

# Bag of Negatives for Siamese Architectures

Bojana Gajić<sup>123</sup>  
bgajic@cvc.uab.es

Ariel Amato<sup>1</sup>  
aamato@vintra.io

Ramon Baldrich<sup>23</sup>  
ramon@cvc.uab.es

Carlo Gatta<sup>1</sup>  
cgatta@vintra.io

<sup>1</sup> Vintra Inc  
San Jose, California, USA

<sup>2</sup> Autonomous University of Barcelona  
Bellaterra, Barcelona, Spain

<sup>3</sup> Computer Vision Center  
Bellaterra, Barcelona, Spain

---

## Abstract

Training a Siamese architecture for re-identification with a large number of identities is a challenging task due to the difficulty of finding relevant negative samples efficiently. In this work we present Bag of Negatives (BoN), a method for accelerated and improved training of Siamese networks that scales well on datasets with a very large number of identities. BoN is an efficient and loss-independent method, able to select a bag of “high quality negatives”, based on a novel online hashing strategy.

## 1 Introduction

Instance retrieval and re-identification (re-ID) are ranking problems where, given an image of an object, a system is searching for the images of the same object, usually taken by a different camera, under different light, from different angles etc. Several research lines address this problem [8, 16, 17, 21, 22, 23, 30, 31, 32]. Some of these works propose methods that are computationally very expensive, using network architectures that are task-specific [16, 17, 21, 23, 30, 31, 32]. Others focus on keeping architectures simple and flexible to adapt to any ranking problem [9, 22, 25].

Once the network architecture is designed, the question that naturally arises is: what should be the appropriate loss function for training? Several works use classification loss, due to its simplicity and the availability of several tools ready for the task [10, 34]. These methods typically train the network as if every identity in the dataset is a class, while extracting image descriptors by pooling features from the last convolutional layer at testing time. Even though these approaches are very simple, they suffer from serious limitations in terms of the number of classes they can handle. For instance, a dataset with more than 10M identities (such as the FaceNet dataset, as reported in [35]) and with, e.g, a 128 embedding size, would require a fully connected layer with 1.28 billion extra parameters, which would be impossible to train from a practical perspective.

Other approaches propose alternative functions, widely known as ranking losses [9, 22, 24]. The main goal of these losses is to bring the representations of samples from the same ID close to each other, while separating the representations of images that belong to different

classes. To do so, Siamese architectures with several parallel streams, made of the same networks with shared weights, are used [9, 22, 25]. These architectures allow computing compact representations for several input images simultaneously, combining them with a ranking loss. Even though these approaches do not depend on the number of classes in the training set, they are extremely hard to train due to the complexity of finding input pairs, triplets or quadruplets that produce a loss greater than zero. In this paper we propose an online strategy for mining samples, which contributes to a more efficient training of Siamese architectures.

## 1.1 Relevant work

Online hard negative mining (OLHN) methods take advantage of the sample representations available at mini-batch level in order to improve the probability of retrieving relevant negatives for the triplet loss [9, 19, 22]. The latter two methods will be discussed more in detail in the next section since they are the most popular ([9, 18, 25]) and most effective OLHNs, and we aim at improving their performance.

An offline re-weighting of the loss can improve the quality of negative samples, but at non-negligible computational cost [18]. Taking advantage of extra knowledge on sub-categories within the dataset is also advantageous in mining negative samples [27].

In recent years, online methods that provide relevant negative samples prior to triplet formation have been proposed, thus not at mini-batch level.

In [9] the authors propose a strategy that builds a tree of identities to facilitate the sampling of relevant negatives for a given anchor. The method clearly improve the quality of the negative samples but at the cost of updating the tree at every epoch. Also, the tree construction is based on an identity-to-identity distance matrix, which thus scales polynomially with the number of identities.

In [25] the authors explicitly face the problem of training a Siamese network with 100k identities. The basic idea is to generate a representation for each identity, and apply clustering on all the identities to generate clusters or subspaces, wherein identities are similar in each subspace. The authors propose to train a classifier on a subset of identities, then use the classifier to generate image representations, and finally perform k-means clustering on them in order to form the subspaces. The authors do not update the cluster during the training, thus the subspaces could become sub-optimal in later stages of the training.

## 2 Motivation

The triplet loss (see Equation (1)) is based on the construction of triplets  $i \in \mathcal{T}$  formed by an anchor sample  $x_i^a$ , a positive sample  $x_i^p$  (belonging to the same ID as the anchor) and a negative sample  $x_i^n$  (belonging to a different ID). The samples are mapped into an embedding by a given function  $f(\cdot)$ , that is usually a deep convolutional network, whose parameters are learned by minimizing the loss  $\mathcal{L}$ . The goal of the triplet loss is to ensure that the anchor-negative pairs are far from each other by a margin  $\alpha$  with respect to the anchor-positive pair distance. The most challenging part of using the triplet loss is finding triplets that produce a non-zero loss. This is hard, since the number of all possible triplets in the dataset is proportional to the cube of total number of images  $N$ ,  $|\mathcal{T}| \sim N^3$ , and the more the system trains, the less probable it is

to find a relevant negative for a given anchor-positive pair.

$$\mathcal{L} = \sum_{i \in \mathcal{T}} [\|f(x_i^a) - f(x_i^p)\|_2^2 - \|f(x_i^a) - f(x_i^n)\|_2^2 + \alpha]_+ \quad (1)$$

Let  $n$  be the average number of images per ID,  $m$  the mini-batch size,  $b$  number of triplets per mini-batch ( $b = m/3$ ), and  $k$  the number of images of each ID in the mini-batch. We introduce the notation  $\hat{n}$ , the number of negative samples that produce a non-zero loss if used in conjunction with the triplet loss and an anchor-positive pair. The “quality” of the retrieved negatives is also relevant, as pointed out in [18]: negative samples have to be selected such that the anchor-negative distance distribution is almost uniform.

Random negative sampling from the whole dataset (to which we refer as Vanilla sampling) has complexity  $\mathcal{O}(1)$  but does not provide relevant negative samples except in the beginning of the training, since the probability of picking a relevant negative,  $p_{\hat{n}} = \hat{n}/(N - n) \simeq \hat{n}/N$ . From now on we will omit  $n$  from the formula since it is negligible w.r.t.  $N$ .

*Semi hard* loss [27] employs a negative sampling strategy that has an increased cost due to the fact that the additional computed distances scale polynomially with the mini-batch size. The improvement in  $p_{\hat{n}}$  with respect to Vanilla sampling is linearly dependent on the number of triplets  $b$ . For this reason, authors use huge mini-batches in the order of 1800 samples.  $p_{\hat{n}}$  is thus increased at the cost of large mini-batches and additional sample to sample distance computation.

*Batch hard* loss [9] is an improved version of the *semi hard* loss where, thanks to a more controlled mini-batch creation and additional distances computation, the method exhibits  $p_{\hat{n}} = 3b\hat{n}/N$ . This strategy offers a 50% improvement in  $p_{\hat{n}}$  w.r.t. the *semi hard* approach, but still provides a probability that depends on the mini-batch size.

Finally, an offline exhaustive search into the dataset provides  $p_{\hat{n}} = 1$ . This is, of course, not viable for large datasets. Nonetheless, for relatively small datasets and with the proper sampling strategy over the  $mN$  distances, exhaustive search provides excellent negatives samples [18].

### 3 Contribution

The paper’s *main contribution* is a **loss-independent** and **computationally inexpensive online** strategy for **improved negative mining** in large datasets, which we named Bag of Negatives (BoN). The main advantages of BoN w.r.t. previous methods are: (1) faster training due to more non-zero loss triplets; (2) better performance on validation sets due to high quality negatives; and (3) a negligible additional computational cost w.r.t. the Siamese architecture training.

BoN does not require computing additional sample representations nor their respective distances to be able to select an appropriate negative. Also, it does not require additional networks and/or losses to be applied on the input image, nor ad-hoc strategies. Moreover, it can be combined with any loss that requires a negative sample. Nonetheless, for simplicity, we will show the behaviour of BoN with triplets; analogous results are achieved with other ranking losses, such as, e.g, contrastive loss. We also want to emphasize that this approach is devised with large datasets and computational efficiency in mind. Finally, the method has only one relevant meta-parameter, discussed in 5.

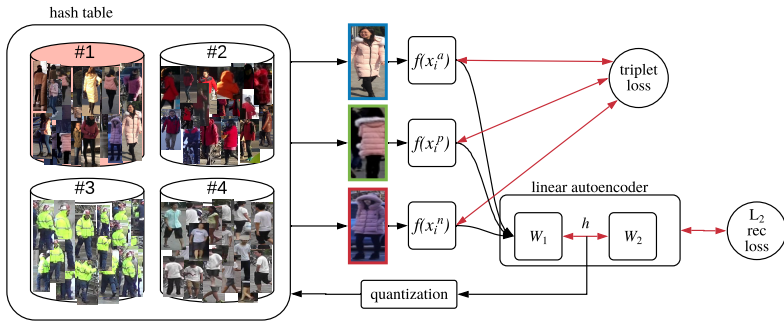


Figure 1: BoN strategy. Triplets with good quality negatives are formed using the information from the hash table. The resulting embedding is used to learn both the deep model  $f$  and a linear projection that, in turn, provides a low-dimensional embedding. Its quantization provides (possibly) new entry positions in the hash table for the input images. The hash table and the linear autoencoder are updated at each training step with minimal overhead.

## 4 Bag of Negatives

A negative sample whose representation is close to the anchor sample provides a triplet that is more likely to produce non-zero loss. The main purpose of BoN is providing these relevant negative samples using an online algorithm that is loss-independent and computationally inexpensive. BoN is inspired by the Spectral Hashing method [29]. Nonetheless, we implemented several changes in order to efficiently adapt to the negative mining problem during training.

Spectral Hashing is a nearest neighbour search algorithm that is shown to be better performing than Product Quantization while being simpler to implement and more efficient at learning the hash function [29]. In terms of performance, it is inferior with respect to methods that address the embedding compression and the quantization as a whole problem, e.g. [4, 6]. However, we have to consider that the embedding is changing during the training, thus a simpler but flexible approach is preferred over methods providing better results at a greater computational cost.

The main approach of the Spectral Hashing method is to (1) learn a linear projection from a high dimensional space (of size  $e$ ) to a lower dimensional space (of size  $s \ll e$ ) by means of a standard PCA; (2) apply the projection to a sample; (3) perform a 1 bit quantization over every dimension by threshold at 0; and (4) group the  $s$  bits into an integer codeword. The codeword represents the entry of an hash table. The underlying assumption is that samples falling into the same bin are neighbours in the high dimensional embedding. Of course, this assumption is overly optimistic and the variation from the optimal behavior is mainly due to the following facts: (1) since  $s \ll e$  we lose some information about the topology of the high dimensional embedding and (2) the quantization is extremely harsh and there is no actual control over the quantization error during the process. Nonetheless, experimental validation shows that the Spectral Hashing method indeed performs well in retrieval tasks [29].

However, the direct application of the Spectral Hashing (or any other nearest neighbour efficient algorithm) method to the problem of retrieving negative samples is not straightforward because the embedding is dynamically changing during the training. One can compute the whole embedding every certain number of steps, plus computing the PCA, and the hash table,

but this naïve strategy does not scale well for large datasets. Consequently, we propose three main modifications to the Spectral Hashing approach in order to have an **online** algorithm that mimics its performance:

1. The PCA is substituted by a linear auto-encoder paired with  $L_2$  reconstruction loss.
2. The quantization threshold is dynamically estimated, per dimension.
3. The hash table is dynamically updated.

All three steps are computed over each input mini-batch at training time. Samples belonging to the mini-batch are updated in the hash table instantaneously.

## 4.1 Linear auto-encoder

Since the online PCA estimation is in general computationally inefficient and potentially numerically unstable [2], we choose to train a linear autoencoder (AE) paired with  $L_2$  reconstruction loss, as in formulas (2), where  $h(x)$  is the projected subspace of dimensionality  $s$ . In order to not influence the embedding, we do not back-propagate the gradients generated by the  $\mathcal{L}_{AE}$  loss through  $f(x)$ .

$$\begin{aligned} h(x) &= W_1 f(x) + b_1 \\ \hat{f}(x) &= W_2 h(x) + b_2 \\ \mathcal{L}_{AE} &= \|f(x) - \hat{f}(x)\|_2^2 \end{aligned} \quad (2)$$

The AE continuously models the projection that provides the codeword to the hash table update procedure. The added cost of learning such an AE is negligible w.r.t. to the Siamese network training. We extract codewords of all samples in the mini-batch and we use them for hash table update in every step of training. At the same time we update the weights of the AE without backpropagating the gradients through the Siamese architecture. The only BoNs parameter is  $s$ , and its choice is related to two factors: (1) the smaller  $s$ , the more difficult to reconstruct (in the  $L_2$  sense) the original sized embedding and (2) the bigger  $s$ , the larger the number of bins obtained after the binarization, more precisely  $B = 2^s$ . A detailed analysis on the behaviour of BoN as a function of  $s$  is presented in the experimental section.

## 4.2 Dynamic quantization thresholds

Since the lower dimensional space  $h(x)$  changes dynamically during the training, the correct thresholds  $\mu$  for its binarization is estimated as a running mean:  $\mu \leftarrow \beta\mu + (1 - \beta)h(x)$  where  $\beta$  controls how quickly the running average forgets old samples. We noticed that varying  $\beta \in [0.95, 0.999]$  does not influence the results so that it is not a critical parameter to tune.

## 4.3 Hash table dynamic update

We maintain a hash table  $L$  with  $2^s$  bins that, for each entry indexed by an integer  $j$ , contains a collection of images identified by pairs  $(v, ID(v))$ , where  $v$  is an integer that uniquely represents an image in the dataset, and  $ID(v)$  is an integer that uniquely represents the ID for the given image. Also, we keep track of the latest hash entry for every image in the dataset, using integer values, such that  $C[v] = j$ . In such a way, updating the hash table has a very

limited computational cost. The slowest part of the update procedure is removing the tuple  $(v, ID(v))$  from the bin to which it had been previously assigned, and it has a cost of  $\mathcal{O}(N/2^s)$ . In term of memory cost, assuming that both the ID and the sample identifiers  $(ID(v)$  and  $v)$  can be represented with 4 bytes integers, we need only a total of  $4(N + 2N)$  bytes to store both the hash table  $L$  and the list of hash entries  $C$ .

In each training step for each image in the mini-batch  $I \in \{I_1, I_2, \dots, I_m\}$  we extract its codeword  $c(I)$  by binarizing the difference between its latent projection from the autoencoder  $h(I)$  and running mean  $\mu$ . We remove the pair  $(v_I, ID(v_I))$  from the hash bin where it had been stored before  $L[C[v_I]]$  and we add the pair to the new hash bin  $L[c(I)]$ . Finally, we update the list of hash entries  $C[v_I]$  with the new hash entry  $c(I)$ .

## 4.4 Bag of Negatives with random sampling

The simplest way of using BoN is creating mini-batches of  $b$  random anchor-positive image pairs. For each pair, we sample a negative image randomly among the images that belong to the same bin as the anchor. In case the anchor belongs to a bin in which there are no other images from a different ID, we sample the negative image randomly from the whole dataset.

## 4.5 Bag of Negatives with *batch hard* loss

As BoN can provide relevant images for batch sampling, it can be easily combined with a loss such as *batch hard*. It is important to create batches of size  $m$ , which contain  $k$  images from  $l$  IDs when training a model with *batch hard* loss. We set  $k = 2$  for all the experiments, as we are focusing on showing the importance of good negative sampling, and we want to avoid the results being influenced by hard positive samples. First we select a random non-empty bin, that contains images from  $r$  different IDs. If  $r$  is 1, we create a set of  $l$  IDs randomly. If  $r$  is greater or equal than  $l$ , we choose  $l$  IDs randomly from the same bin. Finally, if  $r$  is greater than 1 and smaller than  $l$ , we take all the  $r$  IDs from the chosen bin, and pick another random bin from which we sample the remaining  $l - r$  IDs, if possible. We keep repeating the last step until we have  $l$  IDs. Once we have a set of  $l$  IDs, we create a mini-batch of  $k$  random images that belong to each one of  $l$  selected IDs.

# 5 Empirical evidence

## 5.1 Datasets

**Person re-identification large dataset.** We merged eleven publicly available datasets for person re-identification, CUHK01 [15], CUHK02 [14], 3DPeS [10], VIPeR [7], airport [11], MSMT17 [18], Market-1501 [13], DukeMTMC [20]. The merged dataset has 178,002 images of 10,552 identities. We used both training and testing partitions of all the datasets except for Market-1501 and DukeMTMC-reID and we did not use the images that are labeled as distractors or junk.

**Stanford Online Products [19]** is a retrieval dataset which contains 120,053 images of 22,634 products. The dataset is split into two partitions: the training partition, which contains 59,551 images of 11,318 products, and the testing partition which contains 60,502 of 11,316 classes.

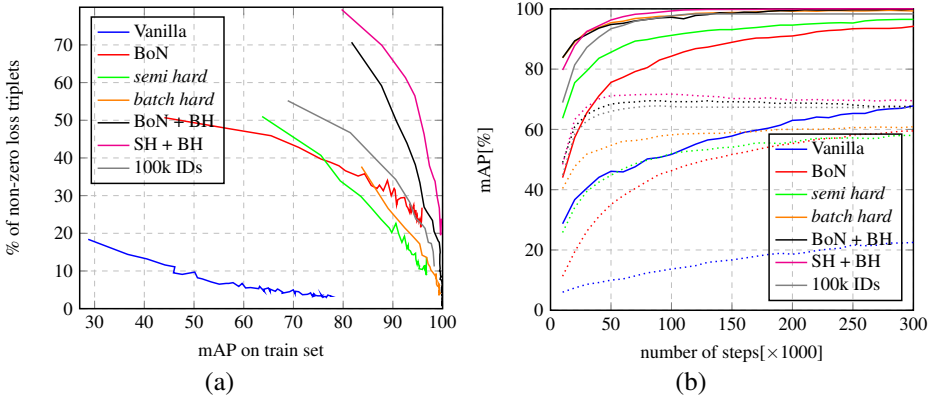


Figure 2: a) Percentage of non-zero loss triplets as a function of mAP on the training set. b) Models evaluated on Training and Market datasets as functions of number of steps.

## 5.2 Pre-trained backbone and training parameters

We use Inception-V3 as a backbone for our model. In particular, we take the convolutional layers and initialize them with weights from a standard network pre-trained on ImageNet. The final descriptors are further max-pooled and  $\ell_2$  normalized. The descriptor size is 2,048. The model is trained using the ADAM optimizer, with the initial learning rate  $10^{-4}$ , and with a learning rate decay of 0.9 each 50k iterations. The images for person re-ID are resized to  $192 \times 384$  pixels, and  $256 \times 256$  for Stanford Online Products. At test time, we extract representations and compare them using dot product.

## 5.3 Analysis of Bag of Negatives

### 5.3.1 Non-zero loss triplets analysis

Figure 2.a shows the percentage of non-zero loss triplets (measured at training time) as a function of the training mean Average Precision (mAP) for the Vanilla sampling, BoN-Random, the *semi hard* loss, the *batch hard* loss, BoN-*batch hard*, Spectral Hashing-*batch hard* and 100k IDs-*batch hard* loss on the person re-ID dataset. The margin is set to  $\alpha = 0.3$ , the mini-batch size is  $m = 48$ , and the leftmost point on the plot for every method is obtained at 10k steps of training in all experiments. As expected, the percentage of non-zero triplets for the Vanilla sampling starts at only 20% and decreases as the mAP increases; at  $\text{mAP} = 77.4$  the non-zero triplets are less than 5% and the training is virtually unable to learn anything else.

BoN-random significantly increases the number of non-zero loss triplets w.r.t. the Vanilla sampling, without modifying the loss nor the way the anchor-positive pairs are formed. The improvement is solely due to the improved sampling of negatives. BoN-random exhibits a behaviour similar to *semi hard* and *batch hard* while providing, in general, more non-zero loss triplets. However, the nature of the improvement provided by our method and *batch hard* is very different: BoN searches for negatives in a local region of the embedding space while *batch hard* forms the triplets seeking non-zero loss triplets in an explicit way within a mini-batch but sampling from the whole embedding.

These two complementary strategies can be easily combined as seen in section 4.5. The



combination inherits the benefits of both approaches: at 10k steps *batch hard* and BoN-*batch hard* have a similar mAP  $\approx 83\%$  but BoN-*batch hard* has about 2 times more non-zero loss triplets, and it has more non-zero loss triplets systematically until the end of the training. As will be seen in the comparison, this behavior not only speeds-up the training, but also provides better triplets, which improves the performance on validation sets.

The combination of Spectral Hashing and *batch hard* requires the following steps: (1) feature extraction on the whole training set, (2) reduction of the feature size by PCA to the size  $s$  ( $s = 18$ ) and (3) hash table construction; we repeat this procedure every 5k steps. Given this hash table, batches are created the same way as explained in 4.5. BoN-*batch hard* shows very similar behavior to the Spectral Hashing - *batch hard*: they both train quickly, obtaining almost the same mAP after 10k steps, with high percentage of non-zero loss triplets. Spectral Hashing - *batch hard* is providing more non-zero loss triplets during the whole training. This is expected, as the hash table is updated at the same moment for all the samples. However, this configuration is far more computationally expensive and does not scale for datasets with a large number of IDs. As an example, each hash table update requires 11.6 minutes of computation every 5k steps, which is the same time required for 1,657 steps of training. In other words, 25% of the training time is spent on hash table update. Thus SH-BH is better, but at a much higher computational cost.

We analyze the behavior of batch creation proposed in [25], using 10 clusters as suggested by the authors. We use these clusters for creating the hash table and we do not update it during the training. In addition to longer training time, this method lacks flexibility in updating the hash table. In other words, samples that are considered relevant negatives to an identity are set at the beginning of the training and are static w.r.t. the training process. Moreover, a possible sub-optimal clustering is going to be seriously detrimental to the training.

### 5.3.2 BoN-Random behaviour varying $s$

It is interesting to note that BoN-Random degenerates to Vanilla sampling for  $s = 0$ . Figure 3.a shows the mAP results on the Market and Duke validation datasets for different values of  $s$  at 200k steps. As it can be seen, the performance increases with  $s$  and it reaches a maximum at  $s = 18$ ; with this value we have  $178k/2^{18} = 0.68$  samples per bin on average. Within a static embedding, and with the standard Spectral Hashing, these figures represent degenerate behaviour in which every bin contains at most one sample, thus defeating the objective of the hashing in providing negative candidates. In reality, the embedding, the linear projection and the thresholds are changing dynamically, such that many bins can be empty. In this case, increasing  $s$  improves the granularity of the hash table without incurring in degenerate hash functions. Nonetheless, with  $s = 22$ , BoN reaches its breaking point and the average number of samples per bins (for non empty bins) is very high, such that BoN-Random starts to perform negative sampling on the whole dataset too frequently.

### 5.3.3 Qualitative results

In Figure 3.b we show images sampled from 2 bins for the person re-identification dataset for BoN-random at 200k steps. As demonstrated, the images belonging to the same bin have some similarity, e.g. same upper body color, same type of clothes etc.



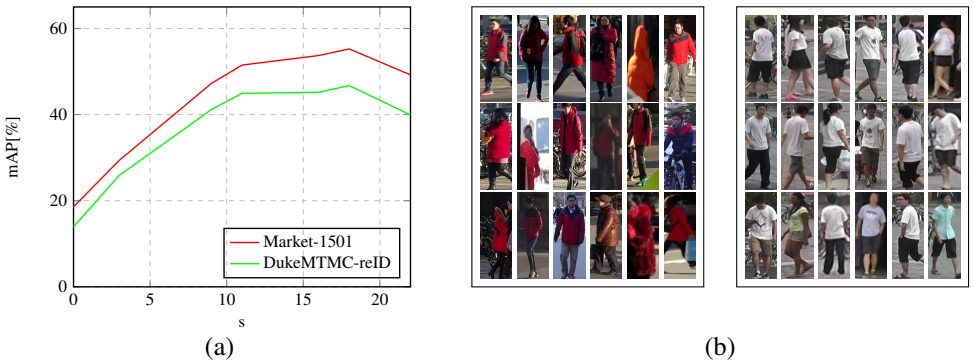


Figure 3: a) Validation mAP as a function of  $s$ , b) Samples from 2 random bins.

## 6 Results and Comparison

In this section we perform a controlled comparison of our proposal with the two most commonly used triplet losses: *semi hard* and *batch hard* and the two batch creation methods: hierarchical tree [5] and 100k IDs [25]. We avoid extra variables (e.g. augmentation, other architectures, etc.) that could mask the empirical results for other reasons not related to negative sampling and triplet construction. For such reasons, we use the same mini-batch size for all the methods, the same pre-trained backbone, the same margin  $\alpha$  and the same embedding size (see subsection 5.2 for the details).

Table 1 shows the results of the comparison on the person re-identification dataset. BoN-random clearly outperforms Vanilla sampling and provides validation mAPs comparable to *semi hard* and *batch hard*. The best results are obtained by Spectral Hashing - *batch hard*, which was expected, taking into account that BoN is an online approximation of Spectral Hashing. The numbers show that the margin between BoN and Spectral Hashing is only 1.5% on average of the two evaluation datasets. However, Spectral Hashing can be used only if the training set is small, so its application on bigger datasets would be infeasible. One can argue that the performance of BoN can be easily reached by just increasing the mini batch size of the *batch hard* method. The "*batch hard* (2x batch)" line shows results from the training in which the mini-batch size is doubled. As expected, in this case, the model trains faster (in terms of number of steps until convergence) and has better performance, but does not outperform BoN-*batch hard*. This experiment shows that BoN is a key component of the accelerated training and improved validation results of BoN-*batch hard*.

We implemented two methods for batch creation known in the literature, Hierarchical Tree (HT) [5] and 100k IDs [25] and combined them with *batch hard*. We followed the procedure described in [5] and computed the distance matrix between all the IDs every 5k steps. We form a batch by randomly selecting one ID, and taking the remaining  $l - 1$  as its closest neighbors. Even though Hierarchical Tree and 100k IDs require a computationally expensive two-phase training, BoN-*batch hard* obtains better results in fewer training steps.

We fine-tuned the BoN-*batch hard* model from the table 1 on Market-1501 and DukeMTMC-reID datasets. To the best of our knowledge, the fine tuned BoN-*batch hard* performs better than all the methods that do not use ad-hoc architectures.

Figure 2.b shows the training (full lines) and validation (dashed lines) performance of the compared methods up to 300k steps. BoN-random exhibits the slower increasing curve in all plots with respect to all methods except the Vanilla sampling. This is due to the fact that the

Table 1: mAP validation results at peak performance for every method. \* stands for the best state-of-the-art result obtained by an ad-hoc architecture.

Method	#steps	Market	Duke
Vanilla	600k	28.1	22.5
BoN-random	440k	61.4	51.3
<i>semi hard</i>	240k	57.0	50.5
<i>batch hard</i>	280k	60.8	53.7
HT [6]- <i>batch hard</i>	310k	65.9	57.5
100k [12]- <i>batch hard</i>	90k	67.8	61.2
BoN- <i>batch hard</i>	<b>80k</b>	<b>69.5</b>	<b>62.1</b>
SH- <i>batch hard</i>	100k	71.7	62.9
<i>batch hard</i> (2x batch)	70k	62.92	56.76
SPReID [11]	-	76.6	63.3
MGN [7]*	-	86.9	78.4
BoN- <i>batch hard</i> - ft	-	77.3	68.6

Table 2: Results on Stanford Online Products [14]. \* stands for the best state-of-the-art result obtained by an ad-hoc architecture.

Method	r1	r10
lifted structured [14]	61.5	80.0
sampling matters [13]	72.7	86.2
hierarchical tree [6]	74.8	88.3
ABE-8 <sup>512</sup> [12]*	76.3	86.4
BoN- <i>batch hard</i>	<b>75.8</b>	<b>88.6</b>

negative is still sampled randomly inside the anchor bin. Nonetheless, it is as good as *semi hard* and *batch hard* on the validation sets at 300k, which means that the negative samples are of sufficiently good quality to provide good generalization capability. BoN-*batch hard* is the fastest method, reaching the over-fitting regime at about 70k-100k steps; this is at least 3 times faster than other methods.

Table 2 shows the comparison of BoN-*batch hard* with state-of-the-art approaches on Stanford Online Products dataset. We trained BoN-*batch hard* using the same methodology as explained in section 5.2 on retrieval, using data augmentation techniques such as random horizontal flipping, blurring, zooming in and out, and cutout. BoN-*batch hard* obtains the state-of-the-art results when compared with all the other methods.

*Batch hard* and BoN are complementary: BoN provides a set of good quality negative samples, while *batch hard* provides the explicit hard negative selection process and the increased number of triplets per mini-batch. The resulting combination is faster and better than both approaches.

## 7 Conclusion

In this paper we introduced Bag of Negatives (BoN), a novel method for accelerated and improved training of Siamese networks that scales linearly on datasets with large numbers of identities. The method is complementary to the popular *batch hard* approach, and their combination provides improved validation results by means of high quality negative candidates. Future work will address a profound analysis of BoN behaviour when combined with other loss functions. Also, a more extensive test on other datasets will be performed. Finally, we will investigate possible solutions on automated estimation of the  $s$  meta-parameter.

## Acknowledgements

This work has been partially supported by the Spanish project RTI2018-095645-B-C21, the grant 2018FI\_B1\_00056 from AGAUR and CERCA, from Generalitat de Catalunya.

## References

- [1] Davide Baltieri, Roberto Vezzani, and Rita Cucchiara. 3dpe: 3d people dataset for surveillance and forensics. In *Proceedings of the 1st International ACM Workshop on Multimedia access to 3D Human Objects*, 2011.
- [2] Hervé Cardot and David Degras. Online principal component analysis in high dimension: Which algorithm to choose? 2018.
- [3] Miguel Á. Carreira-Perpiñán and Ramin Raziperchikolaei. Hashing with binary autoencoders. In *Proc. CVPR*, 2015.
- [4] Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: a deep quadruplet network for person re-identification.
- [5] Weifeng Ge, Weilin Huang, Dengke Dong, and Matthew R Scott. Deep metric learning with hierarchical triplet loss. In *Proc. ECCV*, 2018.
- [6] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *Proc. CVPR*, 2011.
- [7] Douglas Gray and Hai Tao. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *Proc. ECCV*, 2008.
- [8] Yiluan Guo and Ngai-Man Cheung. Efficient and deep person re-identification using multi-level similarity. *arXiv preprint arXiv:1803.11353*, 2018.
- [9] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [10] Mahdi M. Kalayeh, Emrah Basaran, Muhittin Gökmen, Mustafa E. Kamasak, and Mubarak Shah. Human semantic parsing for person re-identification. In *Proc. CVPR*, 2018.
- [11] Srikrishna Karanam, Mengran Gou, Ziyang Wu, Angels Rates-Borras, Octavia Camps, and Richard J. Radke. A systematic evaluation and benchmark for person re-identification: Features, metrics, and datasets, 2016.
- [12] Ira Kemelmacher-Shlizerman, Steven M. Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *Proc. CVPR*, 2016.
- [13] Wonsik Kim, Bhavya Goyal, Kunal Chawla, Jungmin Lee, and Keunjoo Kwon. Attention-based ensemble for deep metric learning. In *Proc. ECCV*, 2018.
- [14] Wei Li and Xiaogang Wang. Locally aligned feature transforms across views. In *Proc. CVPR*, 2013.
- [15] Wei Li, Rui Zhao, and Xiaogang Wang. Human reidentification with transferred metric learning. In *Proc. ACCV*, 2012.
- [16] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *Proc. CVPR*, 2018.

- [17] Hao Liu, Jiashi Feng, Meibin Qi, Jianguo Jiang, and Shuicheng Yan. End-to-end comparative attention networks for person re-identification. In *TIP*, 2017.
- [18] R. Manmatha, Chao-Yuan Wu, Alexander J. Smola, and Philipp Krähenbühl. Sampling matters in deep embedding learning. In *Proc. ICCV*, 2017.
- [19] Hyun Oh Song, Yu Xiang, Stefanie Jegelka, and Silvio Savarese. Deep metric learning via lifted structured feature embedding. In *Proc. CVPR*, 2016.
- [20] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision workshop on Benchmarking Multi-Target Tracking*, 2016.
- [21] M Saquib Sarfraz, Arne Schumann, Andreas Eberle, and Rainer Stiefelhagen. A pose-sensitive embedding for person re-identification with expanded cross neighborhood re-ranking. In *Proc. CVPR*, 2018.
- [22] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proc. CVPR*, 2015.
- [23] Chi Su, Jianing Li, Shiliang Zhang, Junliang Xing, Wen Gao, and Qi Tian. Pose-driven deep convolutional model for person re-identification. In *Proc. ICCV*, 2017.
- [24] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proc. CVPR*, 2014.
- [25] Chong Wang, Xipeng Lan, and Xue Zhang. How to train triplet networks with 100k identities? In *Proc. CVPR Workshop*, 2017.
- [26] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *ACM Multimedia Conference on Multimedia Conference*, 2018.
- [27] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *Proc. CVPR*, 2014.
- [28] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proc. CVPR*, 2018.
- [29] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *Proc. NIPS*, 2009.
- [30] Jing Xu, Rui Zhao, Feng Zhu, Huaming Wang, and Wanli Ouyang. Attention-aware compositional network for person re-identification. *arXiv preprint arXiv:1805.03344*, 2018.
- [31] Haiyu Zhao, Maoqing Tian, Shuyang Sun, Jing Shao, Junjie Yan, Shuai Yi, Xiaogang Wang, and Xiaoou Tang. Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In *Proc. CVPR*, 2017.
- [32] Liming Zhao, Xi Li, Yueting Zhuang, and Jingdong Wang. Deeply-learned part-aligned representations for person re-identification. In *Proc. ICCV*, 2017.

- 
- [33] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proc. ICCV*, 2015.
- [34] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proc. ICCV*, 2017.