FRACKING DEEP CONVOLUTIONAL IMAGE DESCRIPTORS

Edgar Simo-Serra[*]**, Eduard Trulls**[*]
Institut de Robòtica i Informàtica Industrial (CSIC-UPC). Barcelona, Spain
{esimo,etrulls}@iri.upc.edu

Luis Ferraz
Catchoom Technologies S.L. Barcelona, Spain
luis.ferraz@catchoom.com

Iasonas Kokkinos
Ecole Centrale Paris. Chatenay-Malabry, France
iasonas.kokkinos@ecp.fr

Francesc Moreno-Noguer
Institut de Robòtica i Informàtica Industrial (CSIC-UPC). Barcelona, Spain
fmoreno@iri.upc.edu

ABSTRACT

In this paper we propose a novel framework for learning local image descriptors in a discriminative manner. For this purpose we explore a siamese architecture of Deep Convolutional Neural Networks (CNN), with a Hinge embedding loss on the L_2 distance between descriptors. Since a siamese architecture uses pairs rather than single image patches to train, there exist a large number of positive samples and an exponential number of negative samples. We propose to explore this space with a stochastic sampling approach of the training set, in combination with an aggressive mining strategy over both the positive and negative samples which we denote as "fracking". We perform a thorough evaluation of the architecture hyper-parameters, and demonstrate very large performance gains compared to both standard CNN learning strategies and hand-crafted image descriptors like SIFT, up to 2.5x in terms of the area under the Precision-Recall curve.

1 INTRODUCTION

Feature descriptors, i.e. the invariant and discriminative representation of local image patches, is a major research topic in computer vision. The field reached maturity with SIFT (Lowe, 2004), and has since become the cornerstone of a wide range of applications in recognition and registration. While most descriptors use hand-crafted features (Lowe, 2004; Bay et al., 2006; Kokkinos et al., 2012; Trulls et al., 2013; Simo-Serra et al., 2015), there has recently been interest in using machine learning algorithms to learn descriptors from large databases.

In this paper we draw inspiration on the recent success of Deep Convolutional Neural Networks on large-scale image classification problems (Krizhevsky et al., 2012; Szegedy et al., 2013) to build discriminative descriptors for local patches. Specifically, we propose an architecture based on a siamese structure of two CNNs that share the parameters. We compute the L_2 norm on their output, i.e. the descriptors, and use a loss that enforces the norm to be small for corresponding patches and large otherwise. We demonstrate that this approach allows us to learn compact and discriminative representations.

To implement this approach we rely on the dataset of Brown et al. (2011), which contains over 1.5M grayscale 64×64 image patches from different views of 500K different 3D points. With such large datasets it becomes intractable to exhaustively explore all corresponding and non-corresponding pairs. Random sampling is typically used; however, most correspondences are not useful and hinder the learning of a discriminant mapping. We address this issue with aggressive mining of "hard" positives and negatives, a strategy that we denote as "fracking", and which proves fundamental in order to obtain discriminative learned descriptors. In particular, in some of the tests we obtain up to a 169% increase in performance with SIFT as a baseline.

---

[*]Both first authors contributed equally.

2 RELATED WORK

Recent developments in the design of local image descriptors are moving from carefully-engineered features (Lowe, 2004; Bay et al., 2006) to learning features from large volumes of data. This line of works includes unsupervised techniques based on hashing as well as supervised approaches using Linear Discriminant Analysis (Brown et al., 2011; Gong et al., 2012; Strecha et al., 2012), boosting (Trzcinski et al., 2013), and convex optimization (Simonyan et al., 2014).

In this paper we explore solutions based on deep convolutional networks (CNNs). CNNs have been used in computer vision for decades, but are currently experiencing a resurgence kickstarted by the accomplishments of Krizhevsky et al. (2012) on large-scale image classification. The application of CNNs to the problem of descriptor learning has already been explored by some researchers (Jahrer, 2008; Osendorfer et al., 2013). These works are however preliminary, and many open questions remain regarding the practical application of CNNs for learning descriptors, such as the most adequate network architectures and application-dependent training schemes. In this paper we aim to provide a rigorous analysis of several of these topics. In particular, we use a siamese network (Bromley et al., 1994) to train the models, and experiment with different network configurations inspired by the state of the art in deep learning.

Additionally, we demonstrate that aggressive mining of both "hard" positive and negative matching pairs greatly enhances the learning process. Mining hard negatives is a well-known procedure in sliding-window detectors (Felzenszwalb et al., 2010), where the number of negative samples is virtually unlimited and yet most negatives are easily discriminated. Similar techniques have been applied to CNNs for object detection (Girshick et al., 2014; Szegedy et al., 2013).

3 LEARNING DEEP DESCRIPTORS

Given an intensity patch $\mathbf{x} \in \mathbb{R}^N$, the descriptor of $\mathbf{x}$ is a non-linear mapping $D(\mathbf{x})$ that is expected to be discriminative, i.e. descriptors for image patches corresponding to the same point should be similar, and dissimilar otherwise.

In the context of multiple-view geometry, descriptors are typically computed for salient points where scale and orientation can be reliably estimated, for invariance. Patches then capture local projections of 3D scenes. Let us consider that each image patch $\mathbf{x}_i$ has an index p_i that uniquely identifies the 3D point which roughly projects onto the 2D patch, from a specific viewpoint. Therefore, taking the L_2 norm as a similarity metric between descriptors, for an ideal descriptor we would wish that

$$d_D(\mathbf{x}_1, \mathbf{x}_2) = \|D(\mathbf{x}_1) - D(\mathbf{x}_2)\|_2 = \begin{cases} 0 & \text{if } p_1 = p_2 \\ \infty & \text{if } p_1 \neq p_2 \end{cases} \quad (1)$$

We propose learning descriptors using a siamese network (Bromley et al., 1994), i.e. optimizing the model for pairs of corresponding or non-corresponding patches, as shown in Fig. 1-(a). We propagate the patches through the model to extract the descriptors and then compute their L_2 norm, which is a standard similarity measure for image descriptors. We then compute the loss function on this distance. Given a pair of patches $\mathbf{x}_1$ and $\mathbf{x}_2$ we define a loss function of the form

$$l(\mathbf{x}_1, \mathbf{x}_2, \delta) = \delta \cdot l_P(d_D(\mathbf{x}_1, \mathbf{x}_2)) + (1 - \delta) \cdot l_N(d_D(\mathbf{x}_1, \mathbf{x}_2)) \quad (2)$$

where δ is the indicator function, which is 1 if $p_1 = p_2$, and 0 otherwise. l_P and l_N are the partial loss functions for patches corresponding to the same 3D point and to different points, respectively. When performing back-propagation, the gradients are independently accumulated for both descriptors, but jointly applied to the weights, as they are shared.

Although it would be ideal to optimize directly for Eq. (1), we relax it, using a margin m for $l_N(\cdot)$. In particular, we consider the hinge embedding criterion (Mobahi et al., 2009)

$$l_P(d_D(\mathbf{x}_1, \mathbf{x}_2)) = d_D(\mathbf{x}_1, \mathbf{x}_2) \quad \text{and} \quad l_N(d_D(\mathbf{x}_1, \mathbf{x}_2)) = \max(0, m - d_D(\mathbf{x}_1, \mathbf{x}_2)) \quad (3)$$

3.1 CONVOLUTIONAL NEURAL NETWORK DESCRIPTORS

When designing the structure of the CNN we are limited by the size of the input data, in our case 64×64 patches from the dataset of Brown et al. (2011). Note that larger patches would allow us
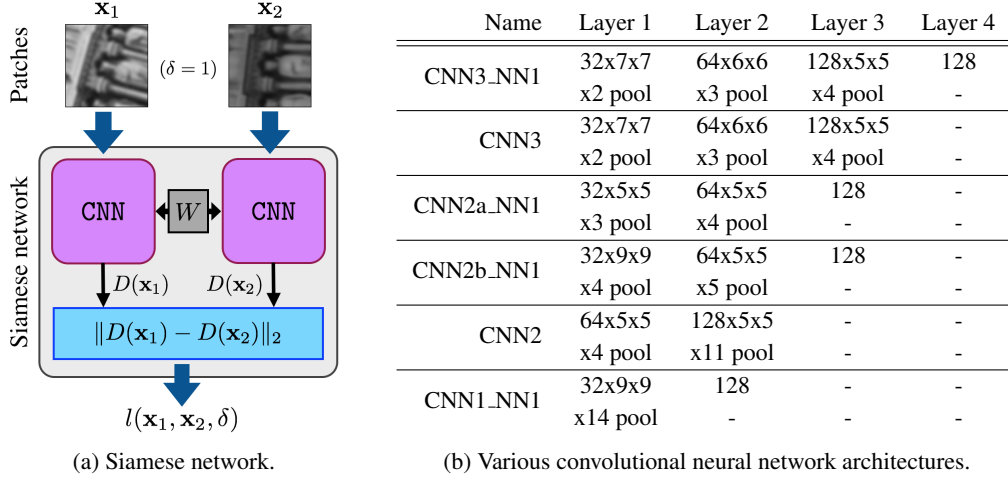
(a) Siamese network.

Name	Layer 1	Layer 2	Layer 3	Layer 4
CNN3_NN1	32x7x7	64x6x6	128x5x5	128
	x2 pool	x3 pool	x4 pool	-
CNN3	32x7x7	64x6x6	128x5x5	-
	x2 pool	x3 pool	x4 pool	-
CNN2a_NN1	32x5x5	64x5x5	128	-
	x3 pool	x4 pool	-	-
CNN2b_NN1	32x9x9	64x5x5	128	-
	x4 pool	x5 pool	-	-
CNN2	64x5x5	128x5x5	-	-
	x4 pool	x11 pool	-	-
CNN1_NN1	32x9x9	128	-	-
	x14 pool	-	-	-

(b) Various convolutional neural network architectures.

Figure 1: **Left:** Schematic of a siamese network, where pairs of input patches are processed by two copies of the same CNN. **Right:** Different CNN configurations evaluated in this paper.

to consider deeper networks, and possibly more informative descriptors, but at the same time they would be also more susceptible to occlusions. We consider networks of up to three convolutional layers, followed by up to a single additional fully-connected layer. We target descriptors of size 128, the same as SIFT (Lowe, 2004); this value also constrains the architectures we can explore.

As usual, each convolutional layer consists four sub-layers: filter layer, non-linearity layer, pooling layer and normalization layer. Since sparser connectivity has been shown to improve performance while lowering parameters and increasing speed (Culurciello et al., 2013), except for the first layer, the filters are not densely connected to the previous layers. Instead, they are sparsely connected at random, so that the mean number of connections each input layer has is constant.

Regarding the non-linear layer, we use hyperbolic tangent (Tanh), as we found it performs better than Rectified Linear Units (ReLU). We use L_2 pooling for the pooling sublayers, which were shown to outperfom the more standard max pooling (Sermanet et al., 2012). Normalization has been shown to be important for deep networks (Jarrett et al., 2009) and fundamental for descriptors (Mikolajczyk & Schmid, 2005). We use subtractive normalization for a 5×5 neighbourhood with a Gaussian kernel. We will justify these decisions empirically in Sec. 4.

An overview of the architectures we consider is given in Fig. 1-(b). We choose a set of six networks, from 2 up to 4 layers. The architecture hyperparameters (number of layers and convolutional/pooling filter size) are chosen so that no padding is needed. We consider models with a final fully-connected layer as well as fully convolutional models, where the last sublayer is a pooling layer. Our implementation is based on Torch7 (Collobert et al., 2011).

3.2 STOCHASTIC SAMPLING STRATEGY AND "FRACKING"

Our goal is to optimize the network parameters from an arbitrarily large set of training patches. Let us consider a dataset with N patches and $M \leq N$ unique 3D patch indices, each with n_i associated image patches. Then, the number of matching image patches or positives N_P and the number of non-matching images patches or negatives N_N in the dataset is

$$N_P = \sum_{i=1}^{M} \frac{n_i(n_i - 1)}{2} \qquad \text{and} \qquad N_N = \sum_{i=1}^{M} n_i(N - n_i) \qquad (4)$$

In general both N_P and N_N are intractable to exhaustively iterate over. We approach the problem with random sampling. For gathering positives samples we can randomly choose a set of B_P 3D point indices $\{p_1, \cdots, p_{B_P}\}$, and choose two patches with corresponding 3D point indices randomly. For negatives it is sufficient to choose B_N random pairs with non-matching indices.

However, when the pool of negative samples is very large random sampling will produce many negatives with a very small loss, which do not contribute to the global loss, and thus stifle the learning
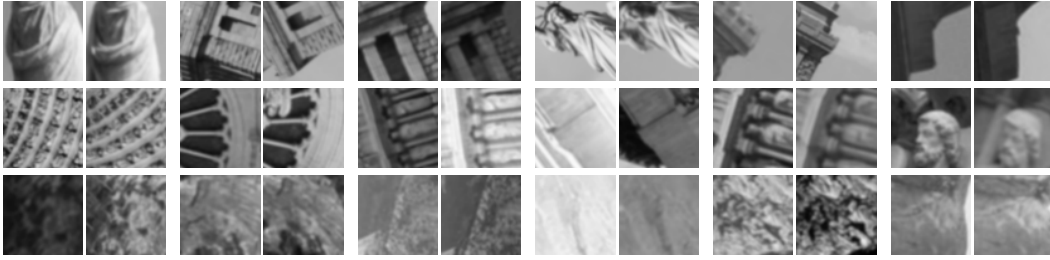
3

Figure 2: Pairs of corresponding samples from the Multi-view Stereo Correspondence dataset. Top row: 'Liberty' (LY). Middle row: 'Notre Dame' (ND). Bottom row: 'Yosemite' (YO).

process. Instead, we can iterate over non-corresponding patch pairs to search for "hard" negatives, i.e. with a high loss. In this manner it becomes feasible to train discriminative models faster while also increasing performance. This technique is commonplace in sliding-window classification.

Therefore, at each epoch we generate a set of B_N randomly chosen patch pairs, and after forward-propagation through the network and computing their loss we keep only a subset of the B_N^M "hardest" negatives, which are back-propagated through the network in order to update the weights. Additionally, the same procedure can be used over the positive samples, i.e. we can sample B_P corresponding patch pairs and prune them down to the B_P^M "hardest" positives. We show that the combination of aggressively mining positive and negative patch pairs, which we call "fracking", allows us to greatly improve the discriminative capability of learned descriptors. Note that extracting descriptors with the learned models does not further require the siamese network and does not incur the computational costs related to mining.

3.3 LEARNING

We normalize the dataset by subtraction of the mean of the training patches and division by their standard deviation. We then learn the weights by performing stochastic gradient descent. We use a learning rate that decreases by an order of magnitude every fixed number of iterations. Additionally, we use standard momentum in order to accelerate the learning process. We use a subset of the data for validation, and stop training when the metric we use to evaluate the learned models converges. Due to the exponentially large pool of positives and negatives available for training and the small number of parameters of the architectures, no techniques to cope with overfitting are used. The particulars of the learning procedure are detailed in the following section.

4 RESULTS

For evaluation we use the Multi-view Stereo Correspondence dataset (Brown et al., 2011), which consists of 64×64 grayscale image patches sampled from 3D reconstructions of the Statue of Liberty (LY), Notre Dame (ND) and Half Dome in Yosemite (YO). Patches are extracted using the Difference of Gaussians detector (Lowe, 2004), and determined as a valid correspondence if they are within 5 pixels in position, 0.25 octaves in scale and $\pi/8$ radians in angle. Fig. 2 shows some samples from each set, which contain significant changes in position, rotation and illumination conditions, and often exhibit very noticeable perspective changes.

We join the data from LY and YO to form a training set with over a million patches. Out of these we reserve a subset of 10,000 unique 3D points for validation (roughly 30,000 patches). The resulting training dataset contains 1,133,525 possible positive patch combinations and 1.117×10^{12} possible negative combinations. This skew is common in correspondence problems such as stereo or structure from motion; we address it with aggressive mining or "fracking".

A popular metric for classification systems is the Receiving Operator Characteristic (ROC), used e.g. in (Brown et al., 2011), which can be summarized by its area under the curve (AUC). However, ROC curves can be misleading when the number of positive and negative samples are very different (Davis & Goadrich, 2006), and is already nearly saturated for the baseline descriptor, SIFT (see Sec. 6). A richer metric is the Precision-Recall curve (PR). We benchmark our models with PR curves and their AUC. In particular, for each of the 10,000 unique points in the validation set we randomly sample two corresponding patches and 1,000 non-corresponding patches, and use them to compute

Architecture	Parameters	PR AUC
SIFT	—	0.361
CNN1_NN1	68,352	0.032
CNN2	27,776	0.379
CNN2a_NN1	145,088	0.370
CNN2b_NN1	48,576	0.439
CNN3_NN1	62,784	0.289
CNN3	46,272	**0.558**

Table 1: Effect of network depth, and fully convolutional networks vs networks with a fully-connected layer. PR AUC on the validation set for the top-performing iteration.

Architecture	PR AUC
SIFT	0.361
CNN3	**0.558**
CNN3 ReLU	0.442
CNN3 No Norm	0.511
CNN3 MaxPool	0.420

Table 2: Fully convolutional CNN3 models with Tanh and ReLU, without normalization, and with max pooling instead of L_2 pooling. The best results are obtained for Tanh units, normalization, and L_2 pooling (i.e. 'CNN3'). PR AUC on the validation set for the top-performing iteration.
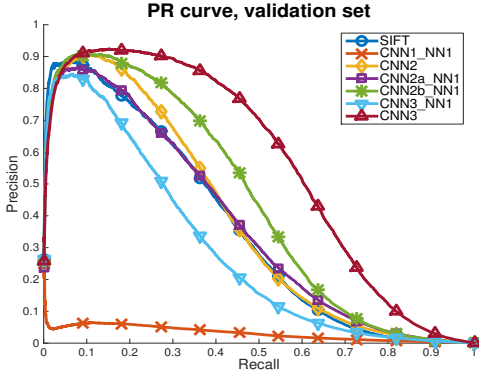
Figure 3: Effect of network depth (up to 3 CNN layers), and fully convolutional networks vs networks with final fully-connected layer (NN1). Fully-convolutional models outperform models with fully-connected neural network at the end.
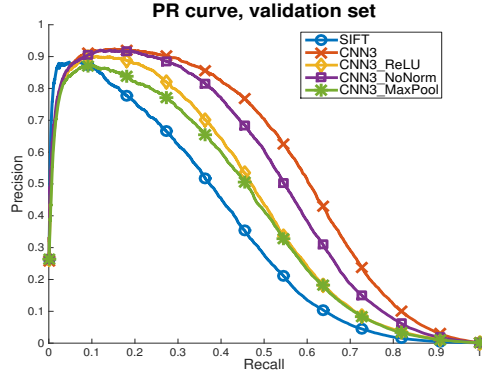
Figure 4: Fully convolutional CNN3 models with Tanh and ReLU, without normalization, and with max pooling instead of L_2 pooling. The best results are obtained for Tanh units, normalization, and L_2 pooling ('CNN3' model).

the PR curve. We rely on the validation set for the LY+YO split to examine different configurations, network architectures and mining techniques; these results are presented in Secs. 4.1-4.4.

Finally, we evaluate the top-performing models over the test set in Sec. 4.5. We follow the same procedure as for validation, compiling the results for 10,000 points with 1,000 non-corresponding matches each, now over 10 different folds (see Sec. 6 for details). We run three different splits, for generalization: LY+YO (tested on ND), LY+ND (tested on YO), and YO+ND (tested on LY).

We will consider all hyperparameters to be the same unless otherwise mentioned, i.e. a learning rate of 0.01 that decreases ever $10,000$ iterations by a factor of 10. We consider negative mining with $B_N = 256$ and $B_N^M = 128$, and no positive mining; i.e. $B_P = B_P^M = 128$.

4.1 DEPTH AND FULLY CONVOLUTIONAL ARCHITECTURES

The network depth is constrained by the size of the patch. We consider only up to 3 convolutional layers (CNN1-3). Additionally, we consider adding a single fully-connected layer at the end (NN1). Fully-connected layers increase the number of parameters by a large factor, which increases the difficulty of learning and can lead to overfitting. We show the results of the various architectures we evaluate in Table 1 and Fig. 3. Deeper networks outperform shallower ones, and architectures with a fully-connected layer at the end do worse than fully convolutional architectures. In the following experiments with consider only models with 3 convolutional layers.

R_P	R_N	Cost	PR AUC
1	1	—	0.366
1	2	20%	0.558
2	2	35%	0.596
4	4	48%	0.703
8	8	67%	**0.746**
16	16	80%	0.538

Table 3: Mining factors indicate the samples considered (B_P, B_N), i.e. 1: 128, 2: 256, 4: 512, 8: 1024, and 16: 2048, of which 128 are used for training. Column 3 indicates the fraction of the computational cost spent mining.

Architecture	Output	Parameters	PR AUC
SIFT	128D	—	0.361
CNN3	128D	46,272	**0.596**
CNN3 Wide	128D	110,496	0.552
CNN3_NN1	128D	62,784	0.456
CNN3_NN1	32D	50,400	0.389

Table 4: Effect of the size of the filters in the CNN3 model, and the fully-connected layers at the end. The best results are obtained with fully-convolutional networks.
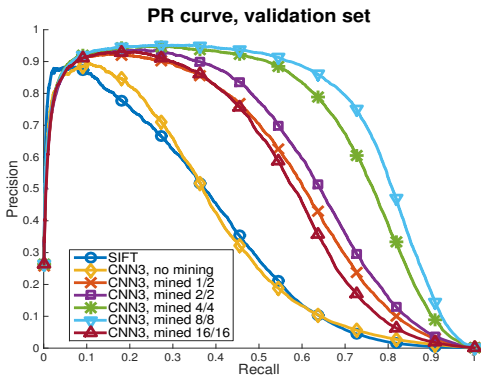
Figure 5: Effect of mining. The best results are obtained with 8/8 factors, i.e. pools of 1024 samples filtered down to 128.
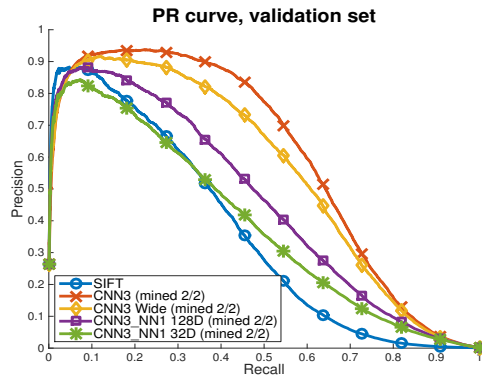
Figure 6: Filter size and fully-connected layers for dimensionality reduction. The best results are obtained with fully-convolutional networks.

4.2 HIDDEN UNITS MAPPING, NORMALIZATION, AND POOLING

It is generally accepted that Rectified Linear Units (ReLU) perform much better in classification tasks (Krizhevsky et al., 2012) than other non-linear functions. They are, however, ill-suited for regression tasks (as the problem in hand), as they can only output positive values. We consider both the standard Tanh and ReLU. For the ReLU case we still use Tanh for the last layer. We also consider not using the normalization sublayer for each of the convolutional layers. Finally, we consider using max pooling rather than L_2 pooling. We show results for the fully-convolutional CNN3 architecture in Table 2 and Fig. 4. The best results are obtained with Tanh, normalization and L_2 pooling ('CNN3' in the table/plot). We will use this configuration in the following experiments.

4.3 FRACKING

We analyze the effect of both positive and negative mining, or "fracking", by training different models in which a large, initial pool of B_P positives and B_N negatives are pruned to a smaller number of 'hard' positive and negative matches, used to update the parameters of the network. We keep this batch size for parameter learning constant, with $B_N^M = 128$ and $B_P^M = 128$, and increase the ratio of both negative mining $R_N = B_N/B_N^M$ and positive mining $R_P = B_P/B_P^M$. We keep all other parameters constant. We use the notation R_P/R_N, for brevity.

Large mining factors have a high computational cost, up to 80% of the total computational cost, which includes mining (i.e. forward propagation of all B_P and B_N samples) and learning (i.e. backpropagating the "hard" positive and negative samples). In order to speed up the learning process we initialize the CNN3 models with positive mining, i.e. 2/2, 4/4, 8/8 and 16/16, with an early iteration of a model trained only with negative mining (1/2).

6

Train	Test	SIFT	CNN3 mine-1/2	CNN3 mine-2/2	CNN3 mine-4/4	CNN3 mine-8/8	PR AUC Increase (Best vs SIFT)
LY+YOS	ND	0.349	0.535	0.555	0.630	**0.667**	91.1%
LY+ND	YOS	0.425	0.383	0.390	0.502	**0.545**	28.2%
YOS+ND	LY	0.226	0.460	0.483	0.564	**0.608**	169.0%

Table 5: PR AUC for the generalized results over the three dataset splits.
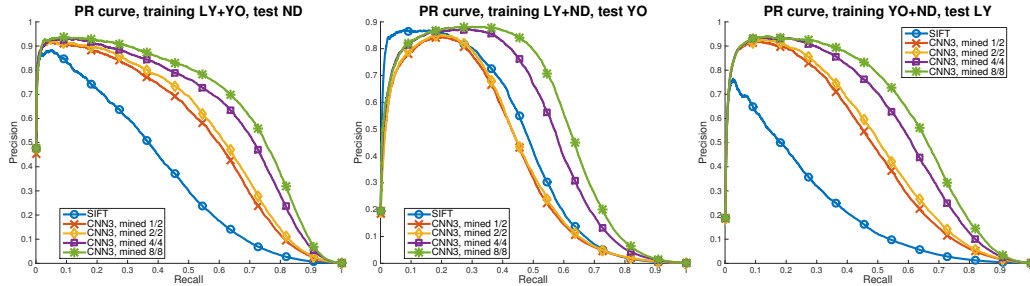
Figure 7: PR curves for the generalized results over the three dataset splits.

Results are shown in Table 3 and Fig. 5. We see that for this particular problem aggressive mining is fundamental. This is likely due to the extremely large number of both negatives and positives in the dataset, in combination with models with a low number of parameters. We observe a drastic increase in performance up to 8/8 mining factors.

4.4 NUMBER OF FILTERS AND DESCRIPTOR DIMENSION

We analyze increasing the number of filters in the CNN3 model, and adding a fully-connected layer that can be used to decrease the dimensionality of the descriptor. We consider increasing the number of filters in layers 1 and 2 from 32 and 64 to 64 and 96, respectively. Additionally, we double the number of internal connections between layers. This more than doubles the number of parameters in this network. To analyze descriptor dimensions we consider the CNN3_NN1 model and change the number of outputs in the last fully-connected layer from 128 to 32. In this case we consider positive mining with $B_P = 256$ (i.e. 2/2). Results can be seen in Table 6 and Fig. 6. The best results are obtained with smaller filters and fully-convolutional networks. Additionally we notice that mining is also instrumental for models the NN1 layer (compare results with Table 1).

4.5 GENERALIZATION

In this section we consider the three dataset splits. We train the best performing models, i.e. CNN3 with different mining ratios, on a combination of two sets, and test them on the remaining set. We select the training iteration that performs best over the validation set. The test datasets are very large (up to 633K patches) and we use the same procedure as for validation, evaluating 10,000 unique points, each with 1,000 random non-corresponding matches. We repeat this process over 10 folds, thus considering 100,000 sets of one corresponding patch vs 1,000 non-corresponding patches. We show results in terms of PR AUC in Table 5, and the corresponding PR curves are pictured in Fig. 7.

We report consistent improvements over the baseline, i.e. SIFT. The performance varies significantly from split to split; this is due to the nature of the different sets. 'Yosemite' contains mostly fronto-parallel translations with illumination changes and no occlusions (Fig. 2, row 3); SIFT performs well on this data. Our learned descriptors outperform SIFT on the high-recall regime (over 20% of the samples; see Fig. 7), and is 28% better overall in terms of PR AUC. The effect is much more dramatic on 'Notredame' and 'Liberty', which contains significant patch translation and rotation, as well as viewpoint changes around outcropping non-convex objects which result in occlusions (see Fig. 2, rows 1-2). Our learned descriptors outperform SIFT by 91% and 169% testing over ND and LY, respectively.
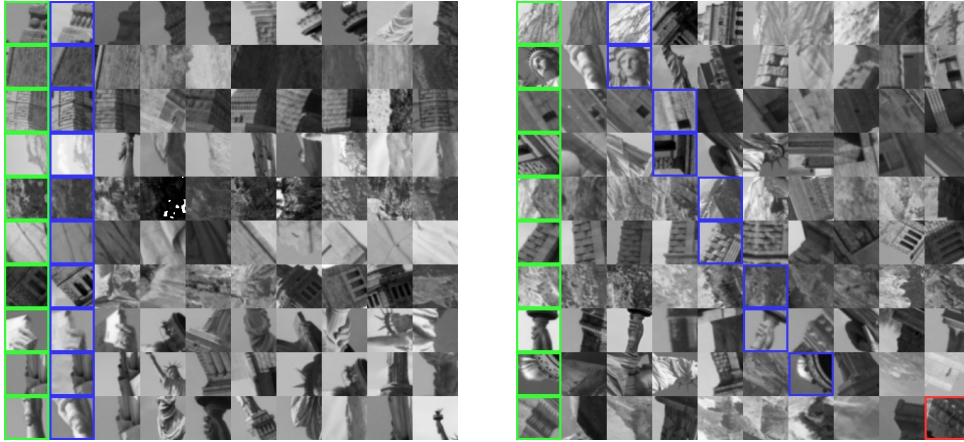
Figure 8: Samples of matches retrieved with our descriptor. Each row depicts the reference patch (in green) and the top matching candidates, sorted from left to right by decreasing similarity. The ground truth match is highlighted in blue. **Left:** Examples where the ground truth match is retrieved in the first position. **Right:** Examples in which the ground truth match is ranked at positions 2-6. The last row shows a failure case, highlighted in red, where the correct match is ranked 632/1000.

4.6   QUALITATIVE ANALYSIS

Fig. 8 shows samples of matches retrieved with our CNN3-mined-4/4 network, over the validation set for the first split. In this experiment the corresponding patches were ranked in the first position in 76.5% of cases; a remarkable result, considering that every true match had to be chosen from a pool of 1,000 false correspondences. The right-hand image shows cases where the ground truth match was not ranked first; notice that most of these patches exhibit significant changes of appearance. We include a failure case (highlighted in red), caused by a combination of large viewpoint and illumination changes; these misdetections are however very uncommon.

5   CONCLUSIONS

We use siamese networks to train deep convolutional models for the extraction of image descriptors for correspondence matching. This problem typically involves small patches and large volumes of data. The former constrains the size of the network, limiting the discriminating power of the models, while the latter makes it intractable to exhaustibly iterate over all the training samples. We address this problem with a novel training scheme based on aggressive mining of both positive and negative correspondences.

Current research in convolutional neural networks for computer vision is focused on classification rather than regression. Despite previous research in this area (Jahrer, 2008; Osendorfer et al., 2013) it remains unclear what is the most adequate architecture for the problem at hand. This is a critical point, as evidenced by the dramatic effect Krizhevsky et al. (2012) had in re-invigorating the field. In this paper we investigate a wide range of architectures (filter size, Tanh and ReLU units, normalization), and consider both fully-convolutional networks and networks with a fully-connected neural network at the end.

We show that the combination of state-of-the-art techniques with aggressive mining in the training stage can result in large performance gains, up to over 2.5x the performance of SIFT in terms of the Precision-Recall AUC. Our experiments over different data splits suggest that learning descriptors is particularly relevant for hard correspondence problems (e.g. 'Liberty' and 'Notredame' sets). The best models are fully convolutional.

We identify multiple directions for future research. Color has proved informative in deep learning for detection problems; our current models are built on grayscale data due to data restrictions. Our networks are likewise restricted by the size of the patches (64×64 pixels), so that we currently limit our models to three convolutional layers. We intend to explore the performance of deeper networks along with larger patches. Additionally, our study indicates that fully-convolutional models outperform models with a fully-connected neural network at the end, but we intend to study this

problem in further detail, particularly with dimensionality reduction in mind. Finally, we have not compared our CNN-based models to other learned descriptors, e.g. (Simonyan et al., 2014), due to time constraints, but we intend to do so in the future.

REFERENCES

Bay, H., Tuytelaars, T., and Van Gool, L. SURF: Speeded Up Robust Features. In *ECCV*, 2006.

Bromley, J., Guyon, I., Lecun, Y., Sckinger, E., and Shah, R. Signature verification using a "siamese" time delay neural network. In *NIPS*, 1994.

Brown, M., Hua, Gang, and Winder, S. Discriminative learning of local image descriptors. *PAMI*, 33(1):43–57, 2011.

Collobert, R., Kavukcuoglu, K., and Farabet, C. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, 2011.

Culurciello, E., Jin, J., Dundar, A., and Bates, J. An analysis of the connections between layers of deep neural networks. *CoRR*, abs/1306.0152, 2013.

Davis, J. and Goadrich, M. The relationship between PR and ROC curves. In *ICML*, 2006.

Felzenszwalb, P., Girshick, R., McAllester, D., and Ramanan, D. Object detection with discriminatively trained part-based models. *PAMI*, 32(9):1627–1645, 2010.

Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.

Gong, Y., Lazebnik, S., Gordo, A., and Perronnin, F. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. In *PAMI*, 2012.

Jahrer, M., Grabner M. Bischof H. Learned local descriptors for recognition and matching. In *Computer Vision Winter Workshop*, 2008.

Jarrett, K., Kavukcuoglu, K., Ranzato, M., and LeCun, Y. What is the best multi-stage architecture for object recognition? In *ICCV*, 2009.

Kokkinos, I., Bronstein, M., and Yuille, A. Dense scale-invariant descriptors for images and surfaces. In *INRIA Research Report 7914*, 2012.

Krizhevsky, A., Sutskever, I., and Hinton, G. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

Lowe, D. Distinctive image features from scale-invariant keypoints. *IJCV*, 60:91–110, 2004.

Mikolajczyk, K. and Schmid, C. A performance evaluation of local descriptors. *PAMI*, 27(10): 1615–1630, 2005.

Mobahi, H., Collobert, R., and Weston, J. Deep learning from temporal coherence in video. In *ICML*, 2009.

Osendorfer, C., Bayer, J., Urban, S., and van der Smagt, P. Convolutional neural networks learn compact local image descriptors. In *ICONIP*, volume 8228. 2013.

Sermanet, P., Chintala, S., and LeCun, Y. Convolutional neural networks applied to house numbers digit classification. In *ICPR*, 2012.

Simo-Serra, E., Torras, C., and Moreno-Noguer, F. DaLI: deformation and light invariant descriptor. *IJCV*, 2015.

Simonyan, K., Vedaldi, A., and Zisserman, A. Learning local feature descriptors using convex optimisation. *PAMI*, 2014.

Strecha, C., Bronstein, A., Bronstein, M., and Fua, P. Lda-hash: Improved matching with smaller descriptors. In *PAMI*, volume 34, 2012.

Szegedy, C., Toshev, A., and Erhan, D. Deep neural networks for object detection. In *NIPS*, 2013.

Trulls, E., Kokkinos, I., Sanfeliu, A., and Moreno-Noguer, F. Dense segmentation-aware descriptors. *CVPR*, 2013.

Trzcinski, T., Christoudias, M., Fua, P., and Lepetit, V. Boosting binary keypoint descriptors. In *CVPR*, 2013.

6 SUPPLEMENTAL MATERIAL

This section contains supplemental material. As we argue in Sec. 4, Precision-Recall (PR) curves are the most appropriate metric for our problem; however, we also consider ROC curves and Cumulative Match Curves (CMC).

ROC curves are created by plotting the true positive rate TPR as a function of the true negative rate TNR, where:

$$\text{TPR} = \frac{\text{TP}}{\text{P}} \qquad \text{TNR} = 1 - \frac{\text{FP}}{\text{N}} \tag{5}$$

Alternatively, the CMC curve is created by plotting the Rank against the Ratio of correct matches. That is, CMC(k) is the fraction of correct matches that have rank$\leq$k. In particular CMC(1) is the percentage of examples in which the ground truth match is retrieved in the first position.

We report these results for either metric in terms of the curves (plots) and their AUC (tables), for the best-performing iteration.

6.1 EXPERIMENTS: (1) DEPTH AND ARCHITECTURES

We extend the results of Sec. 4.1, which are summarized in Table 6. Figs. 9, 10 and 11 show the PR, ROC and CMC curves respectively.

Architecture	Parameters	PR AUC	ROC AUC	CMC AUC
SIFT	—	0.361	0.944	0.953
CNN1_NN1	68,352	0.032	0.929	0.929
CNN2	27,776	0.379	0.971	0.975
CNN2a_NN1	145,088	0.370	**0.987**	**0.988**
CNN2b_NN1	48,576	0.439	0.985	0.986
CNN3_NN1	62,784	0.289	0.980	0.982
CNN3	46,272	**0.558**	0.986	0.987

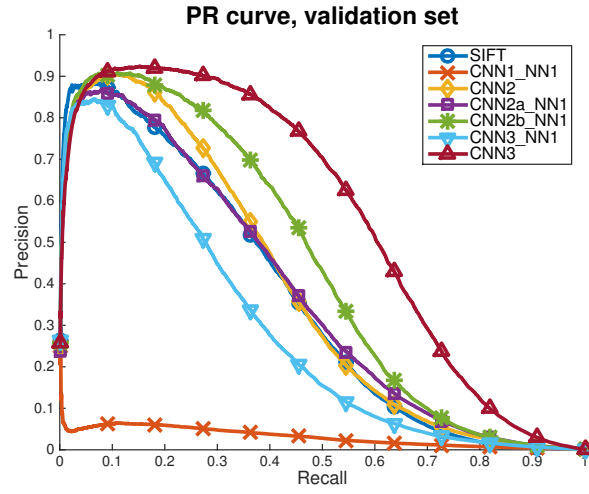Table 6: Extended table for the experiments of Sec. 4.2.

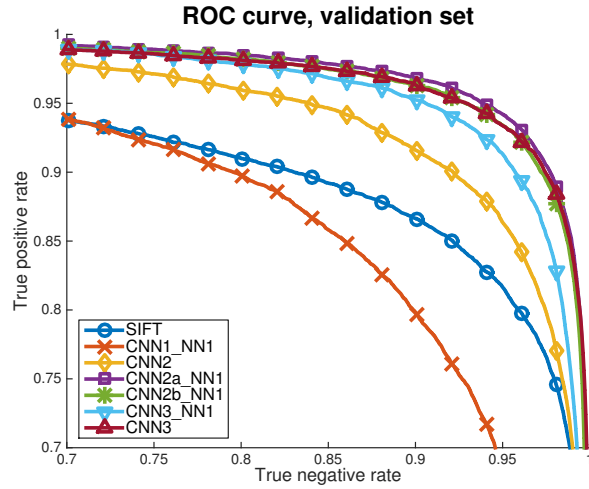Figure 9: PR curves for the experiments of Sec. 4.1 (repeated).

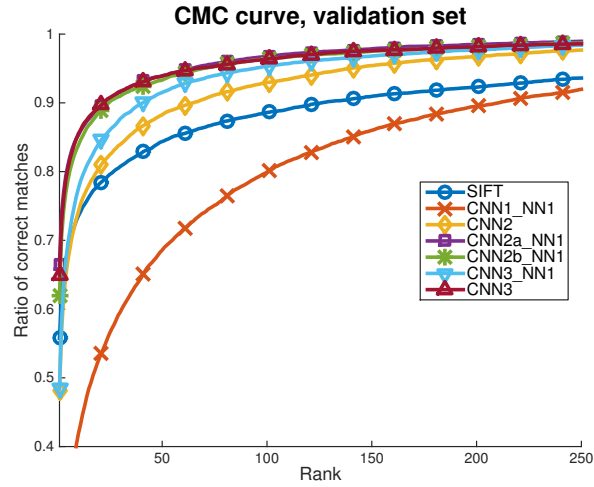Figure 10: ROC curves for the experiments of Sec. 4.1.

Figure 11: CMC curves for the experiments of Sec. 4.1.

6.2 EXPERIMENTS: (2) HIDDEN UNITS, NORMALIZATION AND POOLING

We extend the results of Sec. 4.2, which are summarized in Table 7. Figs. 12, 13 and 14 show the PR, ROC and CMC curves respectively.

Architecture	PR AUC	ROC AUC	CMC AUC
SIFT	0.361	0.944	0.953
CNN3	**0.558**	**0.986**	**0.987**
CNN3 ReLU	0.442	0.973	0.976
CNN3 No Norm	0.511	0.980	0.982
CNN3 MaxPool	0.420	0.973	0.975

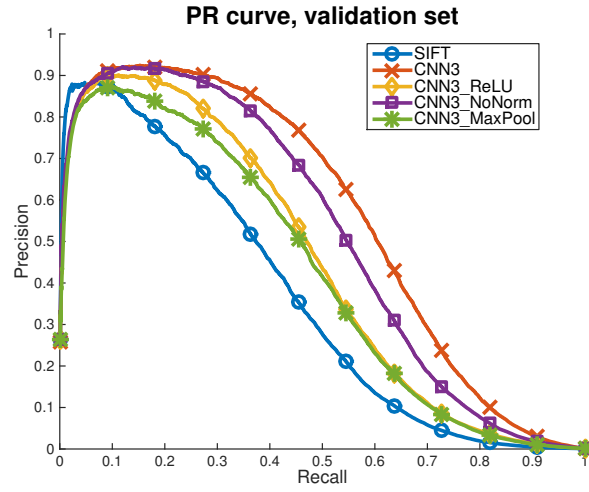Table 7: Extended table for the experiments of Sec. 4.2.

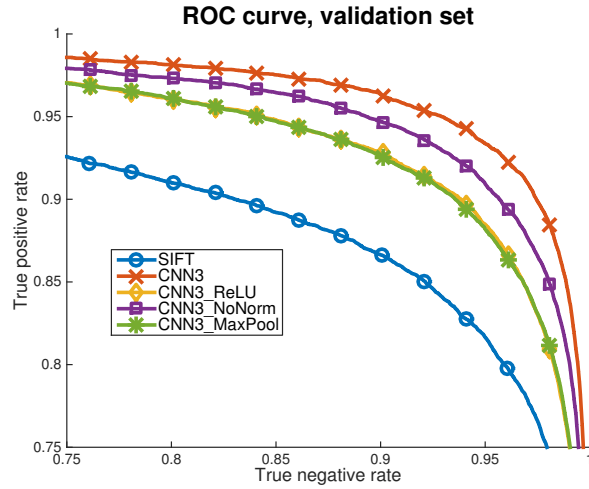Figure 12: PR curves for the experiments of Sec. 4.2 (repeated).

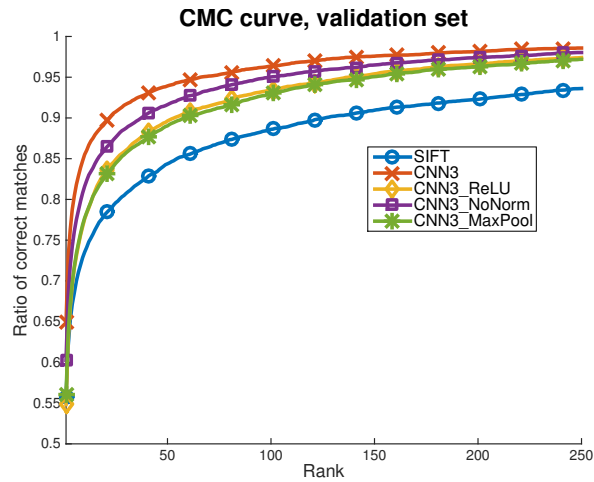Figure 13: ROC curves for the experiments of Sec. 4.2.

Figure 14: CMC curves for the experiments of Sec. 4.2.

6.3 EXPERIMENTS: (3) FRACKING

We extend the results of Sec. 4.3, which are summarized in Table 8. Figs. 15, 16 and 17 show the PR, ROC and CMC curves respectively.

$\frac{B_P}{B_P^M}$	$\frac{B_N}{B_N^M}$	PR AUC	ROC AUC	CMC AUC
1	1	0.366	0.977	0.979
1	2	0.558	0.986	0.987
2	2	0.596	0.988	0.989
4	4	0.703	0.993	0.993
8	8	**0.746**	**0.994**	**0.994**
16	16	0.538	0.983	0.986

Table 8: Extended table for the experiments of Sec. 4.3.

PR curve, validation set
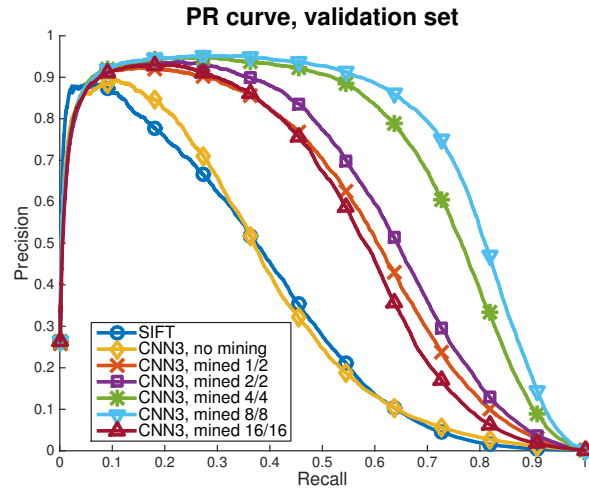
Figure 15: PR curves for the experiments of Sec. 4.3 (repeated).

ROC curve, validation set

Figure 16: ROC curves for the experiments of Sec. 4.3.

CMC curve, validation set

Figure 17: CMC curves for the experiments of Sec. 4.3.

6.4 EXPERIMENTS: (4) NUMBER OF FILTERS AND DESCRIPTOR DIMENSION

We extend the results of Sec. 4.4, which are summarized in Table 9. Figs. 18, 19 and 20 show the PR, ROC and CMC curves respectively.

Architecture	Output	Parameters	PR AUC	ROC AUC	CMC AUC
SIFT	128D	—	0.361	0.944	0.953
CNN3	128D	46,272	**0.596**	**0.988**	**0.989**
CNN3 Wide	128D	110,496	0.552	0.987	0.988
CNN3_NN1	128D	62,784	0.456	0.988	0.988
CNN3_NN1	32D	50,400	0.389	0.986	0.987

Table 9: Extended table for the experiments of Sec. 4.4.

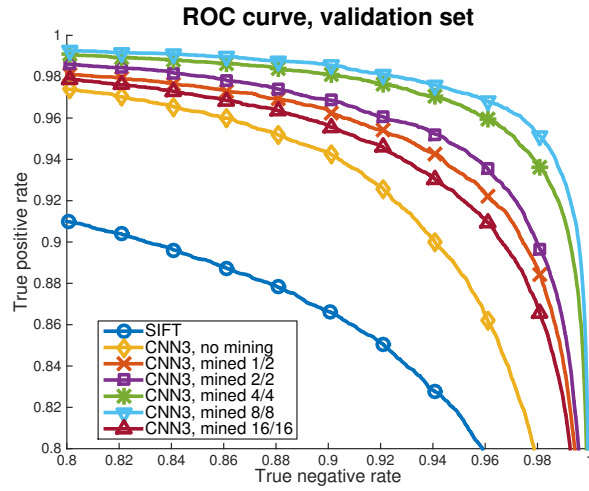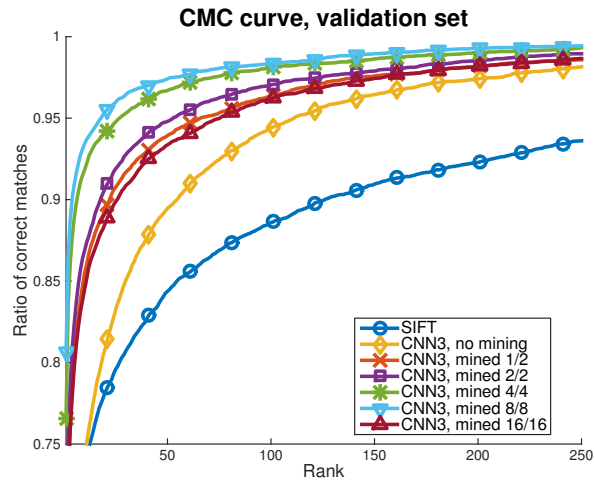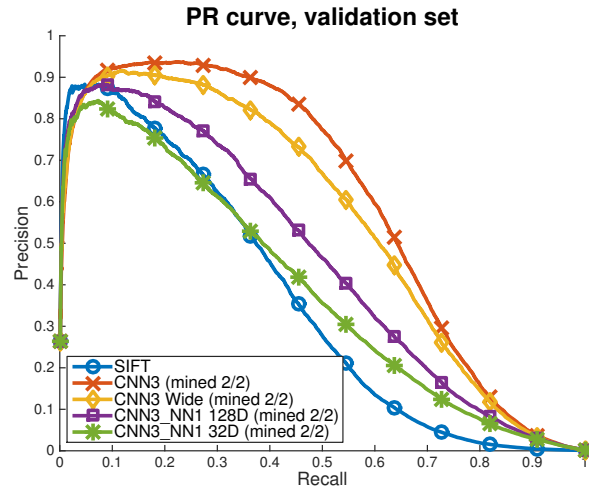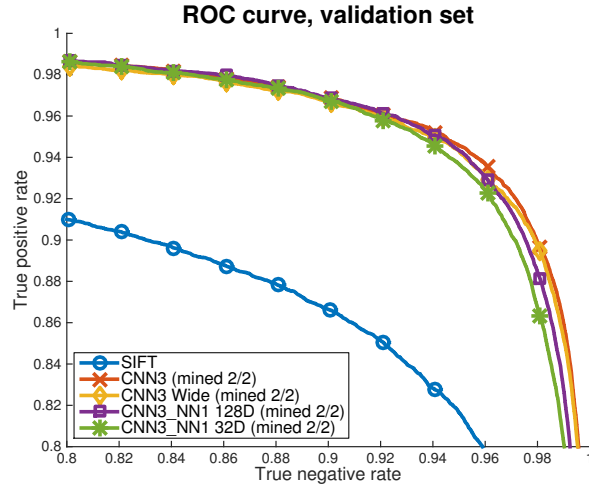Figure 18: PR curves for the experiments of Sec. 4.4 (repeated).

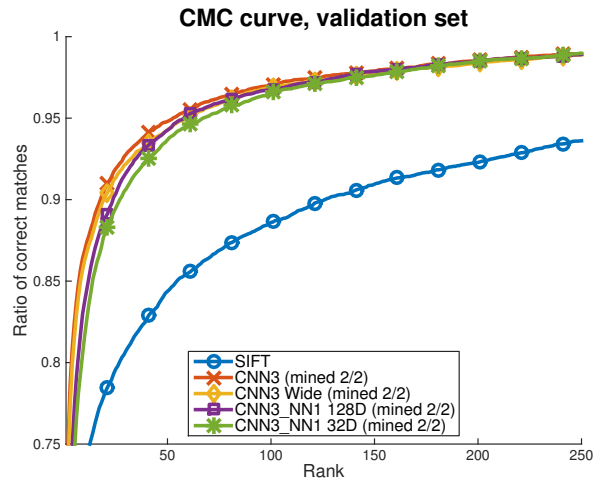Figure 19: ROC curves for the experiments of Sec. 4.4.

Figure 20: CMC curves for the experiments of Sec. 4.4.

6.5 EXPERIMENTS: (5) GENERALIZATION

In this section we extend the results of Sec. 4.5. We summarize the results over three different dataset splits, each with ten test folds of 10,000 randomly sampled positives and 1,000 randomly sampled negatives. We show the PR results in Tables 10-12, and Figs. 21-23, the ROC results in Tables 13-15, and Figs. 24-26, and the CMC results in Tables 16-18, and Figs. 27-29.

explain somewhere that we opt not to compute the auc over only part of the curve even though we could?

Precision-Recall AUC, Train: LY+YOS, Test: ND (10 folds)												
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.	Increase
SIFT	.364	.352	.345	.343	.349	.350	.350	.351	.341	.348	.349	—
CNN3, mine-1/2	.535	.527	.538	.537	.548	.529	.537	.535	.530	.529	.535	53.3%
CNN3, mine-2/2	.559	.548	.560	.556	.566	.554	.557	.554	.550	.549	.555	59.0%
CNN3, mine-4/4	.628	.619	.635	.632	.639	.625	.636	.631	.624	.626	.630	80.5%
CNN3, mine 8/8	.667	.658	.669	.667	.678	.659	.672	.667	.662	.666	**.667**	**91.1%**

Table 10: Generalized results in Precision-Recall. Models trained over LY+YOS and tested on ND.

Precision-Recall AUC, Train: LY+ND, Test: YOS (10 folds)												
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.	Increase
SIFT	.428	.419	.413	.416	.414	.427	.429	.442	.432	.430	.425	—
CNN3, mine-1/2	.381	.385	.367	.386	.366	.390	.393	.401	.383	.376	.383	-9.9%
CNN3, mine-2/2	.388	.395	.377	.393	.376	.397	.401	.405	.388	.381	.390	-8.2%
CNN3, mine-4/4	.502	.504	.483	.509	.485	.515	.513	.516	.499	.489	.502	18.1%
CNN3, mine-8/8	.547	.547	.528	.551	.528	.559	.556	.561	.546	.530	**.545**	**28.2%**

Table 11: Generalized results in Precision-Recall. Models trained over LY+ND and tested on YOS.

Precision-Recall AUC, Train: YOS+ND, Test: LY (10 folds)												
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.	Increase
SIFT	.223	.226	.229	.228	.226	.222	.233	.235	.219	.223	.226	—
CNN3, mine-1/2	.460	.464	.464	.460	.454	.452	.462	.463	.462	.456	.460	103.5%
CNN3, mine-2/2	.482	.487	.490	.485	.478	.472	.484	.488	.486	.478	.483	113.7%
CNN3, mine-4/4	.564	.566	.569	.562	.560	.557	.564	.567	.570	.562	.564	149.6%
CNN3, mine-8/8	.607	.611	.610	.604	.603	.604	.606	.615	.612	.608	**.608**	**169.0%**

Table 12: Generalized results in Precision-Recall. Models trained over LY+ND and tested on YOS.

ROC AUC, Train: LY+YOS, Test: ND (10 folds)											
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.956	.954	.955	.958	.957	.955	.955	.955	.956	.955	.956
CNN3, mine-1/2	.979	.978	.979	.982	.980	.978	.981	.981	.979	.979	.980
CNN3, mine-2/2	.981	.980	.981	.983	.982	.980	.983	.982	.981	.981	.981
CNN3, mine-4/4	.985	.984	.985	.987	.986	.985	.988	.986	.985	.985	.986
CNN3, mine-8/8	.986	.985	.986	.988	.987	.986	.989	.986	.986	.986	**.987**

Table 13: Generalized results in ROC. Models trained over LY+YOS and tested on ND.

ROC AUC, Train: LY+ND, Test: YOS (10 folds)										
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.949	.947	.948	.949	.949	.950	.949	.950	.950	.950	.949
CNN3, mine-1/2	.956	.953	.955	.956	.957	.958	.957	.957	.958	.957	.956
CNN3, mine-2/2	.958	.955	.957	.958	.959	.959	.958	.959	.960	.958	.958
CNN3, mine-4/4	.971	.969	.971	.971	.973	.973	.972	.972	.973	.971	.972
CNN3, mine-8/8	.974	.972	.975	.974	.976	.975	.975	.975	.976	.974	**.975**

Table 14: Generalized results in ROC. Models trained over LY+ND and tested on YOS.

ROC AUC, Train: YOS+ND, Test: LY (10 folds)										
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.938	.939	.936	.938	.933	.935	.936	.938	.937	.936	.937
CNN3, mine-1/2	.973	.973	.973	.971	.972	.972	.972	.973	.973	.972	.972
CNN3, mine-2/2	.975	.976	.975	.974	.976	.975	.974	.976	.976	.974	.975
CNN3, mine-4/4	.980	.980	.980	.979	.980	.980	.979	.982	.981	.979	.980
CNN3, mine-8/8	.983	.983	.983	.981	.983	.982	.982	.984	.983	.982	**.982**

Table 15: Generalized results in ROC. Models trained over LY+ND and tested on YOS.

CMC AUC, Train: LY+YOS, Test: ND (10 folds)										
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.964	.962	.963	.966	.965	.963	.964	.963	.964	.962	.963
CNN3, mine-1/2	.982	.981	.981	.984	.982	.981	.983	.983	.982	.981	.982
CNN3, mine-2/2	.983	.982	.983	.985	.984	.982	.985	.984	.984	.983	.984
CNN3, mine-4/4	.987	.987	.987	.989	.988	.987	.989	.988	.987	.987	.988
CNN3, mine-8/8	.988	.988	.988	.990	.989	.988	.990	.989	.989	.989	**.989**

Table 16: Generalized results in CMC. Models trained over LY+YOS and tested on ND.

CMC AUC, Train: LY+ND, Test: YOS (10 folds)										
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.956	.955	.956	.956	.956	.958	.956	.957	.956	.958	.956
CNN3, mine-1/2	.964	.962	.963	.964	.965	.966	.964	.965	.966	.966	.965
CNN3, mine-2/2	.966	.965	.966	.967	.968	.968	.967	.968	.969	.968	.967
CNN3, mine-4/4	.977	.976	.978	.978	.980	.980	.977	.979	.980	.980	.978
CNN3, mine-8/8	.980	.979	.981	.981	.982	.982	.980	.982	.982	.982	**.981**

Table 17: Generalized results in CMC. Models trained over LY+ND and tested on YOS.

CMC AUC, Train: YOS+ND, Test: LY (10 folds)										
Model	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10	Avg.
SIFT	.948	.949	.947	.948	.945	.945	.948	.949	.948	.947	.948
CNN3, mine-1/2	.976	.975	.976	.974	.975	.975	.975	.977	.976	.975	.975
CNN3, mine-2/2	.978	.979	.978	.977	.978	.978	.978	.979	.979	.977	.978
CNN3, mine-4/4	.983	.983	.983	.982	.983	.982	.982	.984	.984	.982	.983
CNN3, mine-8/8	.985	.985	.985	.984	.985	.984	.985	.986	.986	.985	**.985**

Table 18: Generalized results in CMC. Models trained over LY+ND and tested on YOS.
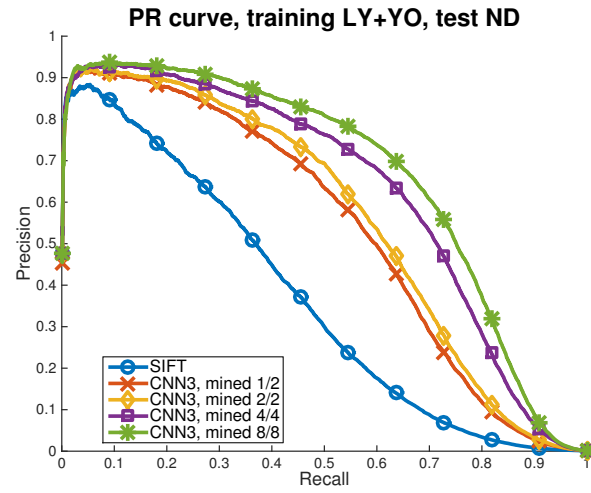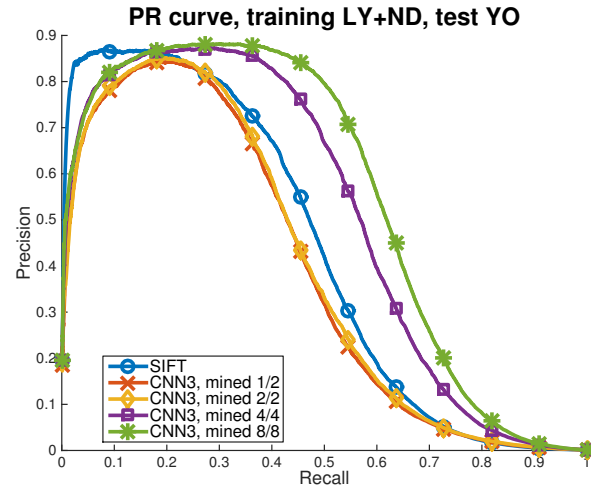
Figure 21: Generalized results in PR, first split.

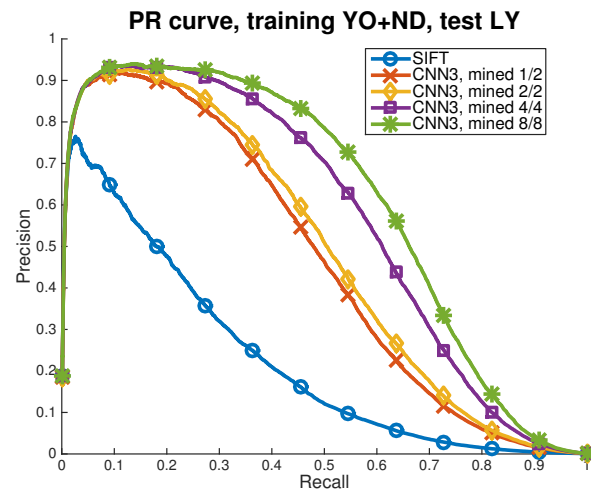Figure 22: Generalized results in PR, second split.

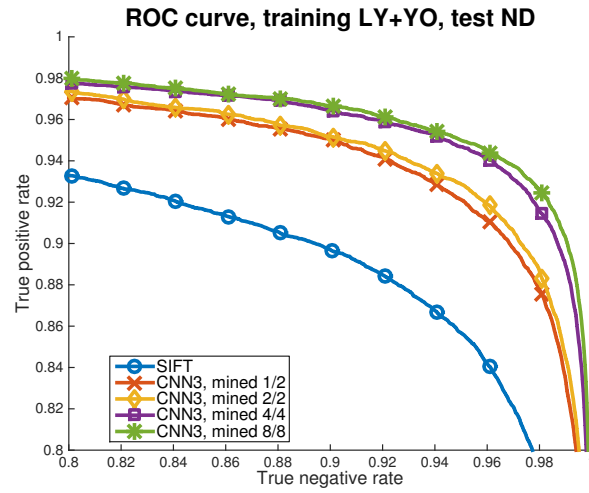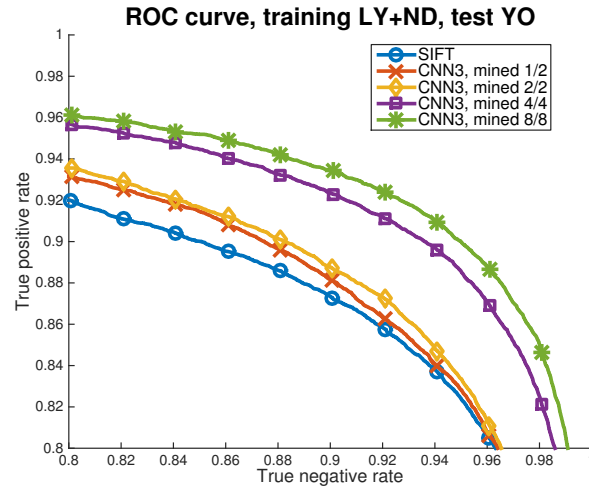Figure 23: Generalized results in PR, third split.
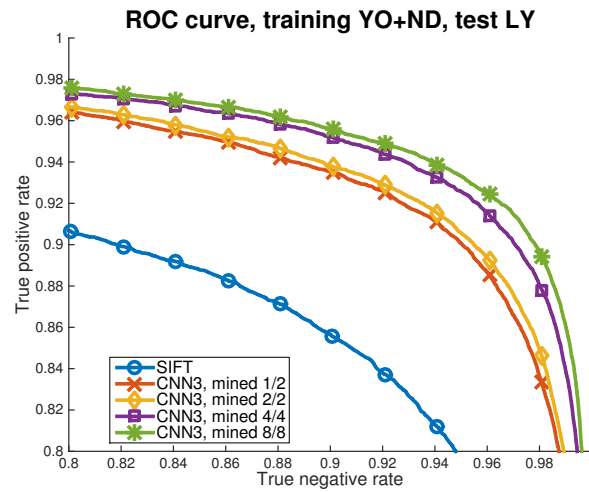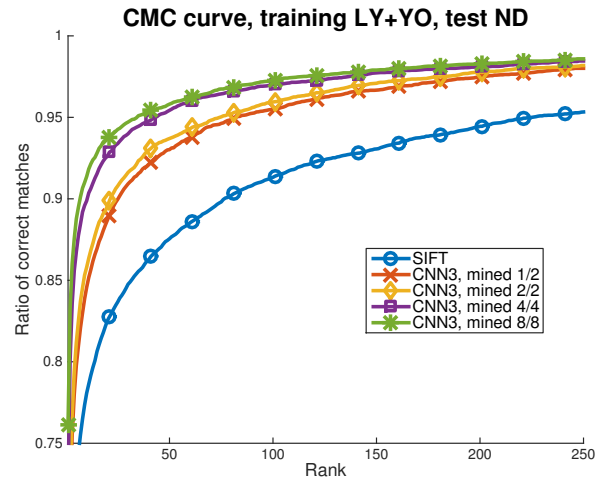
ROC curve, training LY+YO, test ND

Figure 24: Generalized results in ROC, first split.

ROC curve, training LY+ND, test YO

Figure 25: Generalized results in ROC, second split.

ROC curve, training YO+ND, test LY

Figure 26: Generalized results in ROC, third split.

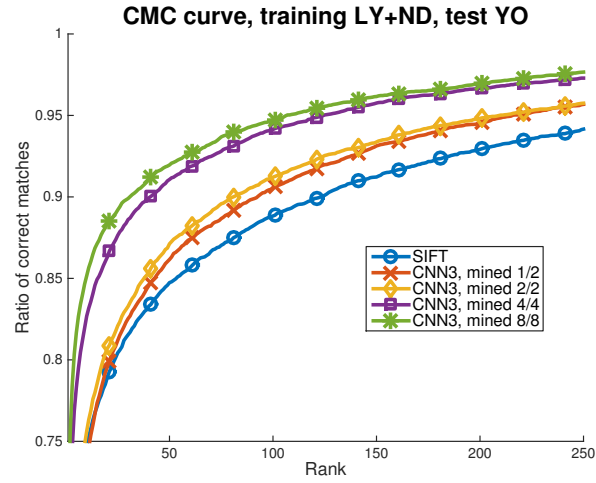Figure 27: Generalized results in CMC, first split.
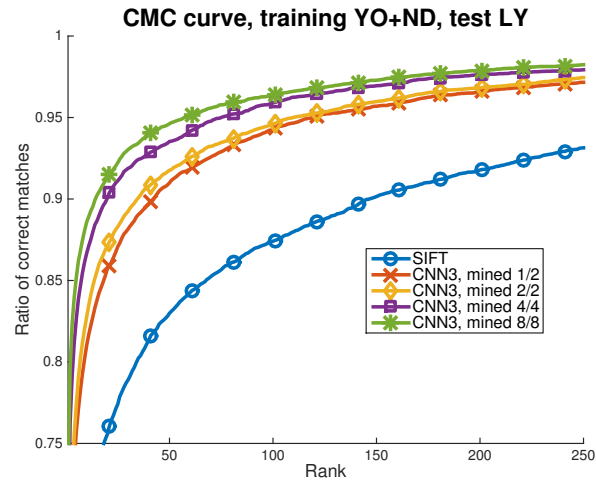

Figure 28: Generalized results in CMC, second split.

Figure 29: Generalized results in CMC, third split.