

# Are Labels Needed for Incremental Instance Learning?

Mert Kilickaya  
 Eindhoven University of Technology  
 kilickayamert@gmail.com

Joaquin Vanschoren  
 Eindhoven University of Technology  
 j.vanschoren@tue.nl

## Abstract

*In this paper, we learn to classify visual object instances, incrementally and via self-supervision (self-incremental). Our learner observes a single instance at a time, which is then discarded from the dataset. Incremental instance learning is challenging, since longer learning sessions exacerbate forgetfulness, and labeling instances is cumbersome. We overcome these challenges via three contributions: i). We propose VINIL, a self-incremental learner that can learn object instances sequentially, ii). We equip VINIL with self-supervision to by-pass the need for instance labelling, iii). We compare VINIL to label-supervised variants on two large-scale benchmarks [6, 33], and show that VINIL significantly improves accuracy while reducing forgetfulness.*

## 1. Introduction

This paper strives for incrementally learning to recognize visual object instances. Visual instance recognition aims to retrieve different views of an input object instance image. It can be seen as fine-grained object recognition, where the goal is to distinguish different instantiations of the same object, such as cup 1 from cup 2. Instance recognition finds applications in many domains, such as in visual search [40], tracking [5, 48, 49] and localization [60].

Distinguishing between different object instances is a challenging task as they often differ only by small nuances. Metric learning [52] is a commonly used approach to learn visual object instances by comparing two views of the same object using a deep convolutional network, such as ResNet [22]. The network is trained to bring representations of the same object close together and separate representations of different objects in a large batch of images.

However, this approach requires iterating over potentially million-scale datasets multiple times to refine the metric space, which can be impractical for privacy reasons (some data may have to be deleted) or scale (when dealing with billions of images). Additionally, using the trained deep net to query a large database of images by comparing

the feature representation of the input image to the database representations is time-consuming and computationally expensive.

This paper builds upon incremental learning to mitigate privacy and scale issues. In incremental learning, the learner observes images from a certain class for a number of iterations. Then, the data of the previous class is discarded, and the learner receives examples from a novel category. Such approach is called class-incremental learning, and receives an increasing amount of attention recently [27, 36, 37, 57].

Existing class-incremental learners are ill-suited for instance-incremental learning for two reasons. First, class-incremental learners rely on full label supervision. Collecting such annotation at the instance level is very expensive. Second, despite years of efforts, class-incremental learners are forgetful, since they lose performance on previously observed categories.

This paper proposes **Visual Self-Incremental Instance Learning**, VINIL, to perform instance-incremental learning, consider Figure 1. VINIL observes multiple views of a single instance at a time, which is then discarded from the dataset. Such examples can be easily captured via turntable cameras [6, 18, 29, 38] or via hand-interactions [15, 34, 50]. Then, VINIL extracts its own supervision via self-supervision [56], therefore self-incremental. Self-incremental learning not only is label-efficient, it also consistently outperforms competitive label-supervised variants, as we will show. In summary, this paper makes three contributions:

- I. In this paper, we study the challenging task of incremental visual instance learning,
- II. We propose VINIL, an incremental instance learner solely guided by self-supervision, by-passing the need for heavy supervision,
- III. Through large-scale experiments on [6, 33], we show that VINIL is more accurate and much less forgetful with respect to competitive label-supervised variants, hence unlocking the potential of large-scale incremental learning for free.

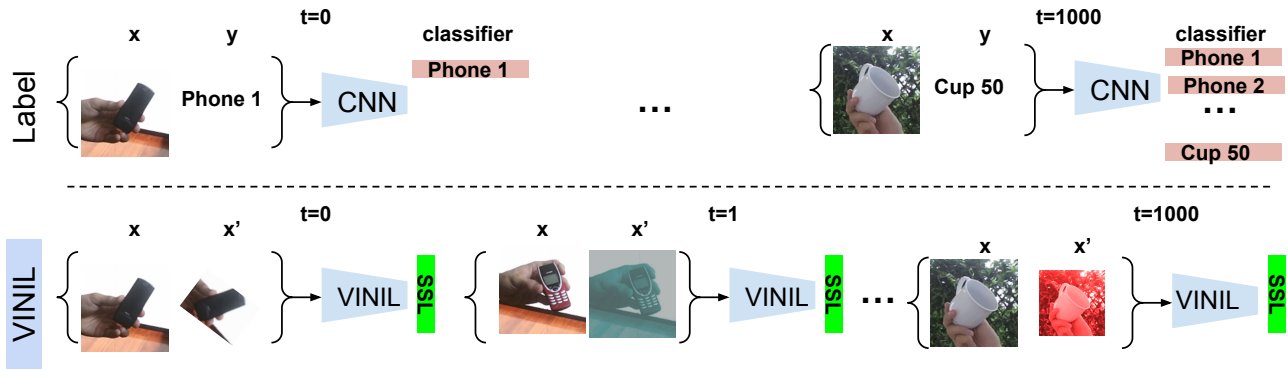


Figure 1. **Top:** Label-incremental learning demands instance-level annotations and trains a new weight per instance. This approach is not suitable for handling large numbers of visual instances, and it is also prone to forgetting previously learned instances. **Bottom:** In this paper, we introduce VINIL, a self-incremental instance learning method. VINIL focuses solely on learning a discriminative embedding and uses Self-Supervised Learning (SSL) to extract supervision from different views of the same instance. As a result, VINIL is label-free, more scalable, and significantly less prone to forgetting compared to label-incremental learning.

## 2. Related Work

**Visual Instance Recognition.** Visual instance recognition has been extensively researched in recent years and has been applied to various computer vision problems, including product retrieval [23, 32, 40, 52], object tracking [5, 48, 49], and geo-localization [31, 46, 51, 53, 59]. The most common approach for these tasks is to induce a discriminative embedding space, often using metric learning techniques [14, 23]. These methods require access to the entire dataset and fine-grained similarity labels. In contrast, this paper presents a novel method for incremental and label-free visual instance recognition.

**Class-Incremental Learning.** Class-incremental learning involves expanding a deep classifier with novel objects, with the goal of maintaining performance on previous categories and avoiding forgetting [36]. Popular techniques to prevent forgetting include regularization, which limits abrupt changes in network weights [28, 30, 43], and memory replay of previous data [4, 24, 45, 47]. Our approach differs from conventional class-incremental learning in two ways. First, while class-incremental learning focuses on object categories, our approach operates at the instance level, presenting new challenges. Second, class-incremental learning requires fully labeled datasets, which is often not possible in instance learning. To overcome these limitations, we use self-supervision and adapt relevant evaluation techniques. Specifically, we use Elastic Weight Consolidation (EwC) as a regularization method [28] and Replay as a memory technique [45] due to their adaptability for label-free learning.

**Self-Supervised Learning.** Self-supervision involves creating pretext tasks to learn deep representations without using labels. Early methods predicted rotations [19] or

patches [39], but contrastive learning has become dominant in recent years [8, 11, 12, 21]. In our work, we utilize self-supervision to extract learning signals in place of instance labels. We experimented with both BarlowTwins [56] and SimSiam [13] due to their high performance and adaptation in incremental learning tasks [16, 35]. We found that BarlowTwins [56] performs better than SimSiam for our incremental learning setup. We believe this is due to its ability to reduce redundancy across different views of the input. Reducing visual redundancy is especially important for different instances of the same object, as visual object instances may only differ in small details.

**Incremental Self-Supervised Learning.** Recently, there has been a surge of interest in use of self-supervision to replace label supervision for incremental learning. We identify three main directions. *i) Pre-training:* Researchers use self-supervised learning either for pre-training prior to incremental learning stage [7, 17, 26] or as an auxiliary loss function to improve feature discrimination [58]. However, these papers still require labels during the incremental learning stage. *ii) Replay:* Second line of techniques propose replay-based methods [10, 35, 42] to supplement self-supervised learners with stored data within the memory. *iii) Regularization:* Third line of work proposes to regularize self-learned representations [16, 20, 35].

In this work, we focus on replay and regularization-based self-incremental learning. More specifically, we closely follow UCL [35] and ask ourselves: What is the contribution of self-supervision for instance incremental learning? Our main observation is that self-supervision consistently yields less forgetful, more accurate and transferable representations, as will be shown via large-scale experiments.

Method	Supervision	Input	Memory	Loss
Fine-Tuning	Label-supervised	$(x, y)$	$\times$	$CE(y, y')$
Fine-Tuning	Self-supervised	$(x)$	$\times$	$BT(x, x')$
EwC	Label-supervised	$(x, y)$	$\times$	$CE(y, y') + Reg(\Theta, y')$
EwC	Self-supervised	$(x)$	$\times$	$BT(x, x') + Reg(\Theta)$
Replay	Label-supervised	$(x, y)$	$(x^m, y^m)$	$CE(y, y') + CE(y^m, y^{m'})$
Replay	Self-supervised	$(x)$	$(x^m)$	$BT(x, x') + BT(x^m, x^{m'})$

Table 1. VINIL performs incremental instance learning via self-supervision, and is compared with label-supervision. We use memory replay [45] and weight regularization [28] as well as simple fine-tuning. *Fine-Tuning* [44] relies on Cross-Entropy (CE) or BarlowTwins (BT) [56] to perform incremental learning. *EwC* [28] penalizes abrupt changes in network weights via regularization ( $Reg(\cdot)$ ). *Replay* [45] replays a part of previous data in the form of input-labels (label-supervised) or input-only (self-supervised).

### 3. VINIL

We present an overview of VINIL in Table 1. The goal of VINIL is to train an embedding network  $f_{\theta_t}(\cdot)$  parameterized by  $\theta_t$ . The network maps an input image  $x$  to a  $D$ -dimensional discriminative embedding,  $h = f_{\theta_t}(x)$  which will then be used to query the database to retrieve different views of the input query for instance recognition. Here,  $t$  denotes the incremental learning step, where the tasks are arriving sequentially:  $\mathbf{T} = (\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_t)$ . We train VINIL via minimizing the following objective:

$$\mathcal{L} = w_c \cdot L_{inst} + (1 - w_c) \cdot L_{incr} \quad (1)$$

where  $w_c$  controls the contribution of instance classification loss  $L_{inst}$  and incremental learning loss  $L_{incr}$ . Incremental learning loss either corresponds to memory replay [45] or weight regularization [28] whereas instance classification loss  $L_{inst}$  is either cross-entropy with labels or a self-supervision objective.

#### 3.1. Incremental Learning

**Fine-Tuning (FT).** A vanilla way to perform incremental instance learning is to apply simple fine-tuning via SGD [44]. In fine-tuning, no incremental learning loss is applied (*i.e.*  $w_c = 1.0$ ) and the sole objective is classification.

In case of label-supervision, a task is defined by a dataset  $D_t^{label} = \{(x_{i,t}, y_{i,t})_{i=1}^{k_t}\}$  where  $k_t$  is the data size at time  $t$ . Then, fine-tuning corresponds to instance discrimination via cross-entropy  $L_{inst} = CE(y_{i,t}, y'_{i,t})$ . Here, instance category prediction for the instance  $i$  at time step  $t$  is obtained with a simple MLP classifier. Notice that this classifier will expand in size linearly with the number of instance categories.

In case of VINIL, a task is defined by a dataset  $D_t^{self} = \{(x_{i,t})_{i=1}^{n_t}\}$  (*i.e.* no labels). Then, fine-tuning corresponds to minimizing the self-supervision objective  $L_{inst} = BT(x_{i,t}, x'_{i,t})$  where  $BT(\cdot)$  is the BarlowTwins [56].

**EwC [28].** EwC penalizes big changes in network weights via comparing the weights in the current and the previous incremental learning step. Originally, EwC re-weights the contribution of each weight to the loss function as a function of instance classification logits (*i.e.* label-supervision). In VINIL, in the absence of labels, we omit this re-weighting and simply use identity matrix.

**Replay [45].** Replay replays a portion of the past data from previous incremental steps to mitigate forgetting. In case of label-supervision, this corresponds to replaying both the input data and their labels via cross-entropy:  $CE(y_{i,t}^m, y_{i,t}^{m'})$  where  $y_{i,t}^{m'}$  is the instance categories for the memory instance  $i$  at time  $t$ . For VINIL, we simply replay the input memory data and its augmented view via self-supervision of BarlowTwins as  $BT(x_{i,t}^m, x_{i,t}^{m'})$ .

#### 3.2. Self-Supervised Learning

In BarlowTwins, the features are extracted from the original and the augmented view of the input image with a siamese deep network, at time step  $t$  as:  $(z_{i,t}, z'_{i,t}) = (f_{\theta^t}(x_{i,t}), f_{\theta^t}(x'_{i,t}))$  where  $x'_{i,t} = aug(x_{i,t})$  is the augmented view of the input. BarlowTwins minimizes the redundancy across views while maximizing the representational information. This is achieved via operating on the cross-covariance matrix via:

$$BT = \sum_i (1 - C_{ii})^2 + w_b \cdot \sum_i \sum_{j \neq i} (C_{ij})^2 \quad (2)$$

where:

$$C_{ij} = \frac{\sum_{\beta} z_{\beta,i} z'_{\beta,j}}{\sum_{\beta} \sqrt{z_{\beta,i}^2} \cdot \sum_{\beta} \sqrt{(z'_{\beta,j})^2}} \quad (3)$$

is the cross-correlation matrix. Here,  $w_b$  controls invariance-redundancy reduction trade-off,  $i$  and  $j$  corresponds to network's output dimensions.

## 4. Experimental Setup

**Implementation.** All the networks are implemented in PyTorch [41]. We use ResNet-18 [22] as the backbone  $f(\cdot)$ , and a single-layer MLP for the instance classifier. We train for 200 epochs for each incremental steps with a learning rate 0.001 decayed via cosine annealing. We use SGD optimizer with momentum 0.9 and batch-size 256. We use random cropping and scaling for augmentation.

We follow the original implementation of BarlowTwins [1]. 10% of the data is stored within the memory for replay [45]. We set  $w_c = 0.7$  and  $w_b = 0.03$ .

**Datasets.** We evaluate VINIL on iLab-20M [6] and Core-50 [33], since they are large-scale, sufficiently different, and widely adopted in incremental learning.

*iLab-20M* is a turntable dataset of vehicles. It consists of 10 objects (*i.e.* bus, car, plane) with varying ([25, 160]) number of instances per category. Objects are captured by varying the background and the camera angle, leading to 14 examples per-instance. We use the public splits provided in [3] with 125k training and 31k gallery images.

*Core-50* is a hand-held object dataset used in benchmarking incremental learning algorithms. The dataset includes 10 objects (*i.e.* phones, adaptors, scissors) with 50 instances per-category. Each instance is captured for 300 frames, across 11 different backgrounds. We use 120k training and 45k gallery images [2].

**Protocol.** We first divide each dataset into 5 tasks, with 2 object categories per-task. Then, each task is subdivided into  $N$  object instance tasks depending on the dataset. We discard the classifier of label-supervised variants after training, and evaluate all models with instance retrieval performance via k-NN with  $k = 100$  neighbors on the gallery set, as is the standard in SSL [8, 11–13, 21].

We use the mean-pooled activations of LAYER4 of ResNet to represent images. All exemplars in the gallery set are used as query.

**Metrics.** We rely on two well established metrics to evaluate the performance of the models, namely accuracy and forgetting.

*i). Accuracy (Acc)* measures whether if we can retrieve different views of the same instance from the gallery set given a query. We measure accuracy for each incremental learning steps, which is then averaged across all sessions.

*ii). Forgetting (For)* measures the discrepancy of accuracy across different sessions. Concretely, it compares the maximum accuracy across all sessions with the accuracy in the last step.

## 5. Experiments

Our experiments address the following research questions: **Q1:** Can VINIL improve performance and reduce forgetting in comparison to label-supervision? **Q2:** Does VINIL learn incrementally generalizable representations across datasets? **Q3:** What makes VINIL effective against label-supervision?

### 5.1. How Does VINIL Compare to Label-Supervision?

First, we compare VINIL’s performance to label-supervision. The results are presented in Table 2.

Method	Core-50		iLab-20M	
	Acc ( $\uparrow$ )	For ( $\downarrow$ )	Acc ( $\uparrow$ )	For ( $\downarrow$ )
FT (Label)	71.450	22.436	89.340	6.500
FT (VINIL)	<b>74.914</b>	<b>4.802</b>	<b>90.398</b>	<b>0.000</b>
Replay (Label)	<b>88.180</b>	<b>6.741</b>	84.464	5.696
Replay (VINIL)	67.677	10.095	<b>90.543</b>	<b>0.000</b>
EwC (Label)	<b>75.117</b>	18.268	87.690	4.535
EwC (VINIL)	73.011	<b>2.167</b>	<b>90.655</b>	<b>0.000</b>

Table 2. Visual Incremental Instance Learning on Core-50 [33] and iLab-20M [6]. VINIL outperforms label-supervised variants for 4 out of 6 settings, while significantly reducing forgetfulness on both datasets. This indicates self-incremental learning is a strong, label-free alternative to label-supervision.

**VINIL Yields Competitive Accuracy.** We first compare the accuracies obtained by VINIL vs. label-supervision. We observe that VINIL yields competitive accuracy against label-supervision: In 4 out of 6 setting, VINIL outperforms label-supervised variants.

**VINIL Mitigates Forgetting.** Secondly, we compare the forget rates of VINIL vs. label-supervision (lower is better). We observe that VINIL consistently leads to much lower forget rates in comparison to label-supervision. On iLab-20M dataset, VINIL results in *no* forgetting. On the more challenging dataset of Core-50, the difference across forget rates are even more pronounced: Label-supervision suffers from 22% forget rate whereas VINIL only by 4%, a relative drop of 80% with fine-tuning.

**Label-supervision Leverages Memory.** Our last observation is that memory improves the accuracy and reduces forgetfulness of label-supervision. In contrast, the use of memory disrupts self-supervised representations. This indicates that replaying both inputs and labels  $((x_i, y_i))$  as opposed to input-only  $((x_i))$ , as in self-supervision) may lead to imbalanced training due to limited memory size [9, 25, 54].

In summary, we conclude that VINIL is an efficient, label-free alternative to label-supervised incremental instance learning. VINIL improves accuracy while reducing forget rate. We also observe that label-supervision closes the gap when an additional memory of past data is present. This motivates further research for improving self-incremental instance learners with memory.

### 5.2. Can VINIL Generalize Across Datasets?

After confirming the efficacy of VINIL within the same dataset, we now move on to a more complicated setting: Cross-dataset generalization. In cross-dataset generalization, we first perform incremental training on Core-50, and then evaluate on iLab-20M. Then, we perform incremental training on iLab-20M and then evaluate on Core-50.

Cross-dataset generalization between Core-50 and iLab-20M is challenging due to the following reasons: *i). Camera:* Core-50 is captured with a hand-held camera whereas iLab-20M is captured on a platform with a turntable camera, *ii). Object Categories:* Object categories are disjoint, as no common objects are present in each dataset, *iii). Object Types:* iLab-20M exhibits toy objects of vehicles whereas Core-50 exhibits hand-interacted daily-life objects.

The results are presented in Table 3. We present iLab-20M to Core-50, and Core-50 to iLab-20M results, along with the relative drop w.r.t training and testing on the same dataset (see Table 2).

Train on $\implies$	iLab-20M	Core-50
Test on $\implies$	Core-50	iLab-20M
Method	$Acc(\% \Delta(\downarrow))$	$Acc(\% \Delta(\downarrow))$
FT (Label)	59.850% <sub>16</sub>	67.249% <sub>24</sub>
FT (VINIL)	<b>66.704%<sub>10</sub></b>	<b>76.302%<sub>15</sub></b>
Replay (Label)	55.692% <sub>36</sub>	69.412% <sub>17</sub>
Replay (VINIL)	<b>61.857%<sub>8</sub></b>	<b>76.125%<sub>15</sub></b>
EwC (Label)	59.030% <sub>21</sub>	70.087% <sub>20</sub>
EwC (VINIL)	<b>70.648%<sub>3</sub></b>	<b>75.793%<sub>16</sub></b>

Table 3. Cross-Dataset Generalization on Core-50 and iLab-20M. We present: *i)* Train on iLab-20M and test on Core-50, *ii)* Train on Core-50 and test on iLab-20M. In addition to accuracy, we also present the relative drop w.r.t training and testing on the same dataset (see Table 2). We observe that VINIL is consistently more robust in cross-dataset generalization when compared with label-supervision. The results indicate that self-supervision is able to extract more domain-agnostic representations, which improves the generality of visual representations, for instance-incremental setup.

**VINIL Yields Generalizable Representations.** We first observe that VINIL consistently yields higher accuracy and

lower drop rate across all 6 settings in both datasets. This indicates that self-supervision extracts more generalizable visual representations from the dataset.

**Label-supervision Overfits with Memory.** Secondly, we observe that label-supervised variants with memory generalizes via overfitting on the training dataset. Replay with label-supervision leads to the biggest drop rate of 36% on Core-50, when trained with iLab-20M. This implies the use of the memory drastically reduces generality of visual representations. A potential explanation is that, since replay utilizes the same set of examples within the limited memory repeatedly throughout learning, this forces the network to over-fit to those examples.

We conclude that VINIL extracts generalizable visual representations from the training source to perform instance incremental training. We also conclude that the astounding performance of label-supervision equipped with memory comes with the cost of overfit, leading to drastic drop in case of visual discrepancies across datasets.

### 5.3. What Factors Affect VINIL’s Performance?

**VINIL Mitigates Bias Towards Recent Task.** We present the heatmaps of the performance for all 5 main tasks, when each task is introduced sequentially, for label-supervision in Figure 2 and for VINIL in Figure 3 on iLab-20M [6]. Each row presents the accuracy for each task, as the tasks are introduced sequentially. For example, the entry (0, 2) denotes the performance on Task-0 when the Task-2 is introduced.

Considering Figure 2 for label-supervision, observe how the tasks achieve their peak performance when they are being introduced to the model, hence the higher numbers within the diagonal. Then, the performance degrades drastically as more and more tasks are being introduced. This indicates label-supervision fails to leverage more data. We call such phenomenon “recency bias”, as the model is biased towards the most recently introduced task.

In contrast, in Figure 3 for VINIL, the performance on each task improves sequentially with the incoming stream of new tasks. This indicates self-supervised representations are less biased towards the recent task, and can leverage data to improve performance. This renders them as a viable option when incremental learning for longer learning steps, such as in incremental instance learning.

**VINIL Focuses on the Object Instance.** We present the activations of the last layer of ResNet, at different incremental time steps, in Figure 4.

Observe how VINIL learns to segment out the target object from the background. This allows the model to accurately distinguish across different instances of the same object sharing identical backgrounds. In contrast, label-supervised variant progressively confuses the object with




Figure 2. Task-level performance of Label-supervision (Fine-tuning). Label-supervision is biased towards recent task.




Figure 3. Task-level performance of VINIL (Fine-tuning). VINIL improves its performance with incoming data, and is less biased towards recent task.

the background. We call such a phenomenon "attentional deficiency" of label-supervised representations.

**VINIL Stores Instance-level Information.** We present nearest neighbors for three queries in Figure 5. We use the average-pooled activations of the last ResNet layer on Core-50 trained with fine-tuning.

Observe how VINIL retrieves the same instance in dif-

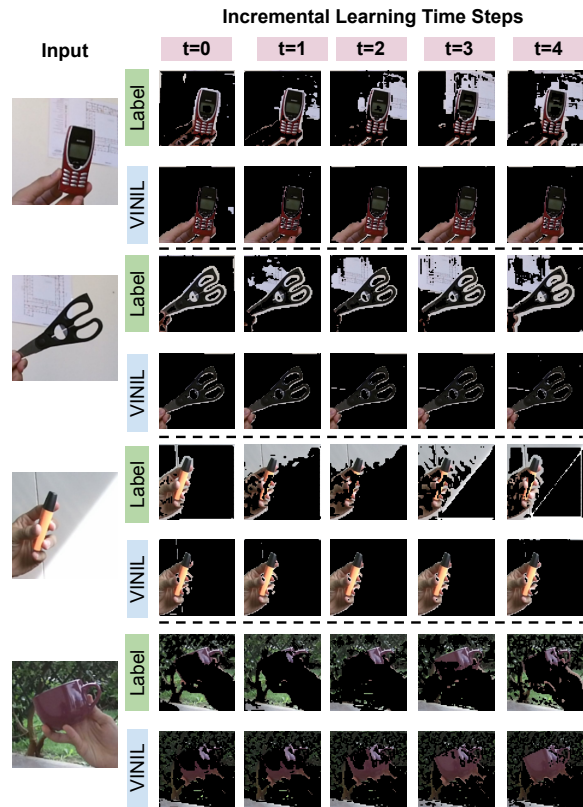


Figure 4. Activations of the last layer of ResNet [22], throughout the incremental learning steps. We compare label-supervision with VINIL (Fine-tuning). Notice how the attention of the label-supervised variant is disrupted after a few learning tasks. Instead, VINIL learns to segment out the target object, successfully suppressing the background context, such as the hand or the background.

ferent viewpoints, such as for the light bulb and can. In contrast, label-supervision is distracted by the background context, as it retrieves irrelevant objects with identical background. This indicates self-supervision generalizes via storing instance-level information. We present a failure case in the last row, as both models fail to represent an object with holes and unfamiliar rotation.

We conclude that VINIL can improve its performance with incoming stream of data, and generalizes via focusing on the target object and storing instance-level details to perform instance-incremental learning.

## 6. Discussion

This paper presented VINIL, a self-incremental visual instance learner. VINIL sequentially learns visual object instances, with no label supervision, via only self-supervision of BarlowTwins [56]. Below, we summarize our main discussion points:

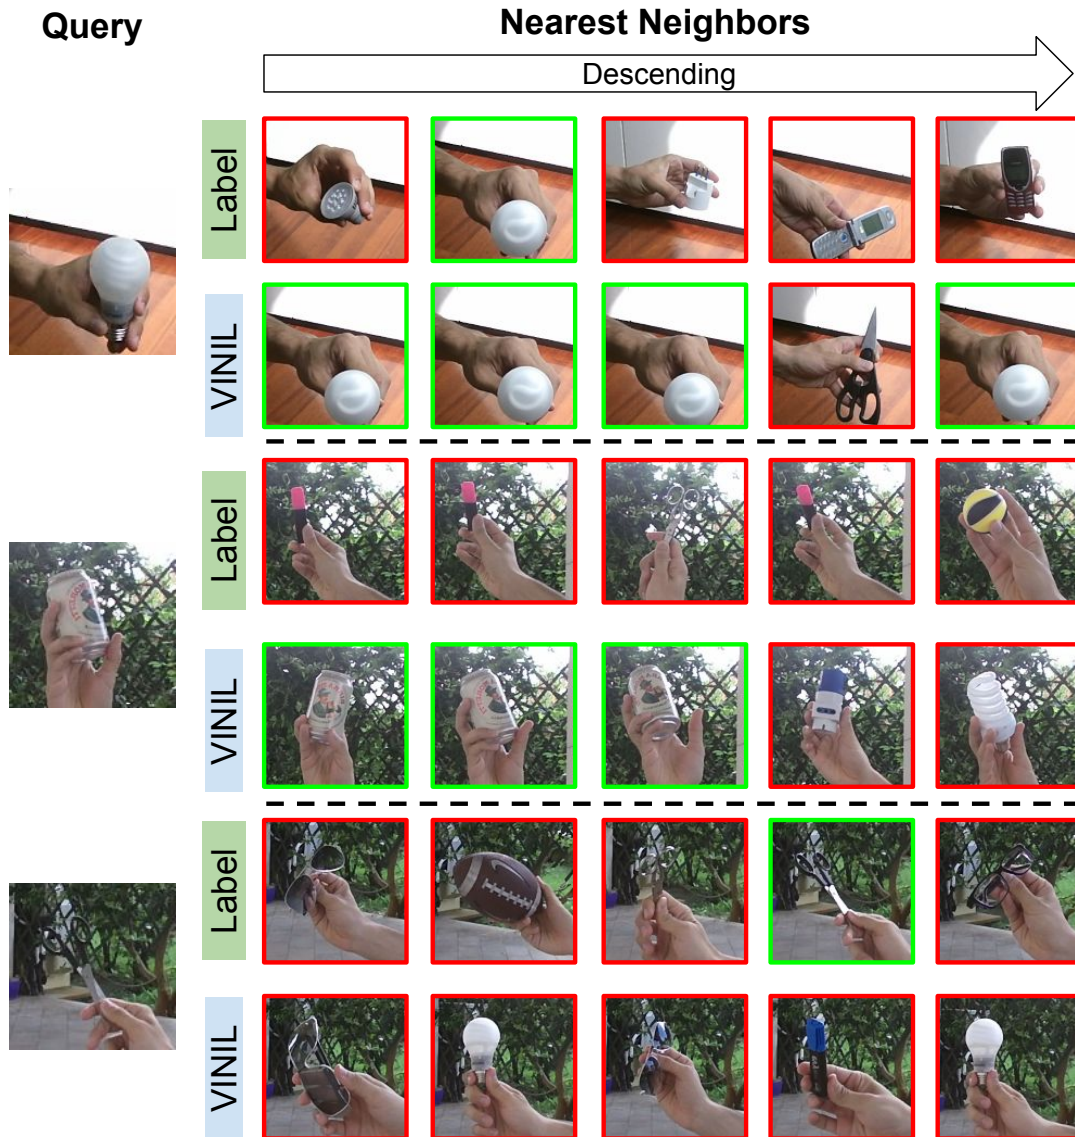


Figure 5. Five nearest neighbors for three object instance queries on Core-50 [33] with fine-tuning. Green is a success, red is a failure. Observe how VINIL retrieves object instances in different views. The last column showcases a failure case, where both models fail to represent an object with holes (scissor).

**Self vs. Label-supervision?** We demonstrate that self-supervision not only omits the need for labels, but it is also more accurate and less forgetful.

**W/ or W/o Memory?** Our results show that the use of memory boosts label-supervised instance incremental learning, however the improvement comes with the cost of over-fitting on the training source.

**Fine-tuning [44] vs. Replay [45] vs. EwC [28]?** We demonstrate that with the use of self-supervision, VINIL closes the gap between simple fine-tuning via SGD and more complicated, compute-intensive techniques like mem-

ory replay or regularization via EwC.

**What Makes VINIL Effective?** VINIL retains representations across tasks, and is able to store and focus on instance-level information, which are crucial for instance-incremental learning.

**Limitation.** VINIL is executed with regularization [28] and memory [45]. One can also consider dynamic networks [55] whose architectures are updated with incoming task data. VINIL is a scalable alternative to dynamic incremental network training due to abundant unlabeled data.

## References

- [1] <https://research.facebook.com/publications/barlow-twins-self-supervised-learning-via-redundancy-reduction/>. 4
- [2] <https://vlomonaco.github.io/core50/index.html.dataset>. 4
- [3] <https://github.com/gyhandy/Group-Supervised-Learning>. 4
- [4] Yogesh Balaji, Mehrdad Farajtabar, Dong Yin, Alex Mott, and Ang Li. The effectiveness of memory replay in large scale continual learning. *arXiv preprint*, 2020. 2
- [5] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *ECCV*, 2016. 1, 2
- [6] Ali Borji, Saeed Izadi, and Laurent Itti. ilab-20m: A large-scale controlled object dataset to investigate deep learning. In *CVPR*, 2016. 1, 4, 5
- [7] Lucas Caccia and Joelle Pineau. Special: Self-supervised pretraining for continual learning. *arXiv preprint*, 2021. 2
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 2020. 2, 4
- [9] Francisco Manuel Castro, Manuel J. Marín-Jiménez, Nicolás Guil Mata, Cordelia Schmid, and Alahari Karteek. End-to-end incremental learning. *ArXiv*, 2018. 4
- [10] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *ICCV*, 2021. 2
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*. PMLR, 2020. 2, 4
- [12] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint*, 2020. 2, 4
- [13] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021. 2, 4
- [14] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 2
- [15] Hehe Fan, Tao Zhuo, Xin Yu, Yi Yang, and Mohan Kankanhalli. Understanding atomic hand-object interaction with human intention. *IEEE TCSVT*, 2021. 1
- [16] Enrico Fini, Victor G Turrissi da Costa, Xavier Alamedapineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *CVPR*, 2022. 2
- [17] Jhair Gallardo, Tyler L Hayes, and Christopher Kanan. Self-supervised training enhances online continual learning. *arXiv preprint*, 2021. 2
- [18] Jan-Mark Geusebroek, Gertjan J Burghouts, and Arnold WM Smeulders. The amsterdam library of object images. *IJCV*, 2005. 1
- [19] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint*, 2018. 2
- [20] Alex Gomez-Villa, Bartłomiej Twardowski, Lu Yu, Andrew D Bagdanov, and Joost van de Weijer. Continually learning self-supervised representations with projected functional regularization. In *CVPR*, 2022. 2
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2, 4
- [22] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 1, 4, 6
- [23] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint*, 2017. 2
- [24] Stella Ho, Ming Liu, Lan Du, Longxiang Gao, and Yong Xiang. Prototypes-guided memory replay for continual learning. *arXiv preprint*, 2021. 2
- [25] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. *CVPR*, 2019. 4
- [26] Dapeng Hu, Shipeng Yan, Qizhengqiu Lu, HONG Lanqing, Hailin Hu, Yifan Zhang, Zhenguo Li, Xinchao Wang, and Jiashi Feng. How well does self-supervised pre-training perform with streaming data? In *ICLR*, 2021. 2
- [27] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *CVPR*, 2022. 1
- [28] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 2017. 2, 3, 7
- [29] Yann LeCun, Fu Jie Huang, and Leon Bottou. Learning methods for generic object recognition with invariance to pose and lighting. In *CVPR*, 2004. 1
- [30] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 2017. 2
- [31] Tsung-Yi Lin, Serge Belongie, and James Hays. Cross-view image geolocalization. In *CVPR*, 2013. 2
- [32] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017. 2
- [33] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In *CoRL*. PMLR, 2017. 1, 4, 7
- [34] Jian Ma and Dima Damen. Hand-object interaction reasoning. *arXiv preprint*, 2022. 1
- [35] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational continuity for unsupervised continual learning. In *ICLR*, 2021. 2
- [36] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint*, 2020. 1, 2
- [37] Sudhanshu Mittal, Silvio Galesso, and Thomas Brox. Essentials for class incremental learning. In *CVPR*, 2021. 1
- [38] Sameer A Nene, Shree K Nayar, Hiroshi Murase, et al. Columbia object image library (coil-100). 1996. 1
- [39] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 2



- [40] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *CVPR*, 2016. 1, 2
- [41] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 4
- [42] Senthil Purushwalkam, Pedro Morgado, and Abhinav Gupta. The challenges of continuous self-supervised learning. *arXiv preprint*, 2022. 2
- [43] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauero. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint*, 2018. 2
- [44] Herbert E. Robbins. A stochastic approximation method. *Annals of Mathematical Statistics*, 2007. 3, 7
- [45] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *NeurIPS*, 2019. 2, 3, 4, 7
- [46] Yujiao Shi and Hongdong Li. Beyond cross-view image retrieval: Highly accurate vehicle localization using satellite image. In *CVPR*, pages 17010–17020, 2022. 2
- [47] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *NeurIPS*, 2017. 2
- [48] Bing Shuai, Andrew Berneshawi, Xinyu Li, Davide Modolo, and Joseph Tighe. Siammot: Siamese multi-object tracking. In *CVPR*, 2021. 1, 2
- [49] Ran Tao, Efstratios Gavves, and Arnold WM Smeulders. Siamese instance search for tracking. In *CVPR*, 2016. 1, 2
- [50] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+o: Unified egocentric recognition of 3d hand-object poses and interactions. In *CVPR*, 2019. 1
- [51] Shruti Vyas, Chen Chen, and Mubarak Shah. Gama: Cross-view video geo-localization. *arXiv preprint*, 2022. 2
- [52] Jian Wang, Feng Zhou, Shilei Wen, Xiao Liu, and Yuanqing Lin. Deep metric learning with angular loss. In *ICCV*, 2017. 1, 2
- [53] Daniel Wilson, Xiaohan Zhang, Waqas Sultani, and Safwan Wshah. Visual and object geo-localization: A comprehensive survey. *arXiv preprint*, 2021. 2
- [54] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Raymond Fu. Large scale incremental learning. *CVPR*, 2019. 4
- [55] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *ArXiv*, 2018. 7
- [56] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 1, 2, 3, 6
- [57] Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *NeurIPS*, 2021. 1
- [58] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *CVPR*, 2021. 2
- [59] Sijie Zhu, Mubarak Shah, and Chen Chen. Transgeo: Transformer is all you need for cross-view image geo-localization. In *CVPR*, 2022. 2
- [60] Sijie Zhu, Taojiannan Yang, and Chen Chen. Vigor: Cross-view image geo-localization beyond one-to-one retrieval. In *CVPR*, 2021. 1