

# Evaluating Contrastive Models for Instance-based Image Retrieval

Tarun Krishna, Kevin McGuinness and Noel O'Connor  
tarun.krishna2@mail.dcu.ie

SFI Insight Centre for Data Analytics, Dublin City University, Ireland

## ABSTRACT

In this work, we evaluate contrastive models for the task of image retrieval. We hypothesise that models that are learned to encode semantic similarity among instances via discriminative learning should perform well on the task of image retrieval, where relevancy is defined in terms of instances of the same object. Through our extensive evaluation, we find that representations from models trained using contrastive methods perform on-par with (and outperforms) a pre-trained supervised baseline trained on the ImageNet labels in retrieval tasks under various configurations. This is remarkable given that the contrastive models require no explicit supervision. Thus, we conclude that these models can be used to bootstrap base models to build more robust image retrieval engines.

## CCS CONCEPTS

• **Computing methodologies** → **Visual content-based indexing and retrieval.**

## KEYWORDS

Deep learning, Contrastive learning, Self-supervised learning

### ACM Reference Format:

Tarun Krishna, Kevin McGuinness and Noel O'Connor. 2021. Evaluating Contrastive Models for Instance-based Image Retrieval. In *Proceedings of ICMR '21: ACM International Conference on Multimedia Retrieval (ICMR '21)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Large scale image retrieval, where the task is to search a large image collection for the most relevant image/content for a given query, is a fundamental task in computer vision. Since their inception, convolutional neural networks (ConvNets) [22, 40] have become the prominent approach for extracting descriptors for image retrieval. These descriptors perform very well in capturing the global semantics of an image and this has led to state-of-the-art results on many benchmark computer vision tasks [7, 18, 37].

The activations in the intermediate layers in ConvNets can be used as a descriptor for an image. These descriptors are often followed by some encoding methods for a compact representation. These encoding techniques range from traditional approaches of

VLAD [21], BoW [25], and Fisher vectors [28], to simple pooling methods like Maximum Activation of Convolution (MAC) [2], Sum Pooling of Convolution (SPoC) [4], Regional-MAC [42], etc. The drawback of these methods is that these (off-the-shelf) network are trained to reduce inter-class variance through supervision on ImageNet classes, this might affect the performance of instance retrieval (i.e. retrieving images that represent the same object or scene as in a query), which is a more fine-grained task.

This drawback has been addressed in the literature by fine-tuning [13, 14, 35, 38]. We hypothesise that a simpler approach could be to retain a traditional off-the-shelf regime but instead use models that are trained based on instance-wise supervision using similarity-based learning. To this end, we investigate contrastive learning based methods, i.e. trained in an unsupervised fashion using contrastive loss [16, 44]. This learning regime relies on learning a meaningful embedding that captures inherent similarity between instances using discriminative approaches [44]. This work investigates the effectiveness of contrastive methods that capture this very idea of instance similarity. To summarize our contributions:

- we extensively evaluate contrastive methods as a fixed feature extractor across different benchmark;
- we provide experimental evidence showing that these models (trained without any explicit supervision) perform on par with a pre-trained supervised baseline (Table 1 and 2);
- we further investigate the role of the dimensionality of the feature embeddings for this task (Table 3).

## 2 RELATED WORK

Conventional image retrieval methods [26, 41] relied on bag-of-words models that exploit local invariant features such as SIFT [24] and large visual vocabularies (e.g. [31]). To aggregate local patches and build a global summary, encoding methods such as Fisher vectors [28] or VLAD [21], have also been proposed [15, 30, 34].

Since the introduction of Deep-ConvNets, [11, 18, 22, 40] there has been a paradigm shift to exploit deep features instead of hand-crafted ones. Intermediate layers in convolutional nets can be used as global or local descriptors. As a result, so-called off-the-shelf [3, 36] features can be used for retrieval. Based on this authors in [4] used sum pooling with a centre prior for aggregating features across spatial dimensions. Other conventional encoding techniques like VLAD [12] or Fisher kernels [29] have also been used in combination with these local feature maps. For example, in [25] proposed BoW-encodings of convolutional features for instance retrieval, whilst [42] proposed R-MAC using max activations over a grid of windows of different scales to obtain compact representations.

Most of the off-the-shelf features are trained on ImageNet [39] to reduce inter-class variance. However, this may degrade the performance of an instance-based retrieval system. One way to address

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICMR '21, July 12–15, 2021, Taipei, Taiwan

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/1122445.1122456>

this is to finetune the model as shown in [5, 13, 14, 35, 38]. In the context of image retrieval, most of the finetuning has been performed on *Landmark* datasets [5], which further requires cleaning of non-related images and potentially expensive post-processing.

Another way is to exploit methods that are trained to reduce intra-class variance, as is the case in contrastive learning. Unlike in supervised learning, these approaches learn to discriminate among individual instances without any concept of categories. This work [44] discusses this notion of instance discrimination. Building on this, a simple formulation is presented in SimCLR [8]. The intuition behind these approaches is to maximize the agreement among augmented views of the same instance using Noise Contrastive Estimation (NCE) [16]. Minimising NCE is equivalent to maximizing mutual information (MI) as was formally shown in CPC [27] as InfoNCE. DIM [19] and AMDIM [6] further extend the idea of InfoNCE across multiple views and scales. One of the downsides of these approaches is that they require large batch sizes on large GPU clusters. To address this drawback, the authors in [9, 17] introduced MoCo, which uses an online and momentum updated offline network that views contrastive learning as a dictionary lookup task. Intuitively, the ability to discriminate among individual instances inherently encoded through the learning makes contrastive learning a good candidate for the task of instance retrieval.

### 3 CONTRASTIVE MODELS

Contrastive learning refers to learning by comparison. This comparison is performed between positive pairs of “similar” and negative pairs of “dissimilar” inputs, which is achieved via a contrastive loss [8, 23] derived from Noise Contrastive Estimation (NCE) [16].

This work targets representative contrastive models for investigation, which are diverse in terms of the way they fuse information. Intuitively, this should lead to each model spanning different feature spaces from an image understanding perspective. The following briefly describes the models considered in this study.

**AMDIM:** *Augmented Multiscale DIM* [6] extends the Deep Info-Max (DIM) framework [19] by learning features extracted from multiple views of a shared context. Local DIM maximizes MI across features extracted independently from augmented views of each input along with features across multiple scales with more powerful encoders. **MoCo:** *Momentum Contrast Representation Learning* [9, 17] alleviates the need for storing offline representations of the entire dataset in memory [44] through the use of a *dynamic* memory *queue*. The samples in the dictionary are progressively replaced. This approach looks at contrastive learning as a way of building a discrete dictionary (*queue*) of inputs (data samples) for a high-dimensional space. We consider MoCo<sub>v1</sub> [17] and MoCo<sub>v2</sub> [9].

**SimCLR:** *Simple framework for Contrastive Learning* [8] is a simplified framework for contrastive learning compared to the previous ones. Here stochastic data augmentations are applied on an input  $\mathbf{x}$  to get two views of  $\mathbf{x}_i$  and  $\mathbf{x}_j$ . Sequentially these augmented inputs are passed through a *base encoder*  $f(\cdot)$  followed by *projection head*, a small network that projects representations from  $f(\cdot)$  to space where a contrastive loss is applied. We investigate SimClr<sub>1x,2x,4x</sub>.

## 4 EXPERIMENTS

This section describes our experimental setup for the evaluation.

### 4.1 Setup

We evaluate models on three standard benchmark datasets: Oxford5k [31], Paris6k [32], and INSTRE [43]. Retrieval performance is measured using mean Average Precision (mAP) following standard procedures for Oxford 5k and Paris 6k benchmarks and for INSTRE evaluating mAP over 1200 images as described in [20]. We further evaluate the performance on revised rOxford 5k and rParis 6k using the new evaluation protocol based on *easy*, *medium*, and *hard* ground truth labels [33]. For the revised benchmarks we report both mAP and mean precision@ (10,5) (mp@10, mp@5).

The goal is not to fine-tune the models but instead evaluate them as a fixed feature extractor to obtain visual descriptors. The base encoder of each of the models is some flavour of ResNet<sup>1</sup> with varying complexity. To this end, we consider the output of the last convolutional layer, i.e. just before the adaptive pooling layer, as our descriptor, which leads to feature maps of size  $\mathbb{R}^{C \times H \times W}$ . To obtain compact representations we use R-MAC ( $L = 3a$ ) [42] over spatial dimensions to get a fixed representation of size  $\mathbb{R}^C$ . We further post-process the vectors by applying  $L_2$  normalization, PCA-whitening, and  $L_2$  normalization again.

We resize our input images to a fixed resolution of  $724 \times 724$  giving a feature spatial dimension of  $23 \times 23$  except in the case of AMDIM where the dimensions are  $40 \times 40$ . However, we downsample this to  $23 \times 23$  to keep uniformity across the evaluation. Also, before running the final evaluation we first run each of the models in training mode (PyTorch `model.train()` just feed-forward) to tune the batch-normalization statistics to the current dataset and then finally test models in evaluation mode (`model.eval()`).

**Baseline.** For comparing across all the contrastive models we use *ResNet50* [18] trained on *ImageNet* as a fixed feature extractor as our pre-trained supervised baseline model. **Note.** For completeness we also evaluate a fine-tuned model [38], which uses Generalized Mean Pooling (instead of R-MAC) trained with Average Precision loss (GeM (AP))<sup>2</sup>. The purpose of this is to provide an indicative upper bound to the evaluation scores.

**Ranking.** We consider global search (G) in this evaluation. We further integrate Global search with Average Query Expansion [10] (AQE), DataBase Augmentation (DBA), [1] and Diffusion [45] (DFS). For AQE we consider nearest neighbour  $N = 10$ , for DBA we consider  $N' = 20$  while combining both of these we consider  $N = 1$  and  $N' = 20$ , based on the findings in [14].

We use a PCA dimension of 512 and evaluate on a global search for R-MAC representations unless otherwise stated.

### 4.2 Results

Table 1 compares different models along with different expansion techniques<sup>3</sup>. For a naive global search on Oxford 5k, the contrastive approach achieves an mAP (%) of 59.40 while the baseline achieves 55.12. Overall best performance is achieved with AQE and DFS for the baseline model (82.64) but this mAP score for SimCLR<sub>2x</sub> (82.34) is in the same range as the former. We also include results

<sup>1</sup>apart from simClr<sub>2x,4x</sub> each uses ResNet50 as a backbone encoder

<sup>2</sup><https://github.com/naver/deep-image-retrieval>

<sup>3</sup>here and in Table 2 GeM (AP) serves as a upper bound indicator rather than a benchmark.

**Table 1: Comparing mAP (%) score across different models. Bold (red) best performing ensemble if it exists.**

Dataset	Method	G	G	G	G	G	G	G	G
			+AQE	+DFS	+DBA	+AQE +DFS	+DBA +AQE	+DBA +DFS	+DBA +AQE +DFS
Oxford 5k	Baseline	55.12	67.85	74.29	62.63	<b>82.64</b>	70.6	77.95	<b>80.4</b>
	SimClr <sub>1x</sub>	51.47	62.96	71.42	61.00	76.40	66.32	72.98	73.38
	SimClr <sub>2x</sub>	58.59	<b>70.49</b>	79.11	67.34	82.34	<b>72.62</b>	<b>79.02</b>	79.96
	SimClr <sub>4x</sub>	<b>59.40</b>	<b>69.80</b>	<b>80.07</b>	<b>67.52</b>	82.00	71.03	78.67	76.98
	MoCo <sub>v1</sub>	56.76	66.89	74.78	64.12	76.44	69.77	73.39	73.95
	MoCo <sub>v2</sub>	58.36	67.49	75.26	65.09	78.07	69.86	72.52	72.86
	AmDim	36.95	43.44	48.13	39.97	52.99	44.69	46.3	47.74
	SimClr <sub>2x</sub> +SimClr <sub>4x</sub>	<b>61.89</b>	<b>72.64</b>	<b>81.66</b>	<b>69.27</b>	<b>86.02</b>	<b>74.43</b>	<b>80.45</b>	<b>81.99</b>
	SimClr <sub>2x</sub> +AmDim	57.06	67.37	75.63	64.69	77.77	69.29	72.84	74.51
	SimClr <sub>2x</sub> +MoCo <sub>v2</sub>	61.33	71.58	78.97	68.47	82.33	72.94	78.36	79.03
	SimClr <sub>4x</sub> +MoCo <sub>v2</sub>	61.87	71.97	80.86	68.20	84.02	72.94	78.56	79.57
	SimClr <sub>4x</sub> +AmDim	58.13	68.76	77.66	64.81	80.05	69.78	75.93	77.32
	GeM (AP)	66.90	78.57	90.36	75.70	94.34	81.92	89.65	91.12
	Paris 6k	Baseline	41.46	61.76	83.74	64.71	90.17	<b>75.56</b>	90.17
SimClr <sub>1x</sub>		42.28	61.15	85.16	64.36	90.17	74.45	89.37	90.49
SimClr <sub>2x</sub>		44.48	59.65	<b>89.39</b>	63.28	<b>91.61</b>	71.91	<b>90.82</b>	<b>91.48</b>
SimClr <sub>4x</sub>		44.96	56.3	88.98	60.40	90.18	68.46	88.18	88.74
MoCo <sub>v1</sub>		43.22	60.34	83.13	63.00	86.9	72.11	87.91	88.73
MoCo <sub>v2</sub>		<b>49.72</b>	<b>64.46</b>	88.02	<b>67.06</b>	89.86	75.01	88.65	89.01
AmDim		34.87	49.2	59.55	47.72	63.7	56.06	64.18	65.08
SimClr <sub>2x</sub> +SimClr <sub>4x</sub>		45.27	56.84	89.97	59.88	91.60	68.07	89.19	90.28
SimClr <sub>2x</sub> +AmDim		44.45	60.06	86.71	63.02	89.06	71.12	88.46	88.76
SimClr <sub>2x</sub> +MoCo <sub>v2</sub>		47.93	61.46	<b>90.92</b>	64.96	<b>92.47</b>	72.37	90.42	91.08
SimClr <sub>4x</sub> +MoCo <sub>v2</sub>		47.44	58.44	90.40	61.88	91.97	69.48	89.00	89.93
SimClr <sub>4x</sub> +AmDim		44.74	57.05	87.59	60.66	88.97	68.60	87.06	87.45
GeM (AP)		49.87	63.46	93.99	67.99	95.56	76.78	90.15	94.07
INSTRE		Baseline	36.64	<b>64.42</b>	71.7	<b>63.96</b>	<b>76.43</b>	<b>74.32</b>	16.41
	SimClr <sub>1x</sub>	27.51	47.85	55.38	47.34	60.61	57.73	21.00	61.15
	SimClr <sub>2x</sub>	36.68	57.04	66.47	56.89	69.48	65.27	24.93	67.18
	SimClr <sub>4x</sub>	<b>44.76</b>	63.73	<b>73.51</b>	63.26	75.02	70.17	28.79	70.54
	MoCo <sub>v1</sub>	33.44	54.88	66.31	54.84	70.1	64.65	29.25	69.8
	MoCo <sub>v2</sub>	33.36	51.79	60.66	50.47	62.56	58.83	22.38	59.77
	AmDim	24.34	37.92	45.47	36.69	47.14	42.74	<b>29.47</b>	47.71
	SimClr <sub>2x</sub> +SimClr <sub>4x</sub>	45.58	63.94	73.54	63.54	74.81	70.54	28.73	70.75
	SimClr <sub>2x</sub> +AmDim	39.37	59.02	68.85	59.09	71.14	66.68	26.19	67.93
	SimClr <sub>2x</sub> +MoCo <sub>v2</sub>	41.99	60.99	70.42	60.64	72.24	68.41	26.09	68.69
	SimClr <sub>4x</sub> +MoCo <sub>v2</sub>	<b>47.07</b>	<b>65.23</b>	<b>74.64</b>	<b>64.93</b>	76.12	71.42	27.09	71.26
	SimClr <sub>4x</sub> +AmDim	45.83	63.88	74.25	63.84	75.28	70.83	<b>30.07</b>	71.60
	GeM (AP)	20.78	31.37	34.85	30.29	38.06	35.93	18.48	37.54

of ensembling contrastive methods<sup>4</sup>, which seems to give a further performance boost (in red in Table 1). A similar inference could be drawn for Paris 6k where the best mAP for global is achieved by MoCo<sub>v2</sub> (49.72) while the overall best (91.61) is achieved with AQE and DFS for SimClr<sub>2x</sub>. A further boost can be observed for the ensemble. In the case of the INSTRE dataset, we see a similar pattern for global search with mAP 44.76 corresponding to SimClr<sub>4x</sub> while the overall best is achieved with AQE and DFS for the baseline

(76.43) while the best result achieved for contrastive models is 75.02 for SimClr<sub>4x</sub>. This clearly indicates that contrastive methods trained to reduce intra class variance capture the notion of instance similarity which is being reflected in this evaluation. Also expansion techniques further boosts the performance over global search.

To further consolidate our findings, we also conducted an evaluation on the revised rOxford 5k and rParis 6k datasets as depicted in Table 2. On rOxford 5k SimClr<sub>2x</sub> gives the best performance on all labels. mP@10 is almost 70% for the *easy* category, with the drop in

<sup>4</sup>The R-MAC representations are concatenated and dimensionally reduced via PCA

**Table 2: Comparison on global ranking across different model. Bold (red) best performing ensemble if it exists.**

Dataset	Method	Easy			Medium			Hard		
		mAP	mp@5	mp@10	mAP	mp@5	mp@10	mAP	mp@5	mp@10
rOxford 5k	Baseline	45.65	67.55	61.81	32.89	64.00	58.57	12.10	23.71	19.43
	SimClr <sub>1x</sub>	47.17	69.78	63.90	31.98	65.24	57.52	9.52	20.57	15.86
	SimClr <sub>2x</sub>	<b>54.95</b>	<b>76.69</b>	<b>69.49</b>	38.54	73.14	<b>65.33</b>	<b>14.41</b>	<b>31.21</b>	<b>22.93</b>
	SimClr <sub>4x</sub>	54.65	74.80	68.37	<b>38.57</b>	<b>73.71</b>	65.05	14.05	30.07	22.79
	MoCo <sub>v1</sub>	48.29	68.48	64.36	33.57	64.43	58.58	9.27	20.05	15.62
	MoCo <sub>v2</sub>	52.69	72.33	65.86	36.49	67.24	59.24	10.72	23.38	18.67
	AmDim	21.24	38.31	31.35	17.54	36.38	32.57	4.11	5.92	6.33
	SimClr <sub>2x</sub> +SimClr <sub>4x</sub>	<b>55.23</b>	76.15	68.87	39.60	<b>74.29</b>	<b>65.64</b>	15.01	29.43	<b>23.97</b>
	SimClr <sub>2x</sub> +AmDim	52.08	72.28	65.51	36.19	66.86	58.57	10.20	21.64	17.64
	SimClr <sub>2x</sub> +MoCo <sub>v2</sub>	55.11	76.47	69.35	38.46	70.52	62.57	12.51	27.12	21.83
	SimClr <sub>4x</sub> +MoCo <sub>v2</sub>	53.98	72.40	67.90	38.31	70.10	64.69	13.38	28.43	23.13
SimClr <sub>4x</sub> +AmDim	49.23	70.76	64.00	35.72	67.62	61.33	11.72	23.79	18.80	
GeM (AP)	64.07	84.93	80.56	51.03	89.43	83.86	30.30	54.86	44.00	
rParis 6k	Baseline	47.20	91.14	87.00	31.63	93.71	91.71	10.44	57.14	43.43
	SimClr <sub>1x</sub>	49.35	93.14	89.52	31.61	95.71	92.00	8.93	46.57	38.43
	SimClr <sub>2x</sub>	54.28	<b>94.00</b>	91.14	34.47	96.00	92.86	10.87	62.00	49.71
	SimClr <sub>4x</sub>	54.97	93.71	<b>91.71</b>	35.39	<b>96.29</b>	<b>94.86</b>	12.34	<b>66.57</b>	<b>56.29</b>
	MoCo <sub>v1</sub>	50.47	92.29	89.71	31.48	95.14	92.57	8.13	46.57	36.29
	MoCo <sub>v2</sub>	<b>55.77</b>	92.86	89.24	<b>36.32</b>	94.86	92.14	<b>11.71</b>	60.00	50.43
	AmDim	38.62	80.29	73.71	25.92	82.57	75.71	6.36	25.14	19.57
	SimClr <sub>2x</sub> +SimClr <sub>4x</sub>	55.54	94.00	<b>91.86</b>	35.59	96.00	94.00	12.37	66.29	56.00
	SimClr <sub>2x</sub> +AmDim	53.71	93.71	90.43	34.57	95.14	93.14	10.46	55.71	46.00
	SimClr <sub>2x</sub> +MoCo <sub>v2</sub>	<b>57.00</b>	93.43	91.00	<b>36.66</b>	95.14	93.29	11.88	64.00	52.43
	SimClr <sub>4x</sub> +MoCo <sub>v2</sub>	56.86	<b>94.29</b>	91.71	36.61	96.29	94.86	<b>12.68</b>	<b>69.43</b>	55.57
SimClr <sub>4x</sub> +AmDim	54.65	92.86	90.57	35.06	95.71	93.43	11.41	61.14	48.71	
GeM (AP)	54.90	94.38	91.95	37.36	98.86	97.00	14.65	76.86	63.71	

**Table 3: Comparison of mAP (%) across different PCA dimension and the true dimension.**

Dataset	Method	PCA-Whitening						True dim.	
		32	64	128	256	512	1024		2048
Oxford 5k	Baseline	48.31	54.61	56.68	57.52	55.12	49.44	37.23	<b>58.47</b>
	SimClr <sub>1x</sub>	34.96	44.96	54.17	<b>54.79</b>	51.47	44.16	34.94	50.63
	SimClr <sub>2x</sub>	36.79	47.96	60.32	<b>61.68</b>	58.59	49.82	24.62	53.72
	SimClr <sub>4x</sub>	35.32	45.86	57.77	<b>61.70</b>	59.40	53.83	44.71	41.94
	MoCo <sub>v1</sub>	29.90	41.95	54.62	<b>57.38</b>	56.76	50.20	39.37	38.73
	MoCo <sub>v2</sub>	39.92	47.88	57.92	<b>60.43</b>	58.36	52.14	40.05	51.76
	AmDim	11.52	15.30	21.71	28.71	36.95	<b>37.58</b>	30.73	16.10
Paris 6k	Baseline	<b>72.22</b>	71.83	63.65	53.40	41.46	28.99	17.89	68.36
	SimClr <sub>1x</sub>	69.90	<b>73.29</b>	66.62	54.58	42.28	30.7	20.06	66.60
	SimClr <sub>2x</sub>	75.21	<b>77.90</b>	69.25	57.83	44.48	33.05	23.27	72.20
	SimClr <sub>4x</sub>	77.16	<b>78.04</b>	69.19	57.29	44.96	35.02	25.79	72.89
	MoCo <sub>v1</sub>	56.84	<b>63.63</b>	61.87	54.07	43.22	32.57	21.63	53.66
	MoCo <sub>v2</sub>	70.08	<b>75.11</b>	71.24	61.66	49.72	35.73	22.98	69.99
	AmDim	21.49	33.96	<b>41.19</b>	40.85	34.87	26.82	17.85	25.50
INSTRE	Baseline	26.44	34.92	<b>38.68</b>	38.55	36.64	29.25	20.20	33.03
	SimClr <sub>1x</sub>	16.47	23.27	29.01	<b>30.35</b>	27.51	21.85	15.97	21.85
	SimClr <sub>2x</sub>	18.42	26.83	35.88	<b>39.67</b>	36.68	29.64	21.57	25.94
	SimClr <sub>4x</sub>	18.92	28.97	40.45	<b>46.99</b>	44.76	36.68	27.65	28.98
	MoCo <sub>v1</sub>	20.08	27.85	33.77	<b>36.24</b>	33.44	26.66	18.46	23.01
	MoCo <sub>v2</sub>	19.00	27.38	34.34	<b>36.22</b>	33.36	26.86	18.79	26.22
	AmDim	10.65	15.88	20.55	24.15	<b>24.34</b>	20.45	14.45	10.25

performance for *hard* label. Similarly, on rParis 6k, SimClr<sub>4x</sub> gives the best performance with mP@10 over 90 for easy and medium, but again this drops off for *hard* labels. As in the Table 1, here

ensembling further boosts performance. Again, contrastive models surpass the baseline pre-trained supervised model.

**Effect of descriptor dimension on performance.** Table 3 reports our findings for global search<sup>5</sup>. Interestingly, true dimensions ( $L_2$  normalized R-MAC representations) appear to perform worse for almost all the models. The best dimension varies across the dataset but it is never the true dimension. This could be attributed to dimensions with small principal components being noisy and redundant and adversely affecting performance.

## 5 CONCLUSION

This work evaluated contrastive models for the task of instance-based image retrieval. Our evaluation found that these methods are on par with those trained on class labels. In fact, in many settings in Table 1, 2 contrastive approaches surpass the supervised model. The quantitative evaluation shows that these contrastive methods can easily surpass supervised models without any explicit supervision.

## 6 ACKNOWLEDGMENT

This work has emanated from research supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289\_P2, co-funded by the European Regional Development Fund and Xperi FotoNation.

<sup>5</sup>Comparison is across the horizontal dimension (columns)

## REFERENCES

- [1] Relja Arandjelović and Andrew Zisserman. 2012. Three things everyone should know to improve object retrieval. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2911–2918.
- [2] Hossein Azizpour, Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, and Stefan Carlsson. 2015. From Generic to Specific Deep Representations for Visual Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [3] Hossein Azizpour, Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, and Stefan Carlsson. 2015. From generic to specific deep representations for visual recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 36–45.
- [4] Artem Babenko and Victor Lempitsky. 2015. Aggregating deep convolutional features for image retrieval. *arXiv preprint arXiv:1510.07493* (2015).
- [5] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. 2014. Neural codes for image retrieval. In *European conference on computer vision*. Springer, 584–599.
- [6] Philip Bachman, R Devon Hjelm, and William Buchwalter. 2019. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910* (2019).
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2014. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062* (2014).
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.), PMLR, 1597–1607. <http://proceedings.mlr.press/v119/chen20j.html>
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. 2020. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297* (2020).
- [10] Ondrej Chum, James Philbin, Josef Sivic, Michael Isard, and Andrew Zisserman. 2007. Total recall: Automatic query expansion with a generative feature model for object retrieval. In *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 1–8.
- [11] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*. PMLR, 647–655.
- [12] Yunchao Gong, Liwei Wang, Ruiqi Guo, and Svetlana Lazebnik. 2014. Multi-scale orderless pooling of deep convolutional activation features. In *European conference on computer vision*. Springer, 392–407.
- [13] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. 2016. Deep image retrieval: Learning global representations for image search. In *European conference on computer vision*. Springer, 241–257.
- [14] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. 2017. End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision* 124, 2 (2017), 237–254.
- [15] Albert Gordo, Jose A Rodriguez-Serrano, Florent Perronnin, and Ernest Valveny. 2012. Leveraging category-level labels for instance-level image retrieval. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 3045–3052.
- [16] Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 297–304.
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9729–9738.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [19] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. 2018. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670* (2018).
- [20] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondrej Chum. 2017. Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2077–2086.
- [21] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. 2010. Aggregating local descriptors into a compact image representation. In *2010 IEEE computer society conference on computer vision and pattern recognition*. IEEE, 3304–3311.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012), 1097–1105.
- [23] Phuc H Le-Khac, Graham Healy, and Alan F Smeaton. 2020. Contrastive representation learning: A framework and review. *IEEE Access* (2020).
- [24] David G Lowe. 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision* 60, 2 (2004), 91–110.
- [25] Eva Mohedano, Kevin McGuinness, Noel E O’Connor, Amaia Salvador, Ferran Marques, and Xavier Giró-i Nieto. 2016. Bags of local convolutional features for scalable instance search. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*. 327–331.
- [26] David Nister and Henrik Stewenius. 2006. Scalable recognition with a vocabulary tree. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, Vol. 2. Ieee, 2161–2168.
- [27] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [28] Florent Perronnin and Christopher Dance. 2007. Fisher kernels on visual vocabularies for image categorization. In *2007 IEEE conference on computer vision and pattern recognition*. IEEE, 1–8.
- [29] Florent Perronnin and Diane Larlus. 2015. Fisher vectors meet neural networks: A hybrid classification architecture. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3743–3752.
- [30] Florent Perronnin, Yan Liu, Jorge Sánchez, and Hervé Poirier. 2010. Large-scale image retrieval with compressed fisher vectors. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 3384–3391.
- [31] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. 2007. Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE conference on computer vision and pattern recognition*. IEEE, 1–8.
- [32] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. 2008. Lost in quantization: Improving particular object retrieval in large scale image databases. In *2008 IEEE conference on computer vision and pattern recognition*. IEEE, 1–8.
- [33] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. 2018. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5706–5715.
- [34] Filip Radenović, Hervé Jégou, and Ondrej Chum. 2015. Multiple measurements and joint dimensionality reduction for large scale image search with short vectors. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval*. 587–590.
- [35] Filip Radenović, Giorgos Tolias, and Ondrej Chum. 2016. CNN image retrieval learns from BoW: Unsupervised fine-tuning with hard examples. In *European conference on computer vision*. Springer, 3–20.
- [36] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. 2014. CNN features off-the-shelf: an astounding baseline for recognition. 2014. *arXiv preprint arXiv:1403.6382* (2014).
- [37] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* 39, 6 (2016), 1137–1149.
- [38] Jérôme Revaud, Jon Almazán, Rafael Sampaio de Rezende, and César Roberto de Souza. 2019. Learning with Average Precision: Training Image Retrieval with a Listwise Loss. *CoRR abs/1906.07589* (2019). [arXiv:1906.07589](http://arxiv.org/abs/1906.07589) <http://arxiv.org/abs/1906.07589>
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.
- [40] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [41] Josef Sivic and Andrew Zisserman. 2003. Video Google: A text retrieval approach to object matching in videos. In *Computer Vision, IEEE International Conference on*, Vol. 3. IEEE Computer Society, 1470–1470.
- [42] Giorgos Tolias, Ronan Sifre, and Hervé Jégou. 2015. Particular object retrieval with integral max-pooling of CNN activations. *arXiv preprint arXiv:1511.05879* (2015).
- [43] Shuang Wang and Shuqiang Jiang. 2015. Instre: a new benchmark for instance-level object retrieval and recognition. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 11, 3 (2015), 1–21.
- [44] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3733–3742.
- [45] Fan Yang, Ryota Hinami, Yusuke Matsui, Steven Ly, and Shin’ichi Satoh. 2019. Efficient image retrieval via decoupling diffusion into online and offline processing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 9087–9094.