This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

ELF: Embedded Localisation of Features in Pre-Trained CNN

Assia Benbihi UMI2958 GeorgiaTech-CNRS Centrale Supélec Université Paris-Saclay Metz, France Matthieu Geist Google Research Brain Team Cédric Pradalier GeorgiaTech Lorraine-UMI2958 GeorgiaTech-CNRS Metz, France

abenbihi@georgiatech-metz.fr

Abstract

This paper introduces a novel feature detector based only on information embedded inside a CNN trained on standard tasks (e.g. classification). While previous works already show that the features of a trained CNN are suitable descriptors, we show here how to extract the feature locations from the network to build a detector. This information is computed from the gradient of the feature map with respect to the input image. This provides a saliency map with local maxima on relevant keypoint locations. Contrary to recent CNN-based detectors, this method requires neither supervised training nor finetuning. We evaluate how repeatable and how 'matchable' the detected keypoints are with the repeatability and matching scores. Matchability is measured with a simple descriptor introduced for the sake of the evaluation. This novel detector reaches similar performances on the standard evaluation HPatches dataset, as well as comparable robustness against illumination and viewpoint changes on Webcam and photo-tourism images. These results show that a CNN trained on a standard task embeds feature location information that is as relevant as when the CNN is specifically trained for feature detection.

1. Introduction

Feature extraction, description and matching is a recurrent problem in vision tasks such as Structure from Motion (SfM), visual SLAM and scene recognition. The extraction consists in detecting image keypoints, then the matching pairs the nearest keypoints based on their descriptor distance. Even though hand-crafted solutions, such as SIFT [19], prove to be successful, recent breakthroughs on local feature detection and description rely on supervised deeplearning methods [12, 25, 40]. They detect keypoints on saliency maps learned by a Convolutional Neural Network (CNN), then compute descriptors using another CNN or a



Figure 1. (1-6) Embedded Detector: Given a CNN trained on a standard vision task (classification), we backpropagate the feature map back to the image space to compute a saliency map. It is thresholded to keep only the most informative signal and keypoints are the local maxima. (7-8): Proxy-descriptor.

separate branch of it. They all require strong supervision and complex training procedures: [40] requires groundtruth matching keypoints to initiate the training, [25] needs the ground-truth camera pose and depth maps of the images, [12] circumvents the need for ground-truth data by using synthetic one but requires a heavy domain adaptation to transfer the training to realistic images. All these methods require a significant learning effort. In this paper, we show that a trained network already embeds enough information to build State-of-the-Art (SoA) detector and descriptor.

The proposed method for local feature detection needs only a CNN trained on standard task, such as ImageNet [11] classification, and no further training. The detector, dubbed ELF, relies on the features learned by such a CNN and extract their locations from the feature map gradients. Previous work already highlights that trained CNN features are relevant descriptors [13] and recent works [6, 15, 31] specifically train CNN to produce features suitable for keypoint description. However, no existing approach uses a pre-trained CNN for feature detection.

ELF computes the gradient of a trained CNN feature map with respect to *w.r.t* the image: this outputs a saliency map with local maxima on keypoint positions. Trained detectors learn this saliency map with a CNN whereas we extract it with gradient computations. This approach is inspired by [32] which observes that the gradient of classification scores *w.r.t* the image is similar to the image saliency map. ELF differs in that it takes the gradient of feature maps and not the classification score contrary to existing work exploiting CNN gradients [30, 34, 35, 37]. These previous works aim at visualising the learning signal for classification specifically whereas ELF extracts the feature locations. The extracted saliency map is then thresholded to keep only the most relevant locations and standard Non-Maxima Suppression (NMS) extracts the final keypoints (Figure 2).



Figure 2. Saliency maps thresholding to keep only the most informative location. Top: original image. (Left-Right: Webcam [39], HPatches [5], COCO[18]) Middle: blurred saliency maps. Bottom: saliency map after threshold. (Better seen on a computer.)

ELF relies only on six parameters: 2×2 Gaussian blur parameters for the automatic threshold level estimation and for the saliency map denoising; and two parameters for the (NMS) window and the border to ignore. Detection only requires one forward and one backward passes and takes ~0.2s per image on a simple Quadro M2200, which makes it suitable for real-time applications.

ELF is compared to individual detectors with standard *repeatability* [24] but results show that this metric is not discriminative enough. Most of the existing detectors can extract keypoints repeated across images with similar repeatability scores. Also, this metric does not express how 'useful' the detected keypoints are: if we sample all pixels as keypoints, we reach 100% of *rep*. but the matching may not be perfect if many areas look alike. Therefore, the detected keypoints are also evaluated on how 'matchable' they are with the *matching score* [24]. This metric requires

to describe the keypoints so we define a simple descriptor: it is based on the interpolation of a CNN feature map on the detected keypoints, as in [12]. This avoids biasing the performance by choosing an existing competitive descriptor. Experiments show that even this simple descriptor reaches competitive results which comforts the observation of [13], on the relevance of CNN features as descriptors. (See 4.1.)

ELF is tested on five architectures: three classification networks trained on ImageNet classification: AlexNet, VGG and Xception [17, 33, 9], as well as SuperPoint [12] and LF-Net [25] descriptor networks. Although outside the scope of this paper, this comparison provides preliminary results of the influence of the network architecture, task and training data on ELF's performance. Metrics are computed on HPatches [5] for general performances, scale and rotation robustness analysis; on Strecha and Webcam [36, 39] for light and 3D robustness analysis. We also test ELF on 3D reconstruction from images (CVPR19 Image Matching challenge [1]). Our contributions are the following:

- Extensive experiments show that a CNN trained on a standard vision task embeds feature location in its feature gradients. This data is as relevant for feature detection as when the CNN is specifically trained for it. This novel detector is on par with other sparse ones.
- We define a systematic method for local feature detection. They also update the previous result from [13]: self-taught CNN features provide SoA descriptors in spite of recent improvements in CNN descriptors [10].
- We release the python-based evaluation code to ease future comparison together with ELF code¹.

2. Related work

Early methods rely on hand-crafted detection and description : SIFT [19] detects 3D spatial-scale keypoints on difference of gaussians and describes them with a 3D Histogram Of Gradients (HOG). SURF [7] uses image integral to speed up the previous detection and uses a sum of Haar wavelet responses for description. KAZE [4] extends the previous multi-scale approach by detecting features in non-linear scale spaces instead of the classic Gaussian ones. ORB [27] combines the FAST [26] detection, the BRIEF [8] description, and improves them to make the pipeline scale and rotation invariant. MSER-based detector hand-crafts desired invariance properties for keypoints, and designs a fast algorithm to detect them [21]. Even though these handcrafted methods have proven to be successful and to reach state-of-the-art performance for some applications, recent research focus on learning-based methods.

One of the first learned detector is TILDE [39], trained under drastic changes of light and weather on the Webcam

¹ELF code:https://github.com/abenbihi/elf

dataset. They learn saliency maps which maxima are keypoint locations. Ground-truth saliency maps are generated with 'good keypoints' i.e. SIFT keypoints repeated in more than 100 images. One drawback of this method is the need for supervision that relies on another detector. However, there is no universal explicit definition of what a good keypoint is. This lack of specification inspires Quad-Networks [28] to adopt an unsupervised approach: they train a neural network to rank keypoints according to their robustness to random hand-crafted transformations. They keep the top/bottom quantile of the ranking as keypoints. ELF is similar in that it does not requires supervision but differs in that it does not need to further train the CNN.

Other learned detectors are trained within full detection/description pipelines such as LIFT [40], SuperPoint [12] and LF-Net [25]. LIFT contribution lies in their original training method of three CNNs. The detector CNN learns a saliency map which maxima are keypoints. They then crop patches around them, compute their orientations and descriptors with two other CNNs. They first train the descriptor with patches around ground-truth matching points with contrastive loss, then the orientation CNN together with the descriptor and finally with the detector. One drawback of this method is the need for ground-truth matching keypoints to initiate the training. In [12], the problem is avoided by pre-training the detector on a synthetic geometric dataset made of polygons on which they detect mostly corners. The detector is then finetuned during the descriptor training on image pairs from COCO [18] with synthetic homographies and the correspondence contrastive loss introduced in [10]. LF-Net relies on another type of supervision: it uses ground-truth camera poses and image depth maps that are 'easy' to compute with laser or standard SfM. Its training pipeline builds over LIFT and employs the projective camera model to project detected keypoints from one image to the other. These keypoint pairs form the groundtruth matching points to train the network. ELF differs in that the CNN model is already trained on a standard task. It then extracts the relevant information embedded inside the network for local feature detection, which requires no training nor supervision.

The detection method of this paper is mainly inspired from the initial observation in [32]: given a CNN trained for classification, the gradient of a class score *w.r.t* the image is the saliency map of the class object in the input image. A line of works aims at visualizing the CNN representation by inverting it into the image space through optimization [20, 14]. Following works use these saliency maps to better understand the CNN training process and justify the CNN outputs. Efforts mostly focus on the gradient definitions [30, 34, 35, 37, 42]. ELF simply backpropagates the feature map back to the image space. As far as we know, this is the first work to exploit such gradients for feature detection.

The simple descriptor used in the matchability evaluation is borrowed from UCN [10]. Given a feature map and the keypoints to describe, it interpolates the feature map on the keypoints location. Trained CNNs have long been known to provide relevant feature descriptors [13]. Recent research has taken on specifically training the CNN for description with patch-based approaches, (*e.g.* [31, 22, 15, 41]) or image-based approaches [10, 38]. We favor [10], also used by SuperPoint, for its simplicity over [38] and its O(1) complexity compared to O(N) patch-based approaches with N the number of keypoints.

3. Method

This section defines ELF, a detection method valid for any trained CNN. Keypoints are local maxima of a saliency map computed as the feature gradient *w.r.t* the image. We use the data adaptive Kapur method [16] to automatically threshold the saliency map and keep only the most salient locations, then run NMS for local maxima detection.



Figure 3. (Bigger version in Appendix.) Saliency maps computed from the feature map gradient $\left| {}^{T}F^{l}(x) \cdot \frac{\partial F^{l}}{\partial \mathbf{I}} \right|$. Enhanced image contrast for better visualisation. Top row: gradients of VGG *pool*₂ and *pool*₃ show a loss of resolution from *pool*₂ to *pool*₃. Bottom: $(pool_{i})_{i \in [1,2,5]}$ of VGG on Webcam, HPatches and Coco images. Low level saliency maps activate accurately whereas higher saliency maps are blurred.

3.1. Feature Specific Saliency

We generate a saliency map that activates on the most informative image region for a specific CNN feature level *l*. Let **I** be a vector image of dimension $D_I = H_I \cdot W_I \cdot C_I$. Let F^l be a vectorized feature map of dimension $D_F =$ $H_l \cdot W_l \cdot C_l$. The saliency map S^l , of dimension D_I , is $S^l(\mathbf{I}) = |{}^t F^l(\mathbf{I}) \cdot \nabla_I F^l|$, with $\nabla_I F^l$ a $D_F \times D_I$ matrix.

The saliency activates on the image regions that contribute the most to the feature representation $F^{l}(\mathbf{I})$. The term $\nabla_I F^l$ explicits the correlation between the feature space of F^l and the image space in general. The multiplication by $F^l(\mathbf{I})$ applies the correlation to the features $F^l(\mathbf{I})$ specifically and generate a visualisation in image space $S^l(\mathbf{I})$. From a geometrical point of view, this operation can be seen as the projection $\nabla_I F^l$ of a feature signal $F^l(\mathbf{I})$ into the image space. From a signal processing approach, $F^l(\mathbf{I})$ is an input signal filtered through $\nabla_I F^l$ into the image space. If $C_I > 1$, S^l is converted into a grayscale image by averaging it across channels.

3.2. Feature Map Selection

We provide visual guidelines to choose the feature level l so that F^l still holds high resolution localisation information while providing a useful high-level representation.

CNN operations such as convolution and pooling increase the receptive field of feature maps while reducing their spatial dimensions. This means that F^l has less spatial resolution than F^{l-1} and the backpropagated signal $S^{\bar{l}}$ ends up more spread than S^{l-1} . This is similar to when an image is too enlarged and it can be observed in Figure 3, which shows the gradients of the VGG feature maps. On the top row, pool₂'s gradient (left) better captures the location details of the dome whereas pool₃'s gradient (right) is more spread. On the bottom rows, the images lose their resolution as we go higher in the network. Another consequence of this resolution loss is that small features are not embedded in F^l if l is too high. This would reduce the space of potential keypoint to only large features which would hinder the method. This observation motivates us to favor lowlevel feature maps for feature detection. We chose the final F^{l} by taking the highest l which provides accurate localisation. This is visually observable by sparse high intensity signal contrary to the blurry aspect of higher layers.

3.3. Automatic Data-Adaptive Thresholding

The threshold is automatic and adapts to the saliency map distribution to keep only the most informative regions. Figure 2 shows saliency maps before and after thresholding using Kapur's method [16], which we briefly recall below. It chooses the threshold to maximize the information between the image background and foreground *i.e.* the pixel distribution below and above the threshold. This method is especially relevant in this case as it aims at maintaining as much information on the distribution above the threshold as possible. This distribution describes the set of local maxima among which we choose our keypoints. More formally, for an image I of N pixels with n sorted gray levels and $(f_i)_{i \in n}$ the corresponding histogram, $p_i = \frac{f_i}{N}$ is the empirical probability of a pixel to hold the value f_i . Let $s \in n$ be a threshold level and A, B the empirical background and foreground distributions. The level s is chosen to maximize the information between A and B and the threshold value is

set to f_s : $A = \left(\frac{p_i}{\sum_{i < s} p_i}\right)_{i < s}$ and $B = \left(\frac{p_i}{\sum_{i > = s} p_i}\right)_{i > s}$. For better results, we blur the image with a Gaussian of parameters $(\mu_{thr}, \sigma_{thr})$ before computing the threshold level.

Once the threshold is set, we denoise the image with a second Gaussian blur of parameters $(\mu_{noise}, \sigma_{noise})$ and run standard NMS (the same as for SuperPoint) where we iteratively select decreasing global maxima while ensuring that their nearest neighbor distance is higher than the window $w_{\text{NMS}} \in \mathbb{N}$. Also we ignore the $b_{\text{NMS}} \in \mathbb{N}$ pixels around the image border.

3.4. Simple descriptor

As mentioned in the introduction, the repeatability score does not discriminate detectors anymore. So they are also evaluated on how 'matchable' their detected keypoints are with the matching score. The ELF detector is completed with a simple descriptor inspired by UCN [10]: a CNN feature map is interpolated on the detected keypoints. Its simplicity over existing competitive ones avoids unfairly boosting ELF's perfomance. Although simple, experiments show that this description completes ELF into a competitive feature detection/description method.

The feature map used for description may be different from the one for detection. High-level feature maps have wider receptive field hence take higher context into account for the description of a pixel location. This leads to more informative descriptors which motivates us to favor higher level maps. However we are also constrained by the loss of resolution previously described: if the feature map level is too high, the interpolation of the descriptors generate vector too similar to each other. For example, the VGG $pool_A$ layer produces more discriminative descriptors than pool₅ even though $pool_5$ embeds information which higher level semantics. Empirically, we observe there exists a layer level l' above which the description performance stops increasing before decreasing. This is measured through the matching score [24]. The final choice of the feature map is done by testing some layers l' > l and select the lowest feature map before the descriptor performance stagnates.

Detectors are evaluated with both their original descriptor and this simple one. The motivation is that detectors may be biased to sample keypoints that their respective descriptor can describe 'well' [40]. So it is fair to compute the matching score with the original detector/descriptor pairs. However, a detector can sample 'useless points' (e.g. sky pixels for 3d reconstructions) that its descriptor can characterise 'well'. In this case, the descriptor 'hides' the detector default. This motivates the integration of a common independent descriptor with all detectors to evaluate them.

4. Experiments

This section describes the evaluation metrics and datasets as well as the method's tuning. Our method is compared to detectors with available public code: the fully hand-crafted SIFT [19], SURF [7], ORB [27], KAZE [4], the learning-based LIFT [40], SuperPoint [12], LF-Net [25], the individual detectors TILDE [39], MSER [21].

4.1. Metrics

We follow the standard validation [24] that evaluates the detection performance with *repeatability* (*rep*). It measures the percentage of keypoints common to both images. We also compute the *matching score* (*ms*) as an additional *detector* metric. It captures the percentage of keypoint pairs that are nearest neighbours in both image space and descriptor space i.e. the ratio of keypoints correctly matched. For completeness, the metric definitions are recalled in Appendix and their implementation in the released code.

A way to reach perfect *rep* is to sample all the pixels or sample them with a frequency higher than the distance threshold ϵ_{kp} of the metric. One way to prevent the first flaw is to limit the number of keypoints but it does not counter the second. Since detectors are always used together with descriptors, another way to think the detector evaluation is: 'a good keypoint is one that can be discriminatively described and matched'. One could think that such a metric can be corrupted by the descriptor. But we ensure that a detector flaw cannot be hidden by a very performing descriptor with two guidelines. One experiment must evaluate all detector with one fixed descriptor (the simple one defined in 3.4). Second, ms can never be higher than rep so a detector with a poor rep leads to a poor ms.

Here the number of detected keypoints is limited to 500 for all methods. As done in [12, 25], we replace the overlap score in [24] to compute correspondences with the 5-pixel distance threshold. Following [40], we also modify the matching score definition of [24] to run a greedy bipartitegraph matching on all descriptors and not just the descriptor pairs for which the distance is below an arbitrary threshold. We do so to be able to compare all state-of-the-art methods even when their descriptor dimension and range vary significantly. (More details in Appendix.)

4.2. Datasets

All images are resized to the 480×640 pixels and the image pair transformations are rectified accordingly.

General performances. The HPatches dataset [5] gathers a subset of standard evaluation images such as DTU and OxfordAffine [2, 23]: it provides a total of 696 images, 6 images for 116 scenes and the corresponding homographies between the images of a same scene. For 57 of these scenes, the main changes are photogrammetric and the remaining



Figure 4. Left-Right: HPatches: planar viewpoint. Webcam: light. HPatches: rotation. HPatches: scale. Strecha: 3D viewpoint.

59 show significant geometric deformations due to viewpoint changes on planar scenes.

Illumination Robustness. The Webcam dataset [39] gathers static outdoor scenes with drastic natural light changes contrary to HPatches which mostly holds artificial light changes in indoor scenes.

Rotation and Scale Robustness. We derive two datasets from HPatches available with the code. For each of the 116 scenes, we keep the first image and rotate it with angles from 0° to 210° with an interval of 40° . Four zoomed-in version of the image are generated with scales [1.25, 1.5, 1.75, 2].

3D Viewpoint Robustness. We use three Strecha scenes [36] with increasing viewpoint changes: *Fountain, Castle entry, Herzjesu-P8*. The viewpoint changes proposed by HPatches are limited to planar scenes which does not reflect the complexity of 3D structures. We use COLMAP [29] to get 'ground-truth' scaleless depth that we release. ELF is also tested in the CVPR19 Image Matching Challenge [1].

4.3. Baselines

We describe the rationale behind the evaluation. The tests run on a QuadroM2200 with Tensorflow 1.4, Cuda8, Cudnn6 and Opencv3.4. We use the OpenCV implementation of SIFT, SURF, ORB, KAZE, MSER with the default parameters and the author's code for TILDE, LIFT, SuperPoint, LF-Net with the provided models and parameters. When comparing detectors in the feature matching pipeline, we measure their matching score with both their original descriptor and ELF simple descriptor. For MSER and TILDE, we use the VGG simple descriptor.

Architecture influence. ELF is tested on five networks: three classification ones trained on ImageNet (AlexNet, VGG, Xception [17, 33, 9]) as well as the trained Super-Point's and LF-Net's descriptor ones. We call each variant with the network's names prefixed with ELF as in saliency. The paper compares the influence of i) architecture for a fixed task (ELF-AlexNet [17] vs. ELF-VGG [33] v.s. ELF-Xception [9]), ii) the task (ELF-VGG vs. ELF-SuperPoint (SP) descriptor), iii) the training dataset (ELF-LFNet on phototourism vs. ELF-SP on MS-COCO). This study is being refined with more independent comparisons of tasks, datasets and architectures in a journal extension.

We use the author's code and pre-trained models which

we convert to Tensorflow [3]. We search the blurring parameters $(\mu_{thr}, \sigma_{thr}), (\mu_{noise}, \sigma_{noise})$ in the range $[3, 21]^2$; the NMS parameters (w_{NMS}, b_{NMS}) in $[4, 13]^2$.

Individual components comparison. Individual detectors are compared with the matchability of their detection together with the description of the simple VGG-pool3 descriptor. This way, the *m.s.* only depends on the detection performance since the description is fixed for all detectors. This experiment also raises the question of whether triplet-like losses are relevant to train CNN descriptors. Indeed, these losses constrain the CNN features directly so that matching keypoints are near each other in descriptor space. Simpler loss, such as cross-entropy for classification, only constrain the CNN output on the task while leaving the representation up to the CNN. The ELF-VGG detector is also integrated with existing descriptors. This evaluates how useful the CNN self-learned feature location compares with the hand-crafted and the learned ones.

Gradient Baseline. Visually, the feature gradient map is reminiscent of the image gradients computed with the Sobel or Laplacian operators. We run two variants of our pipeline where we replace the feature gradient with them. This aims at showing whether CNN feature gradients embed more information than image intensity gradients.

5. Results

Experiments show that ELF compares with the SoA on HPatches and demonstrates robustness similar to recent learned methods. It generates saliency maps visually akin to a Laplacian on very structured images (HPatches) but proves to be more robust on outdoor scenes with natural conditions (Webcam). When integrated with existing feature descriptors, ELF boosts the matching score. Even integrating ELF simple descriptor improves it with the exception of SuperPoint for which results are equivalent. This sheds new light on the representations learnt by CNNs and suggests that deep description methods may underexploit the information embedded in their trained networks.

Numerical results show that all methods can detect repeatable keypoints with similar performances. Even though the matchability of the points (m.s) is a bit more discriminative, neither express how 'useful' the kp are for the end-goal task. One way to do so would be to evaluate an end-goal task (*e.g.* Structure-from-Motion). However, for the evaluation to be rigorous all the other steps should be fixed for all papers. Recently, the Image Matching CVPR19 workshop proposed such an evaluation but is not fully automatic yet. These results also challenge whether current descriptor-training loss are a strong enough signal to constrain CNN features better than a simple cross-entropy.

Unless mentioned otherwise, we compute repeatability for each detector, and the matching score of detectors with their respective descriptors, when they have one. We use



Figure 5. Top-Down: HPatches-Webcam. Left-Right: repeatability, matching score. Better seen on digital screen with colors.

ELF-VGG- $pool_4$ descriptor for TILDE, MSER, ELF-VGG, ELF-SuperPoint, and ELF-LFNet. We use AlexNet and Xception feature maps to build their respective simple descriptors. The meta-parameters for each variants are provided in Appendix.

General performances. Figure 5 (top) shows that the rep variance is low across detectors whereas ms is more discriminative, hence the validation method (Section 4.1). On HPatches, SuperPoint (SP) reaches the best rep-ms [68.6, 57.1] closely followed by ELF (e.g. ELF-VGG: [63.8, 51.8]) and TILDE [66.0, 46.7]. In general, we observe that learning-based methods all outperform hand-crafted ones. Still, LF-Net and LIFT curiously underperform on HPatches: one reason may be that the data they are trained on differs too much from this one. LIFT is trained on outdoor images only and LF-Net on either indoor or outdoor datasets, whereas HPatches is made of a mix of them. We compute metrics for both LF-Net models and report the highest one (indoor). Even though LF-Net and LIFT fall behind the top learned methods, they still outperform handcrafted ones which suggests that their framework learn feature specific information that hand-crafted methods can not capture. This supports the recent direction towards trained detectors and descriptors.

Light Robustness Again, *ms* is a better discriminant on Webcam than *rep* (Figure 5 bottom). ELF-VGG reaches top *rep-ms* [53.2, 43.7] closely followed by TILDE [52.5, 34.7] which was the state-of-the-art detector.

Overall, there is a performance degradation ($\sim 20\%$) from HPatches to Webcam. HPatches holds images with standard features such as corners that state-of-the-art methods are made to recognise either by definition or by supervision. There are less such features in the Webcam dataset because of the natural lighting that blurs them. Also there are strong intensity variations that these models do not handle well. One reason may be that the learning-based methods never saw such lighting variations in their training set. But this assumption is rejected as we observe that even Super-Point, which is trained on Coco images, outperforms LIFT and LF-Net, which are trained on outdoor images. Another justification can be that what matters the most is the pixel distribution the network is trained on, rather than the image content. The top methods are classifier-based ELF and SuperPoint: the first ones are trained on the huge ImageNet dataset and benefit from heavy data augmentation. Super-Point also employs a considerable data strategy to train their network. Thus these networks may cover a much wider pixel distribution which would explain their robustness to pixel distribution changes such as light modifications.

Architecture influence ELF is tested on three classification networks as well as the descriptor networks of Super-Point and LF-Net (Figure 5, bars under 'ELF').

For a fixed training task (classification) on a fixed dataset (ImageNet), VGG, AlexNet and Xception are compared. As could be expected, the network architecture has a critical impact on the detection and ELF-VGG outperforms the other variants. The rep gap can be explained by the fact that AlexNet is made of wider convolutions than VGG, which induces a higher loss of resolution when computing the gradient. As for *ms*, the higher representation space of VGG may help building more informative features which are a stronger signal to backpropagate. This could also justify why ELF-VGG outperforms ELF-Xception that has less parameters. Another explanation is that ELF-Xception's gradient maps seem smoother. Salient locations are then less emphasized which makes the keypoint detection harder. One could hint at the depth-wise convolution to explain this visual aspect but we could not find an experimental way to verify it. Surprisingly, ELF-LFNet outperforms the original LF-Net on both HPatches and Webcam and ELF-SuperPoint variant reaches similar results as the original.



Figure 6. Top-Down: Scale, Rotation. Left-Right: rep, ms.

Scale Robustness. ELF-VGG is compared with SoA detectors and their respective descriptors (Figure 6, top). *Rep.* is mostly stable for all methods: SIFT and SuperPoint are the most invariant whereas ELF follows the same variations

as LIFT and LF-Net. Once again, *ms* better assesses the detectors performance: SuperPoint is the most robust to scale changes, followed by LIFT and SIFT. ELF and LF-Net lose 50% of their matching score with the increasing scale. It is surprising to observe that LIFT is more scale-robust than LF-Net when the latter's global performance is higher. A reasonable explanation is that LIFT detects keypoints at 21 scales of the same image whereas LF-Net only runs its detector CNN on 5 scales. Nonetheless, ELF outperforms LF-Net without manual multi-scale processing.

Rotation Robustness. Even though rep shows little variations (Figure 6, down), all learned methods' ms crash while only SIFT survives the rotation changes. This can be explained by the explicit rotation estimation step of SIFT. However LIFT and LF-Net also run such a computation. This suggests that either SIFT's hand-crafted orientation estimation is more accurate or that HOG are more rotation invariant than learned features. LF-Net still performs better than LIFT: this may be because it learns the keypoint orientation on the keypoint features representation rather than the keypoint pixels as done in LIFT. Not surprisingly, ELF simple descriptor is not rotation invariant as the convolutions that make the CNN are not. This also explains why SuperPoint also crashes in a similar manner. These results suggest that the orientation learning step in LIFT and LF-Net is needed but its robustness could be improved.



Figure 7. Robustness analysis: 3D viewpoint.

3D Viewpoint Robustness. While SIFT shows a clear advantage of pure-rotation robustness, it displays similar degradation as other methods on realistic rotation-andtranslation on 3D structures. Figure 7 shows that all methods degrade uniformly. One could assume that this small data sample is not representative enough to run such robustness analysis. However, we think that these results rather suggest that all methods have the same robustness to 3D viewpoint changes. Even though previous analyses allows to rank the different feature matching pipelines, each has advantages over others on certain situations: ELF or SuperPoint on general homography matches, or SIFT on rotation robustness. This is why this paper only aims at showing ELF reaches the same performances and shares similar properties to existing methods as there is no generic ranking criteria. The evaluation run after the submission by the CVPR19 Image Matching Challenge [1] supports the previous conclusions. See numerical results in Appendix.



Figure 8. Left-Middle-Right bars: original method, integration of ELF detection, integration of ELF description.

Individual components performance. First, all methods' descriptor are replaced with the simple ELF-VGG $pool_3$ one. We then compute their new *ms* and compare it to ELF-VGG on HPatches and Webcam (Figure 8, stripes). The description is based on $pool_3$ instead of $pool_4$ here for it produces better results for the other methods while preserving ours. ELF reaches higher ms [51.3] for all methods except for SuperPoint [53.7] for which it is comparable. This shows that ELF is as relevant, if not more, than previous hand-crafted or learned detectors. This naturally leads to the question: 'What kind of keypoints does ELF detect ?' There is currently no answer to this complex question. Still, we observe that ELF activates mostly on high intensity gradient areas although not all of them. One explanation is that as the CNN is trained on the vision task, it learns to ignore image regions useless for the task. This results in killing the gradient signals in areas that may be unsuited for matching.

Another surprising observation regards CNN descriptors: SuperPoint (SP) keypoints are described with the SP descriptor in one hand and the simple ELF-VGG one in the other hand. Comparing the two resulting matching scores is one way to compare the SP and ELF descriptors. Results show that both approaches lead to similar ms. This result is surprising because SP specifically trains a description CNN so that its feature map is suitable for keypoint description [10]. In VGG training, there is no explicit constraints on the features from the cross-entropy loss. Still, both feature maps reach similar numerical description performance. This raises the question of whether contrastive-like losses, which input are CNN features, can better constrain the CNN representation than simpler losses, such as cross-entropy, which inputs are classification logits. This also shows that there is more to CNNs than only the task they are trained on: they embed information that can prove useful for unrelated tasks. Although the simple descriptor was defined for evaluation purposes, these results demonstrate that it can be used as a description baseline for feature extraction.

The integration of ELF detection with other methods' descriptor (Figure 8, circle) boosts the *ms*. [40] previously suggested that there may be a correlation between the detector and the descriptor within a same method, i.e. the LIFT descriptor is trained to describe only the keypoints output by its detector. However, these results show that ELF can easily be integrated into existing pipelines and even boost

their performances.



Figure 9. Gradient baseline.

Gradient Baseline The ELF saliency map used is replaced with simple Sobel or Laplacian gradient maps. The rest of the detection stays the same. They are completed with simple ELF descriptors from the VGG, AlexNet and Xception networks and we compute their performance (Figure 9 Left). These new hybrids are then compared to their respective ELF variant (Right). Results show that these simpler gradients can detect systematic keypoints with comparable rep on very structured images such as HPatches. However, the ELF detector better overcomes light changes (Webcam). On HPatches, the Laplacian-variant reaches similar ms as ELF-VGG (55 vs 56) and outperforms ELF-AlexNet and ELF-Xception. These scores can be explained with the images structure: for heavy textured images, high intensity gradient locations are relevant enough keypoints. However, on Webcam, all ELF detectors outperform Laplacian and Sobel with a factor of 100%. This shows that ELF is more robust than Laplacian and Sobel operators.

Qualitative results Available in the video². Green lines show putative matches based only on nearest neighbour matching of descriptors.

6. Conclusion

We have introduced ELF, a novel method to extract feature locations from pre-trained CNNs, with no further training. Extensive experiments show that it performs as well as state-of-the art detectors. It can easily be integrated into existing matching pipelines and proves to boost their matching performances. Even when completed with a simple feature-map-based descriptor, it turns into a competitive feature matching pipeline. These results shed new light on the information embedded inside trained CNNs. This work also raises questions on the descriptor training of deep-learning approaches: whether their losses actually constrain the CNN to learn better features than the ones it would learn on its own to complete a vision task. Preliminary results show that the CNN architecture, the training task and the dataset have substantial impact on the detector performances. A further analysis of these correlations is the object of a future work.

²https://youtu.be/oxbG5162yDs

References

- Cvpr19 image matching challenge. https: //image-matching-workshop.github.io/ challenge/, 2019.
- [2] Henrik Aanæs, Anders Lindbjerg Dahl, and Kim Steenstrup Pedersen. Interesting interest points. *International Journal* of Computer Vision, 97(1):18–35, 2012.
- [3] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In OSDI, volume 16, pages 265–283, 2016.
- [4] Pablo Fernández Alcantarilla, Adrien Bartoli, and Andrew J Davison. Kaze features. In *European Conference on Computer Vision*, pages 214–227. Springer, 2012.
- [5] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Conference* on Computer Vision and Pattern Recognition (CVPR), volume 4, page 6, 2017.
- [6] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallowconvolutional neural networks. In *BMVC*, volume 1, page 3, 2016.
- [7] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.
- [8] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. Brief: Binary robust independent elementary features. In *European conference on computer vision*, pages 778–792. Springer, 2010.
- [9] François Chollet. Xception: Deep learning with depthwise separable convolutions. In 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, pages 1800–1807, 2017.
- [10] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In Advances in Neural Information Processing Systems, pages 2414–2422, 2016.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009. CVPR 2009. IEEE Conference on, pages 248–255. Ieee, 2009.
- [12] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In CVPR Deep Learning for Visual SLAM Workshop, 2018.
- [13] Philipp Fischer, Alexey Dosovitskiy, and Thomas Brox. Descriptor matching with convolutional neural networks: a comparison to sift. arXiv preprint arXiv:1405.5769, 2014.
- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, pages 2414–2423, 2016.
- [15] Xufeng Han, Thomas Leung, Yangqing Jia, Rahul Sukthankar, and Alexander C Berg. Matchnet: Unifying feature

and metric learning for patch-based matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3279–3286, 2015.

- [16] Jagat Narain Kapur, Prasanna K Sahoo, and Andrew KC Wong. A new method for gray-level picture thresholding using the entropy of the histogram. *Computer vision, graphics, and image processing*, 29(3):273–285, 1985.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [19] David G Lowe. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [20] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5188–5196, 2015.
- [21] Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*, 22(10):761– 767, 2004.
- [22] Iaroslav Melekhov, Juho Kannala, and Esa Rahtu. Siamese network features for image matching. In 2016 23rd International Conference on Pattern Recognition (ICPR), pages 378–383. IEEE, 2016.
- [23] Krystian Mikolajczyk and Cordelia Schmid. A performance evaluation of local descriptors. *IEEE transactions on pattern* analysis and machine intelligence, 27(10):1615–1630, 2005.
- [24] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, Jiri Matas, Frederik Schaffalitzky, Timor Kadir, and Luc Van Gool. A comparison of affine region detectors. *International journal of computer vision*, 65(1-2):43–72, 2005.
- [25] Yuki Ono, Eduard Trulls, Pascal Fua, and Kwang Moo Yi. Lf-net: learning local features from images. In Advances in Neural Information Processing Systems, pages 6234–6244, 2018.
- [26] Edward Rosten and Tom Drummond. Machine learning for high-speed corner detection. In *European conference on computer vision*, pages 430–443. Springer, 2006.
- [27] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *Computer Vision (ICCV)*, 2011 IEEE international conference on, pages 2564–2571. IEEE, 2011.
- [28] Nikolay Savinov, Akihito Seki, Lubor Ladicky, Torsten Sattler, and Marc Pollefeys. Quad-networks: unsupervised learning to rank for interest point detection. In Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- [29] Johannes L Schonberger and Jan-Michael Frahm. Structurefrom-motion revisited. In *Proceedings of the IEEE Con-*

ference on Computer Vision and Pattern Recognition, pages 4104–4113, 2016.

- [30] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra, et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, pages 618–626, 2017.
- [31] Edgar Simo-Serra, Eduard Trulls, Luis Ferraz, Iasonas Kokkinos, Pascal Fua, and Francesc Moreno-Noguer. Discriminative learning of deep convolutional feature point descriptors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 118–126, 2015.
- [32] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. arXiv preprint arXiv:1312.6034, 2013.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [34] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [35] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. In *ICLR (workshop track)*, 2015.
- [36] Christoph Strecha, Wolfgang Von Hansen, Luc Van Gool, Pascal Fua, and Ulrich Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *Computer Vision and Pattern Recognition*, 2008. *CVPR 2008. IEEE Conference on*, pages 1–8. Ieee, 2008.
- [37] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International Conference* on Machine Learning, pages 3319–3328, 2017.
- [38] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. Inloc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7199–7209, 2018.
- [39] Yannick Verdie, Kwang Yi, Pascal Fua, and Vincent Lepetit. Tilde: A temporally invariant learned detector. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 5279–5288, 2015.
- [40] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. Lift: Learned invariant feature transform. In *European Conference on Computer Vision*, pages 467–483. Springer, 2016.
- [41] Sergey Zagoruyko and Nikos Komodakis. Learning to compare image patches via convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4353–4361, 2015.
- [42] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.