0.03
LEGIT	B4B	50K	CIFAR10	QUERY	94.25 \pm 0.11	N/A	N/A	N/A
LEGIT	B4B	5K	STL10	QUERY	N/A	98.05 \pm 0.04	N/A	N/A
LEGIT	B4B	73K	SVHN	QUERY	N/A	N/A	69.66 \pm 0.14	N/A
LEGIT	B4B	60K	F-MNIST	QUERY	N/A	N/A	N/A	89.68 \pm 0.01
ATTACK	NONE	50K	IMAGENET	STEAL	67.92 \pm 0.04	66.02 \pm 0.22	61.30 \pm 0.01	89.46 \pm 0.01
ATTACK	B4B	50K	IMAGENET	STEAL	42.02 \pm 0.05	38.91 \pm 0.06	19.94 \pm 0.02	73.33 \pm 0.04
ATTACK	NONE	100K	IMAGENET	STEAL	75.07 \pm 0.01	76.32 \pm 0.02	71.79 \pm 0.06	89.76 \pm 0.01
ATTACK	B4B	100K	IMAGENET	STEAL	19.27 \pm 0.03	21.24 \pm 0.03	19.84 \pm 0.01	71.01 \pm 0.03
SYBIL	B4B	50K+50K	IMAGENET	STEAL	45.56 \pm 0.06	42.50 \pm 0.02	24.25 \pm 0.03	78.01 \pm 0.08

671 F.6 Additional evaluation of transformations

672 Additionally, we show the impact of transformations on the performance of legitimate users in Table 7
 673 (for both SimSiam and DINO).

Table 7: **Impact of Transformations on the Performance for Legitimate Users.** We show that the transformations applied per-account do not harm the performance of legitimate users on their downstream tasks. The victim encoders was trained on the ImageNet dataset using SimSiam and DINO frameworks.

TRANSFORMATION	ENCODER	CIFAR10	STL10	SVHN	F-MNIST
NONE	<i>Victim SimSiam</i>	90.41 \pm 0.02	95.08 \pm 0.13	75.47 \pm 0.04	91.22 \pm 0.11
AFFINE	SIMSIAM	90.24 \pm 0.11	95.05 \pm 0.1	74.96 \pm 0.18	91.42 \pm 0.15
PAD+SHUFFLE	SIMSIAM	90.4 \pm 0.05	95.34 \pm 0.06	75.47 \pm 0.01	91.38 \pm 0.15
AFFINE+PAD+SHUFFLE	SIMSIAM	90.18 \pm 0.06	95.03 \pm 0.05	74.86 \pm 0.1	91.35 \pm 0.1
BINARY	SIMSIAM	88.78 \pm 0.2	94.72 \pm 0.02	68.42 \pm 0.16	88.91 \pm 0.34
NONE	<i>Victim DINO</i>	94.51 \pm 0.08	97.98 \pm 0.04	70.66 \pm 0.16	89.98 \pm 0.03
AFFINE	DINO	94.25 \pm 0.11	98.05 \pm 0.04	69.77 \pm 0.11	89.68 \pm 0.01
PAD+SHUFFLE	DINO	94.72 \pm 0.02	98.07 \pm 0.03	70.44 \pm 0.1	89.91 \pm 0.08
AFFINE+PAD+SHUFFLE	DINO	94.26 \pm 0.06	98.02 \pm 0.01	69.49 \pm 0.2	89.70 \pm 0.1
BINARY	DINO	92.96 \pm 0.1	98.03 \pm 0.03	59.53 \pm 0.27	88.26 \pm 0.04

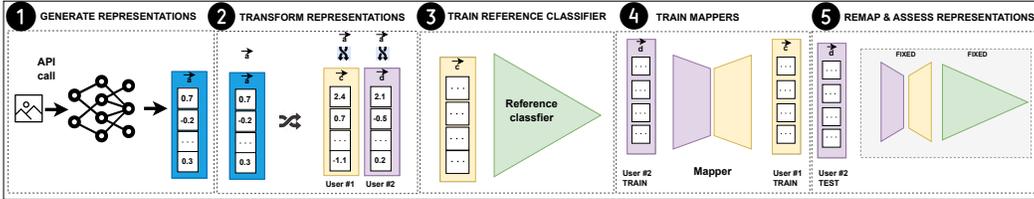


Figure 11: **Protocol to Evaluate the Mapping Between Representations.** We present the protocol of evaluating remappings for two sybil accounts. ① API receives inputs from two sybil accounts and generates corresponding representations. ② Representations are transformed on a per-user basis and returned. ③ Adversary trains a reference classifier on representations from account one. ④ Adversary trains a linear model to find mapping from representations of account two to representations of account one. ⑤ To check the quality of obtained mapping representations from test set of account two are mapped using the fixed mapper (from step 4) to representation space of account one. This enables the calculation of cosine distance between representations from account one and their counterparts from account two shown in Figure 5. Additionally, the fixed reference classifier (from step 3) can be used to measure the accuracy drop caused by remapping .