

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

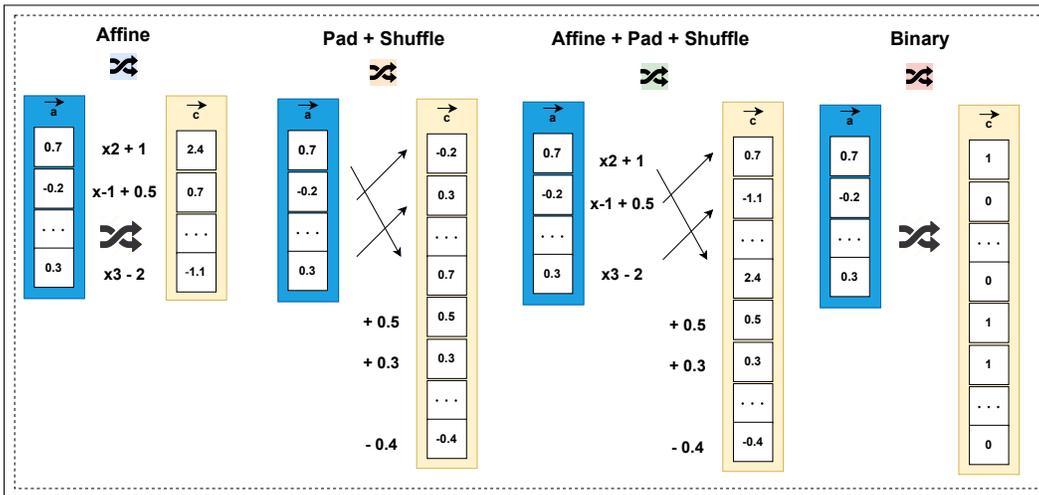


Figure 6: **Overview on Transformations.** We depict the inner-workings of the transformations considered in this work.

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
 580 compression 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.

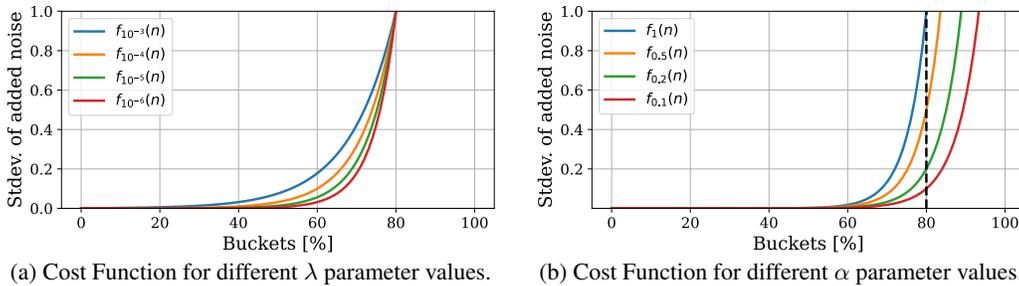


Figure 7: **Effects of λ and α parameters on the Cost Function.** We present the Cost Function for $\alpha=1$, $\beta=80$ and different values of λ (left) and $\lambda = 10^{-6}$, $\beta=80$ and different values of α (right).

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.

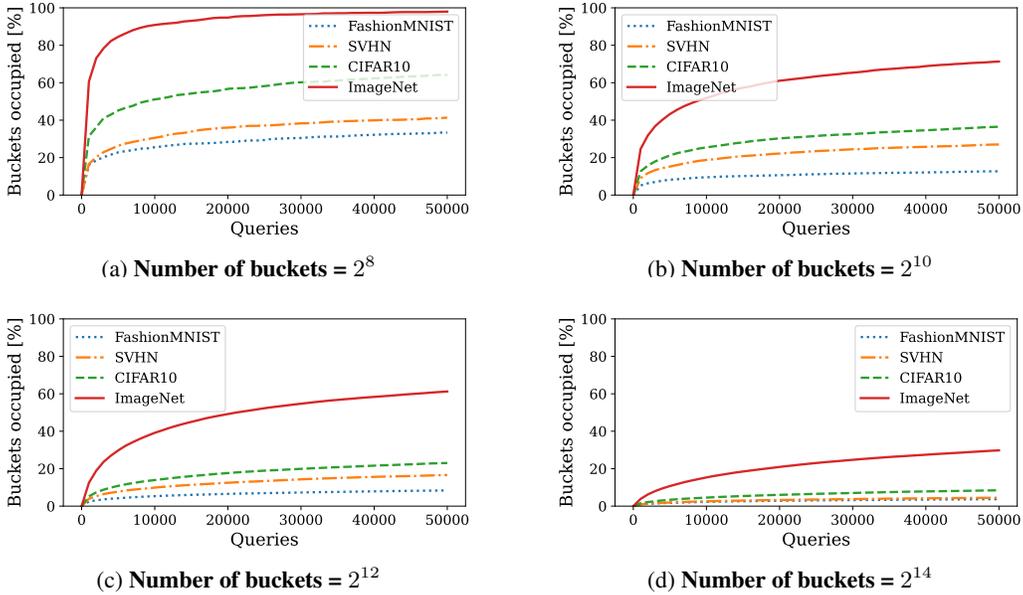


Figure 8: **Estimating Embedding Space Coverage through LSH on SimSiam Encoder.** We extend the results from Figure 7a(a) and present the fraction of buckets occupied by representations of different datasets as a function of the number of queries posed to the encoder. We consider different number of 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

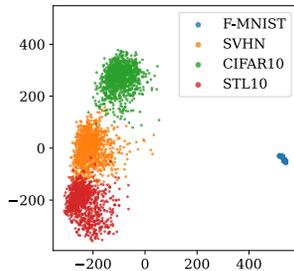


Figure 9: **Representations from Different Tasks Occupy Different Sub-Spaces of the Embedding Space. Presented for FashionMNIST, SVHN, CIFAR10, and STL10.** In this plot, we used the DINO ViT Small encoder trained on ImageNet.

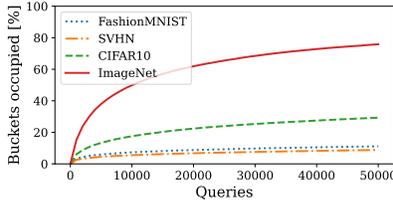


Figure 10: **Estimating Embedding Space Coverage through LSH on the DINO Encoder.** The number of buckets is set to 2^{12} . We also use the same notation and carry out our experiments in the same way as in Figure 7a.

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

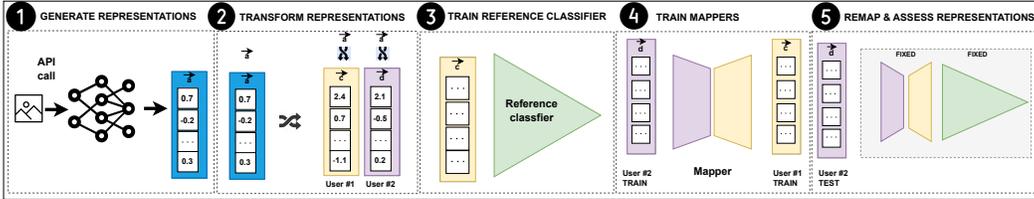


Figure 11: **Protocol to Evaluate the Mapping Between Representations.** We present the protocol of evaluating remappings for two sybil accounts. ① API receives inputs from two sybil accounts and generates corresponding representations. ② Representations are transformed on a per-user basis and returned. 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 .