SImProv: Scalable Image Provenance Framework for Robust Content Attribution

Alexander Black, Tu Bui, Simon Jenni, Zhifei Zhang, Viswanathan Swaminanthan, John Collomosse

Abstract—We present SImProv – a scalable image provenance framework to match a query image back to a trusted database of originals and identify possible manipulations on the query. SImProv consists of three stages: a scalable search stage for retrieving top-k most similar images; a re-ranking and near-duplicated detection stage for identifying the original among the candidates; and finally a manipulation detection and visualization stage for localizing regions within the query that may have been manipulated to differ from the original. SImProv is robust to benign image transformations that commonly occur during online redistribution, such as artifacts due to noise and recompression degradation, as well as out-of-place transformations due to image padding, warping, and changes in size and shape. Robustness towards out-of-place transformations is achieved via the end-to-end training of a differentiable warping module within the comparator architecture. We demonstrate effective retrieval and manipulation detection over a dataset of 100 million images.

Index Terms—Content Attribution, Image Provenance, Manipulation Detection, Near-duplicated Search, Large-scale Retrieval, Comparator Network.
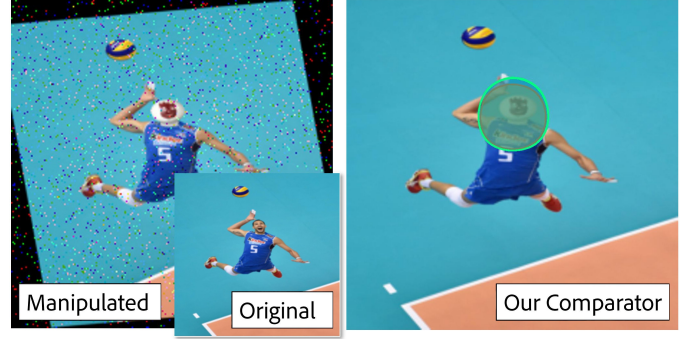

Fig. 1. Our deep image comparator is trained to highlight the differences between a pair of images due to editorial manipulation (here, hands up vs. down), whilst ignoring change due to benign transformations of the image during online distribution (here, warping and blurring). Output in green, ground-truth in yellow.

I. INTRODUCTION

Images are a great way to share stories and spread information. However, images can be easily manipulated to tell altered or even completely false stories. As both the number of images shared online each day and the ease of image manipulation grow, the need for tools to provide content provenance information rises. This is addressed in the recently introduced C2PA standards [1] which specifies how provenance information can be encapsulated as meta-data alongside the image content. If an image follows the C2PA standards, users can extract the entire edit story via its secondary stream meta-data.

This paper addresses a common scenario where meta-data is striped from an image during its online redistribution. It contributes a technique for robustly matching a query (without meta-data) to an original from a trusted database (with full meta-data), followed by an intuitive visualization of the image regions that have been manipulated to differ from the original.

Robust image matching poses many challenges. Images spread online are often subject to *benign* transformations such as changes to quality, resolution, aspect ratio, format *etc*. Additionally, we aim to match images that have been *manipulated* for editorial reasons that alter or falsify their stories (we also call this *editorial changes*, as opposed to *benign changes*). We note that cryptographic (bit-level) hashes cannot be relied for matching, nor can simple pixel difference operations be used to visualize changes due solely to manipulation. We propose SImProv - a robust and scalable content provenance framework that compliments C2PA. SImProv has two technical contributions:

Robust Near-Duplicate Image Search. We learn a visual search embedding that is robust to both benign transformations and content manipulations. We train a convolutional neural network (CNN) using a contrastive learning approach. We use a dataset of original photographs modified in Adobe Photoshop™, combined with data augmentations simulating benign image modifications. This yields a search embedding for robustly matching a near-duplicate *query* image circulating 'in the wild' to a trusted database of original images (hereon, we use the term 'near-duplicate' to refer to images that undergo certain transformations regardless of such transformations being benign or editorial changes).

An earlier version of SImProv was proposed at the CVPR workshop on Media Forensics 2021 [2]. The proposed method improves upon this using instance-level feature pooling methods to improve near-duplicate image search. We show that incorporating these into our image fingerprinting descriptor improves performance scalability, using a corpus of up to 100 million diverse photographic and artistic images from Behance.Net. These adaptations demonstrate the utility of our approach for web-scale content authenticity applications.

Pairwise Image Comparison. We propose a novel CNN architecture for pairwise image comparison that learns a joint image pair representation. We use this architecture to train two models for near-duplicated detection and editorial change localization respectively. In the near-duplicated detection model, the pair representation is used in conjunction with the individual visual search embeddings of both images to decide whether the two input images are two versions of the same image or completely unrelated *distinct* images. In the

editorial change visualization model, the pair representation is used to produce a *heatmap* that localizes visual discrepancies due to editorial manipulation. The network incorporates both a de-warping and image correlation module, and is trained end-to-end to ignore out-of-place transformation of content e.g. due to padding or warping as well as in-place corruption due to noise. In this extension of the earlier proposed pair-wise approach [2] we show that fusing end-to-end features from the image embedding together with the pair-wise embedding model improves the performance of the near-duplicate detection and re-ranking. These tasks were previously trained and applied as two sequential, entirely decoupled processes.

II. RELATED WORK

Visual content authenticity has been explored from the perspectives of both detection, and attribution.

Detection of visual tampering or generative ('deep fake') content [3] is typically a 'blind' detection problem. Given a single image, statistics may be learned to localize manipulated regions [4], [5], identify the use of a generative adversarial network (GAN) [6] or even determine (fingerprint) which GAN synthesized an image [7]. Detection of video manipulation similarly exploits temporal anomalies [8] or GAN limitations such as lack of blinking [9].

Image Attribution methods bind an image to data on its provenance, via embedded metadata [10], [11], watermarking [12]–[15], or perceptual hashing [16]–[19]. Emerging standards securely transport a cryptographically signed edit history within image metadata [1], [10], [11]. Yet image metadata is often stripped by social platforms, and may be replaced to misattribute an image [20]. Watermarking methods similarly embed provenance information, within image content. Both metadata and watermarking methods may instead embed a link to a trusted database (in some cases a blockchain [21]) containing the provenance data.

Perceptual hashing also keys into a trusted database using a robust content-aware hash for visual similarity search [22]. Classical approaches sample the spectral domain using wavelets or DCT coefficients [16], [23]. More recently, deep learning has been applied to learn robust visual hashes. Deep Hashing Networks (DHNs) [24] extended an ImageNet-trained AlexNet [25] feature encoder [26] with a quantization loss, to obtain hashes that retained semantic discrimination. CSQ [27] treats hashing as a retrieval/attribution optimization problem. Both DHN and CSQ but require pairwise labels or semantic annotation unavailable in our use case. Deep Supervised Hashing (DSH) [17] and HashNet [18] train CNNs to learn visual hashes, using a siamese network and ranking loss; such losses are used extensively in visual search [28]. DSDH [29] learns metric ranking and classification directly from the hash code. Our approach is aligned in the sense that we also apply deep metric learning, but differs in that we use contrastive training [30] and data augmentation to learn invariances relevant to benign and editorial image transformation. The image similarity detection challenge and dataset [31] focus on large-scale retrieval of images, subjected to benign transformations, but does not take into account editorial changes.

Localization of image manipulation focuses on blind detection tasks *e.g.* identifying image splicing [32] or use of photo-retouching tools [4]. Uniquely we approach the problem as a combination of perceptual hashing and pair-wise comparison. Our image comparator (the second contribution of this paper) assumes that a trusted 'original' image may be first uncovered by a visual search (the first contribution). Our comparator learns to ignores discrepancies due to benign image transformations, but is sensitized to editorial manipulations. This is achieved through a differential optical flow [33] and dewarping module into our two-stream architecture. Two-stream networks have been employed to predict the kinds of edit operation applied to a pair of images [34]. We differ by producing a heatmap of edit operations, de-sensitized to particular transformation classes. A further feature of our method is a classification score also available at inference to determine whether an image is a benign or manipulated version, or a different image.

III. METHOD

Our approach for image provenance assumes the existence of a trusted database $\mathcal{D} = \{I_1, I_2, ..., I_N\}$ containing N original images and their associated provenance information (*e.g.* curated by a trusted publisher, or via a decentralized immutable data-store such as a blockchain). Given a query image q, our goals are: (i) determining whether there exists an original version of q in $\mathcal{D}$; and (ii) localizing editorial changes if a match is found. The two goals appear to conflict each other since the former requires robustness to both benign and editorial changes while the latter should be sensitive to editorial manipulations. Learning a single model to achieve our goals is therefore extremely challenging. We instead propose a multi-stage framework. Our SImProv consists of 3 stages: (i) a visual search stage followed by (ii) re-ranking and near-duplicated detection, and finally (iii) detection and visualization of editorial changes (Fig. 2).

Firstly, in III-A we describe the representation learning process, used for near-duplicate image search (stage 1). We develop a model that learns 256-D representations of images that are further binarized into a 128-bit hash for scalable search [35]. The search is used to identify the most similar images to a users' query image.

Secondly, III-B describes the Pairwise Embedding Network (PEN) to deliver a pairwise representation of the query and a candidate image. PEN is our core design for the later stages of SImProv. In III-C, PEN is integrated to our stage-2 Pairwise Similarity Evaluation Network (PSEN) to re-rank the top k images (k=100) and identify the likelihood of candidate image being a near-duplicated version of the query, as opposed to being just another distinct image.

Finally, III-D describes how PEN is leveraged to identify whether the query image is a manipulated or benignly transformed version of the original (stage 3). If the query is identified as manipulated, we visualize a heatmap of the manipulated region on top of the image.

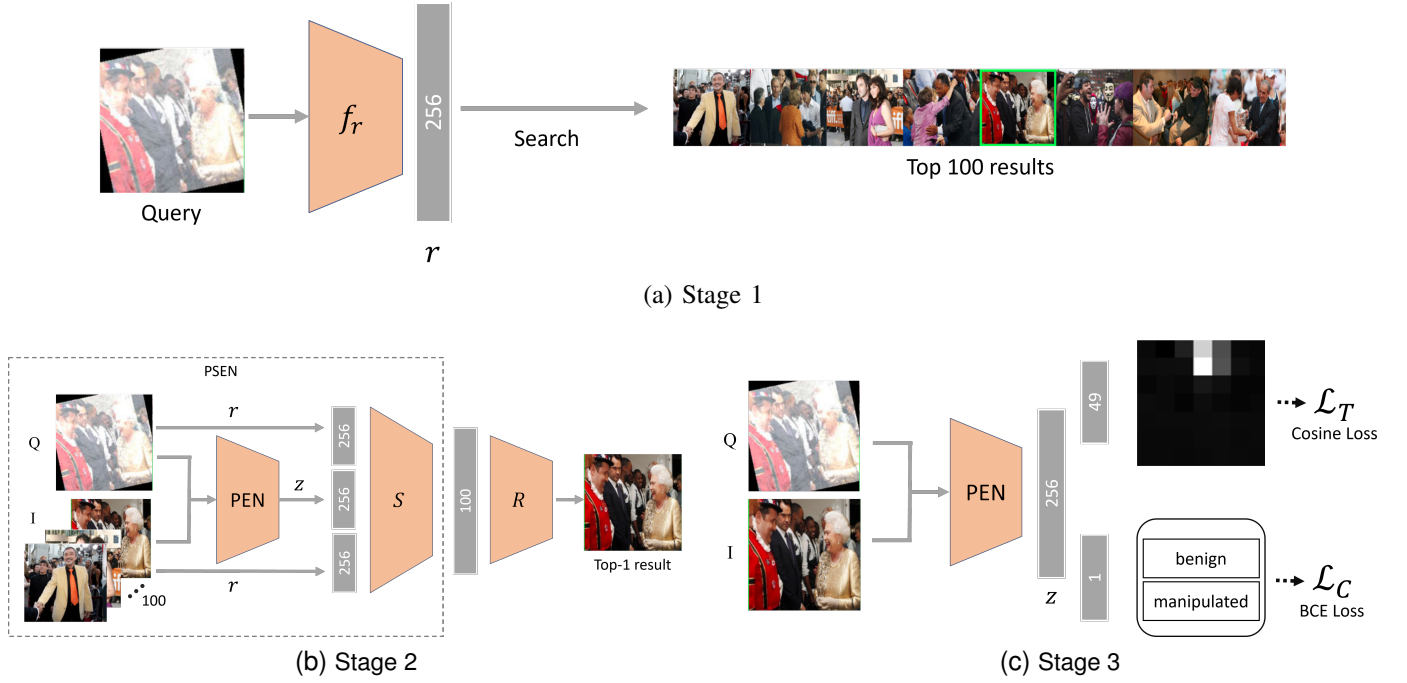(a) Stage 1

(b) Stage 2   (c) Stage 3

Fig. 2. Architecture diagrams of each of three stages of the proposed SImProv framework. Stage 1, near-duplicate image search uses a resolution-agnostic CNN to produce a feature embedding r that is further hashed for fast and scalable retrieval. Stage 2 performs re-ranking of the top-100 retrieved results, utilizing a pairwise embedding z to re-order the results and identify a correct match to the query. Finally, stage 3 uses the Pairwise Embedding Network (PEN) to identify whether the query image has been manipulated and localizes the manipulation with a heatmap.

A. Near-Duplicate Image Search

1) Representation Learning: We train a CNN model $f_r(.)$ to encode an image I into a compact embedding space $r = f_r(I) \in \mathbb{R}^{256}$. We employ the ResNet50 [36] backbone for f_r, replacing the final layer with a 256-D fully connected (fc) layer as the embedding. We use DeepAugMix [37] as the pretrained weight and finetune with loss:

$$\mathcal{L}(r) = -\log \frac{\sum_b e^{d(r,r_b^+)/\tau} + \sum_m e^{d(r,r_m^+)/\tau}}{\sum_i e^{d(r,r_i)/\tau}} \quad (1)$$

$$\text{where } d(u,v) = \frac{E(u) \cdot E(v)}{|E(u)| \, |E(v)|} \quad (2)$$

where $E(.)$ is a linear buffer layer between the embedding and loss; $d(u,v)$ measures the cosine similarity between the intermediate embeddings $E(u)$ and $E(v)$; τ is the contrastive temperature ($\tau = 0.8$). r_b^+ and r_m^+ refer to the embeddings of the benign-transformed and manipulated (also subjected to benign transformations) versions of image I respectively; while r_i is the embeddings of all images other than I in the mini-batch. Our loss resembles SimCLR [30] with 2 key differences: (i) our loss leverages multiple positive images for a given input image I instead of just one pair in [30]; and (ii) we adapt SimCLR in a near-duplicate retrieval problem treating images with editorial changes as positives. During training, we ensure both benign and manipulated versions of any given image I are present in a mini batch.

2) Feature Pooling: During inference, we find our model benefits from geometric pooling (GeM) [38] in two ways. Firstly, model becomes resolution agnostic and can take larger images as input, capturing more information. Secondly, it

allows to focus on local features, which is more beneficial for matching out of place transformed images.

For a set of K spatial feature map activations $\Phi = [\phi_1, \ldots \phi_K]$, the GeM [38] pooling operation G is defined as

$$G(\Phi) = [g_1, \ldots g_k, \ldots g_K] \quad (3)$$

$$g_k = \left(\frac{1}{|\phi_k|} \sum_{x \in \phi_k} x^{p_k}\right)^{\frac{1}{p_k}} \quad (4)$$

where p_k is a hyper-parameter (p_k is fixed at a default value of 3 in our experiments).

3) Hashing: Although our 256-D embedding r is already compact, it is difficult to scale search to millions of images and retain interactive speed at the same time. Inspired from [35], we further binarize the embedding features via a 2-step quantization:

$$b = q_1(r) + q_2(r - q_1(r)) \in \{0,1\}^D \quad (5)$$

where $q_1(.)$ is a coarse quantizer to allocate the feature r into one of several clusters, and $q_2(.)$ is a fine quantizer encoding the residual of z and its corresponding centroid. $q_1(.)$ behaves like an inverted list enabling search within a fraction of the database, while $q_2(.)$ delivers a compact binary code efficient for search in the Hamming space. We use KMeans with 1024 clusters for $q_1(.)$ also extending the search to nearby 10 clusters, and Product Quantization for $q_2(.)$ resulting in total a 128-bit descriptor.
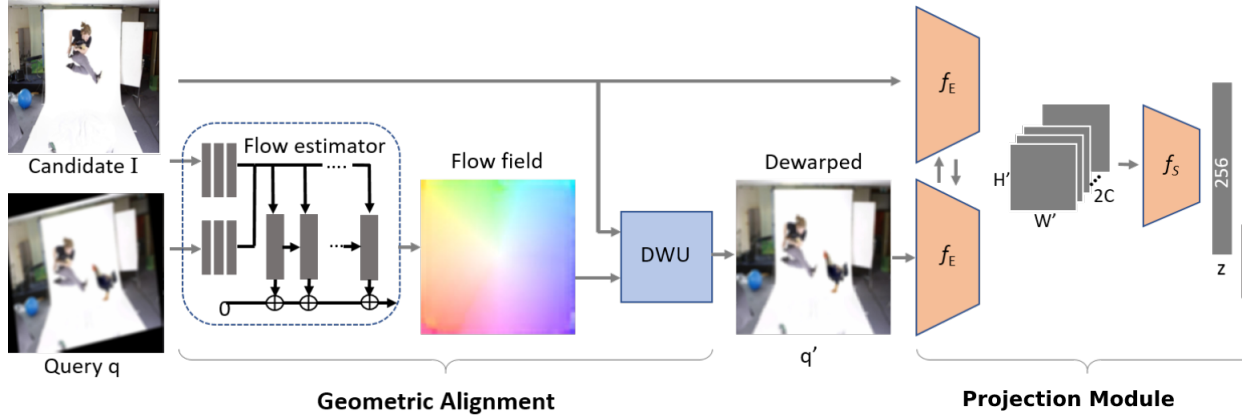
Fig. 3. Architecture of the proposed Pairwise Embedding Network (PEN). A candidate match to the user queried image is obtained from near-duplicate search (III-A, not shown). Image alignment is performed via differentiable de-warping unit (DWU) based on a dense optical flow estimate provided by the flow estimator. The resulting image pair are separately encoded via a feature extractor $f_E(.)$ and the concatenated features passed through $f_S(.)$ to obtain the combined feature z.

B. Pairwise Embedding Network

We propose an Pairwise Embedding Network (PEN) that learns a joint representation of two input images (Fig. 3). This architecture is later utilized for two purposes: near-duplicated detection (III-C1) and localization of editorial change (III-D), which corresponds to stage 2 and 3 of SImProv respectively.

The PEN accepts a pair of query-candidate images as input and outputs a n-dimensional ($n = 256$ in our experiments) representation of the image pair. The PEN architecture consists of 2 modules: a geometrical alignment module, $\mathcal{F}_A$, followed by a projection module, $\mathcal{F}_P$ (Fig. 3). Below we describe our designs for $\mathcal{F}_A$ and $\mathcal{F}_P$.

Geometric Alignment Module is used to account for the fact that the query q may undergo through out-of-place transformations which alter the pixel placement e.g. affine transformations or padding. We correct its alignment prior to joint representation learning. This is a crucial step especially for manipulation localization (as we show later in III-D). $\mathcal{F}_A$ comprises an optical flow estimator and a de-warping unit (DWU). We utilize the works in optical flow estimation field which aims to estimate the optical flow between video frames. Here we determine the alignment between two images instead. In theory, any flow estimation network that can be trained end-to-end is acceptable. In this work we use RAFT [33]. Supposed the query image q and the retrieved original image I are both resized to a fixed height (H) and width (W), RAFT identifies a dense pixel displacement field $\{\rho^x, \rho^y\} \in \mathbb{R}^{H \times W}$ from q to I by computing correlation between the per-pixel features from all pairs of pixels (see [33] for more details).

Our DWU then applies the predicted optical flow to the query for the best alignment to the candidate image:

$$M : (x, y) \mapsto (x + \rho^x(x), y + \rho^y(y)) \qquad (6)$$
$$\text{DWU}(q|\rho^x, \rho^y) = \mathcal{S}(M) \in \mathbb{R}^{H \times W} \qquad (7)$$

where (x, y) refers to the pixel coordinates in the query q which are mapped into its estimated correspondence M according to the optical flow $\{\rho^x, \rho^y\}$. $\mathcal{S}(.)$ is a bilinear sampler that effectively fits a local grid around M: $\mathcal{S}(M) =$

$\{M + \Delta M | \Delta M \in \mathbb{R}^2, |\Delta M| <= 1\}$ where output coordinates are computed by linear interpolation.

Projection Module takes the candidate I and the aligned query, $q' = \mathcal{F}_A(q|I)$ and outputs a single feature z. We first extract local features of each image using a shared CNN module:

$$z_q = f_E(q'); \ z_I = f_E(I) \in \mathbb{R}^{H' \times W' \times C} \qquad (8)$$

where H', W' and C are the new height, width and feature dimension respectively. Our feature extractor $f_E(.)$ is 3 convolution layers separated by ReLU, batch norm and max pooling. It outputs features at $\frac{1}{4}$ resolution ($H' = H/4, W' = W/4$ and we set $C = 128$). The combined features feed another CNN to learn a fusion representation z:

$$z = f_S([z_q, z_I]) \in \mathbb{R}^{256} \qquad (9)$$

where $[,]$ is concatenation, and $f_S(.)$ is formed from 4 ResNet residual blocks [36] followed by average pooling and a FC layer outputting 256-D features. PEN output is used for both re-ranking, near-duplicated detection and manipulation localization as described below.

C. Near-duplicate detection and Re-ranking

Our near-duplicate image search (III-A) method is designed to produce compact descriptors to enable interactive speeds in search through millions of images. However, the increase in speed comes at a cost of precision. The correct image could end up near the top of retrieval results, but in many cases might not in the first place. We propose a re-ranking model based on pairwise comparison of the query image with each of the top-k retrieval candidates. Such pairwise comparison is much slower and is not feasible for search through millions of images, but allows to identify the most likely match within a shortlist. The re-ranking consists of two steps: pairwise similarity evaluation and final reordering. Similarity evaluation produces a similarity confidence score for each of the 100 query-candidate image pairs. Re-ordering looks at the full list of 100 confidence scores and decides which of the candidate

images is the most likely match to the query. Below we describe similarity evaluation (III-C1) and reordering (III-C2).

1) Image Similarity Evaluation: We propose a Pairwise Similarity Evaluation Network (PSEN) that uses two images as input: the query image q and a candidate image c, retrieved by the near-duplicate search model (III-A). The PSEN uses the previously obtained individual embeddings of the images, as well as a PEN joint embedding learned from stack of two images together (Fig. 2b).

The final output of the model is a confidence score s, indicating the likelihood that the query and candidate images are the same image under different transformations:

$$s = S([f_r(q), f_r(c), PEN([q,c])]) \in [0,1] \qquad (10)$$

where $[,]$ is concatenation, $f_r(.)$ is the search embedding (III-A1), $PEN([.,.])$ is the pair representation (III-B) and $S(.)$ is a binary classification fully connected layer. The model is trained with binary cross-entropy loss.

2) Re-ordering: The role of re-orderer is to decide which of the candidate images is the most likely match to the query, based on two pieces of information: the initial raking of near-duplicate image search (distance between candidate and query embeddings) and similarity confidence scores. The index n of the most likely match is defined as

$$n = R(s_0, \ldots, s_{99}) \in [0, 100] \qquad (11)$$

where $s_i = S([f_r(q), f_r(c_i), PEN([q,c_i])])$ is the similarity confidence score between the query image and i-th retrieved candidate image; $R(.)$ is a neural network, consisting of three fully connected layers of sizes $[8192, 1024, 101]$, respectively. Re-orderer is trained as a 101-way classifier with cross-entropy loss. First 100 classes correspond to indices of candidate images and the final class indicates that the correct match to the query is not present within the candidate images.

D. Detecting and Localizing Editorial Change

This stage assumes a near-duplicated image to a query has been found after phase 2 (III-C). To predict the benign-editorial relationship and visualize the possible manipulated regions, we train a second PEN model with a combination of two losses (Fig. 2c). The first loss is a binary cross entropy predicting whether the pair is benign (*i.e.* the query q is either identical or a benign transformed version of the candidate I) or manipulated (*i.e.* z is a manipulated version of I):

$$c = E_c(z) \in \mathbb{R}^2 \qquad (12)$$

$$\mathcal{L}_C = -\log \frac{e^{c_y}}{\sum_{i=1}^2 e^{c_i}} \qquad (13)$$

where $E_c(.)$ is a FC layer projecting z to a 2-D feature c, and y is the classification target of the pair (q, I).

The second loss minimizes the cosine distance between the manipulation heatmap derived from z and the ground truth heatmap. We produce the heatmap at resolution $t \times t$ from z via a FC layer, $E_t(z) \in \mathbb{R}^{t^2}$, and compute loss:

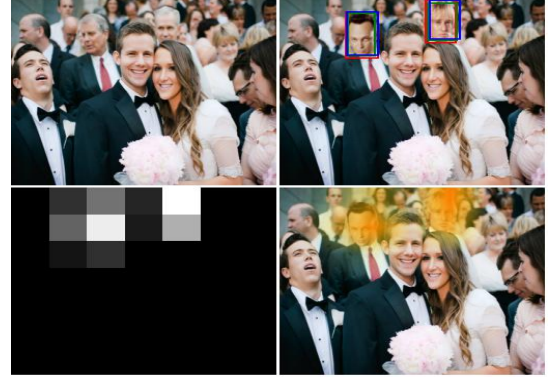$$\mathcal{L}_T = 1 - \frac{E_t(z) \cdot T}{|E_t(z)| \, |T|} \qquad (14)$$

Fig. 4. Examples of the crowd-annotation we collected on PSBattles to identify ground-truth (g-t) manipulated regions: Original image (top-left); manipulated image with MTurk annotations via bounding boxes (top-right); 7x7 ground truth (bottom-left), manipulated superimposed with the g-t heatmap (bottom-right).

where T is the ground truth manipulation heatmap. T is a matrix of zeros if the pair (q, I) is benign and if a manipulated pair $T \in [0,1]$ derived from human ground truth annotations (subsec. IV-A). We define the output heatmap resolution $t = 7$ during training. At test time, the 7×7 heatmap is interpolated to the original resolution $H \times W$ and super-imposed on the query image. The heat map is continuous but can be thresholded for more intuitive visualization.

The total loss is $\mathcal{L}(.) = w_c \mathcal{L}_C(.) + w_t \mathcal{L}_T(.)$ where loss weight $w_c = w_t = 0.5$ is set empirically.

IV. EXPERIMENTS AND DISCUSSION

We evaluate near-duplicate search, re-ranking and heatmap localization, as well as classification performance.

A. Datasets

We train and evaluate on PSBattles [39]; a dataset of images manipulated in Adobe Photoshop™, collected from the 'Photoshopbattles' forum on Reddit. The dataset contains more than 10k original images and, for each of these, several manipulated variants; in total 102,028 variants contributed by 31K artists. We remove the original-manipulated pairs that are obviously different, retaining only similar pairs $\mathcal{D} = \{(O_i, P_i)| \, ||f(O_i) - f(P_i)||_2 < \beta\}$ where $f(.)$ is a pretrained ImageNet ResNet50 feature extractor and $\beta = 150$ is the distance threshold. That leaves 7,171 originals and 24,157 manipulated images. The data is split into training (**PSBat-Train**) and test (**PSBat-Test**) sets, the former has 6,364/21,197 and the latter has 807/2,960 original/manipulated images. The PSBat-Train is used to train our image retrieval, similarity evaluation and edit localization models while PSBat-Test is used for the two benchmarks below.

PSBat-Ret. We construct a database of 807 original images from PSBat-Test plus 2 million diverse distractor images scraped from the Adobe Stock website. Next we created 3 query sets: (i) *Manip* contains 2,960 manipulated images in PSBat-Test; (ii) *Benign* contains 29.6k images created by transforming the PSBat-Test original images; and (iii)

Manip+Benign also contains 29.6k images but via transforming the manipulated set instead. To obtain Benign and Manip+Benign, we applied a suite of benign transformations common in online image re-distribution. These include JPEG compression (40%-90%), random crop (90% area), padding (max 10% each side), rotation (max 15 degree), flipping and ImageNet-C [40] transformations containing various additive noise (*e.g.* Gaussian, shot, impulse noise) and blur (*e.g.* Gaussian, motion, defocus blur) and enhancement (*e.g.* brightness, contrast, snow) for all 5 severity levels in ImageNet-C [40]. To test different input resolutions, we apply augmentations to 1024x1024 images and then produce all other resolutions by downscaling the resulting image. However, ImageNet-C parameters are set up with 224x224 resolution in mind, so we adjust the augmentation specific parameters in such a way that applying them to 1024x1024 image and downscaling it to 224x224 produces an output that is as close as possible to directly applying standard augmentation directly to 224x224 images. We divide benign transformations to 3 groups: the *primary* group contains resize and JPEG re-compression; the *in-place* group contains in-place transformations from ImageNet-C transformations; and the *out-place* group contains those transformations that change pixel coordinates such as padding and affine warps. When transforming an image we apply all those in the primary group, followed by a random transformation in either the in-place group, or out-place group, or both.

Training and evaluating edit localization model requires labeled manipulated regions. We identify these regions via crowd sourced annotation. For each original-manipulated pair in PSBat-Train/Test, 3 workers draw bounding boxes around the manipulated areas, obtaining a binary heatmap each $G^k = \{0,1\}^{H \times W}, k = 1, 2, 3$ where $G^k(x, y) = 1$ if pixel (x, y) is contained in a bounding box drawn by worker k. We normalize G^k w.r.t 7x7 image size and combine to a ground truth heatmap $T = \{\sum_k G^k(x, y)/3 \in \mathbb{R}^{7 \times 7}\}$ (Fig. 4).

Additionally, we also evaluate the retrieval efficacy of SImProv on BAM-100M, a large scale dataset consisting of 100M artworks from Behance [1]. BAM-100M is significantly larger and more diverse than ImageNet, since its collection spans many fields beyond photography, such as paintings, graphic designs, advertising and graffiti. We note this is the largest experiment in term of dataset size for image provenance to date. We create two query sets, BAM-Q-Res and BAM-Q-Aug from 1K images sampled at random from BAM-100M. To make BAM-Q-Res, we downscale images at random ratio in range 0.1-0.9 (up to 10x downscaling) with bilinear interpolation keeping aspect ratio. To make BAM-Q-Aug, we apply the same augmentation strategy as in PSBat-Ret.

B. Metrics

To evaluate near duplicate search, we use Instance Retrieval IR@k metric which measures the ratio of queries that returns the relevant images within top-k retrieval. Formally, $IR@k = \frac{1}{Q} \sum_{i=1}^{Q} \sum_{j=1}^{k} r(q_i, j)$ where Q is number of queries, relevance function $r(q_i, j) = 1$ if the returned image at rank j

Fig. 5. Heatmaps showing manipulation of an original image (inset) at threshold 0.35. The region of manipulation is correctly identified both without (top) and with (bottom) benign transformation of the manipulated version.

is relevant to the query q_i (there is only one such image in PSBat-Ret), otherwise 0.

We use Average Precision (AP) to measure the accuracy of both classifiers: the same/different similarity evaluation network and benign/manipulated classifier branch of the edit localization network.

For the generated heatmap, we up-sample the 7x7 heatmap to the image resolution $H \times W$, convert to binary with a threshold and compute Intersection over Union (IoU) with the ground truth, $IoU = \frac{1}{Q} \sum_{i=1}^{Q} \frac{S(U_i) \cap T_i}{S(U_i) \cup T_i}$ where T_i is the $H \times W$ binary ground truth heatmap, U_i is the predicted heatmap after interpolation and thresholding. We leverage the image pair classification result to improve the heatmap with $S(U_i) = U_i$ if the query is classified as manipulated, $\{0\}^{H \times W}$ if benign and $\{1\}^{H \times W}$ if distinct.

C. Evaluating Near-Duplicate Search

We compare our retrieval method (both before and after re-ranking) against 9 baselines. **ICN** [2] is our initial workshop version of SImProv. **ImageNet** [36], **MSResNet** [41] and **DeepAugMix** [37] are 3 public pretrained CNN models, all use ResNet50 architecture. The classic ImageNet model is trained on ILSVRC2012 [42], MSResNet is built by Microsoft to power its Bing image search engine while DeepAugMix reports state of art performance on the ImageNet-C benchmark. **ImageNet fine.** and **MSResNet fine.** are the finetuned models on PSBat-Train using our training strategy stated in subsec. III-A, as compared with finetuning DeepAugMix for **SImProv (stage 1)**. **CSQ** [27] and **HashNet** [18] are two supervised class-level online hashing methods. For fair comparison, we train these models using the same CNN backbone (ResNet50) with the same data augmentation strategy (sec. IV-A). **pHash** [43] is a classical image hashing method using relative DCT coefficients. All methods produce 128-bit hash code except pHash (64-bit).

Tab. I compares retrieval performance. The two online hashing methods, CSQ [27] and HashNet [18], are among the worst performers. CSQ and HashNet struggle to cope with strong ImageNet-C transformations present during training

TABLE I
RETRIEVAL PERFORMANCE (ON 2M IMAGES, PSBAT-RET) REPORTED AS IR SCORE AT RANKS [1,10,100], FOR QUERY IMAGES SUBJECTED TO BENIGN TRANSFORMS, MANIPULATION, OR BOTH. STAGE 1 REFERS TO NEAREST-NEIGHBOR SEARCH ONLY.

| Method	Benign			Manip			Manip+Benign			Average		
	IR@1	IR@10	IR@100	IR@1	IR@10	IR@100	IR@1	IR@10	IR@100	IR@1	IR@10	IR@100
SImProv (stage 1 + 2)	**0.9746**	**0.9849**	0.9849	**0.9170**	**0.9564**	0.9564	**0.9142**	**0.9271**	**0.9271**	**0.9353**	**0.9561**	**0.9561**
SImProv (stage 1)	0.9450	0.9749	0.9849	0.8838	0.9307	**0.9564**	0.8064	0.8845	**0.9271**	0.8784	0.9300	**0.9561**
ICN [2] (stage 1 + 2)	0.9423	0.9811	**0.9867**	0.9154	0.9453	0.9453	0.8705	0.9099	0.9164	0.9094	0.9454	0.9494
ICN [2] (stage 1)	0.9305	0.9725	**0.9867**	0.8557	0.9061	0.9453	0.7662	0.8582	0.9164	0.8508	0.9123	0.9494
MSResNet fine.	0.7532	0.7931	0.8258	0.9199	0.9500	0.9655	0.6754	0.7406	0.7884	0.7828	0.8279	0.8599
ImageNet fine.	0.7709	0.8988	0.9577	0.7791	0.8689	0.9159	0.5641	0.7348	0.8420	0.7047	0.8342	0.9052
MSResNet [41]	0.7635	0.8121	0.8450	0.9551	0.9753	0.9797	0.6703	0.7408	0.7895	0.7963	0.8428	0.8714
DeepAugMix [37]	0.6944	0.7743	0.8336	0.8956	0.9385	0.9615	0.5629	0.6725	0.7553	0.7176	0.7951	0.8501
ImageNet [36]	0.6954	0.7310	0.7564	0.9649	0.9703	0.9726	0.6134	0.6620	0.6956	0.7579	0.7878	0.8082
CSQ [27]	0.1390	0.1803	0.3292	0.2095	0.2584	0.4122	0.0628	0.0957	0.2214	0.1371	0.1781	0.3209
HashNet [18]	0.2020	0.2521	0.2924	0.2841	0.3324	0.3733	0.0894	0.1294	0.1697	0.1918	0.2380	0.2785
pHash [23]	0.3674	0.3731	0.3768	0.3693	0.3753	0.3764	0.1739	0.1821	0.1857	0.3035	0.3102	0.3130
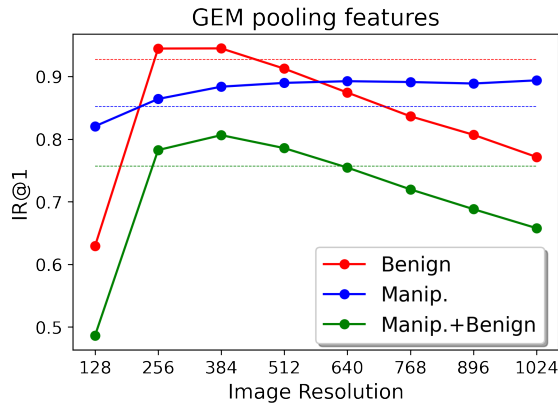
Fig. 6. Top-1 performance of GeM features on different input image resolutions. Dash-lines represent the performance without using feature pooling (*i.e.* using only the output of the last FC layer [2]) on the corresponding test sets.

and test, resulting in lower performance than the classical pHash. ImageNet [36], MSResNet [41] and DeepAugMix [37] perform strongly on the Manip set but poorly when they undergo benign transformations. When trained via our contrastive loss (eq. 1), all models gain with our proposed *SImProv (stage 1)* achieving 25% improvement on Benign IR@1 and 24% on Manip+Benign versus the pretrained DeepAugMix model. *SImProv (stage 1)* also outperforms the finetuned ImageNet/MSResNet by a large margin on all top-k scores and query sets. The improvement of SImProv compared to the one presented in the workshop paper (*ICN (stage 1)*) [2] can be attributed to geometric pooling, which allows higher resolution input. We demonstrate significant performance improvement of the proposed re-ranking method (*SImProv (stage 1+2)*) compared to the naive re-ranking approach used in *ICN (stage 1+2)*. Fig. 10(a) shows retrieval examples for SImProv and its closest competitor ICN, for both benign (first row) and manipulated (second row) queries.

D. Feature Pooling

We show the superiority of GeM features against the traditional output features from the last FC layer of the retrieval model [2] as well as its dependence on input image resolution

TABLE II
TOP-1 RETRIEVAL PERFORMANCE OF GEM AND RMAC FEATURES AND OUTPUT FROM THE LAST FC LAYER ON THE 384x384 RESOLUTION TEST SETS.

	RMAC		GeM		CNN
	L=3	L=4	L=3	L=4	
Benign	0.9503	0.9444	0.9450	0.9441	0.9272
Manip.	0.8777	0.8777	0.8838	0.8926	0.8520
Manip.+Benign	0.8025	0.8027	0.8064	0.8068	0.7570

in Fig. 6. GeM features work best at 384x384 resolution, outperforming [2] by 4% on the challenging Manip.+Benign test set. The presence of benign transformations hampers GeM performance as the resolution increases, underperforming [2] from 512x512 resolution on Benign set and from 640x640 on the Manip.+Benign set.

Tab. II compare GeM with [2] and a similar feature pooling method - RMAC [44], at two pooling levels, L=3 and L=4. It can be seen that the pooling level does not affect much the performance of both GeM and RMAC. Additionally, RMAC is comparable to GeM, slightly outperforming GeM on the Benign set at L=3 but underperforming on the Manip. set at L=4. However, we choose GeM as the proposed method since it is significantly faster than RMAC. It takes RMAC 28.72 seconds to perform 1000 iterations, while GeM is $\sim$ 18 times faster, with just 1.63 for the same setup.

E. Evaluating Localization of Editorial Changes

We compare the localization performance of the proposed method against four baselines. **Sum of Squared Distances (SSD)** - we simply compute SSD between two images at pixel level, resize it to 7×7 then resize it back before thresholding to create continuity in the detected heatmap. **ResNetConv** - we extract $7 \times 7 \times 2048$ features from pretrained ImageNet ResNet50 model for both query and original images. These are averaged across channels to produce a 7×7 heatmap. **ErrAnalysis** - inspired from the blind detection technique in [5], we perform JPEG compression on the query image and compare with itself. **MantraNet** - is a supervised blind detection method [8] that detects anomalous regions. Additionally we evaluate baselines with images passed through our alignment module.

Manipulated (Original inset)   SImProv (ours)   ICN   ResNetConv   SSD   MantraNet   ErrAnalysis
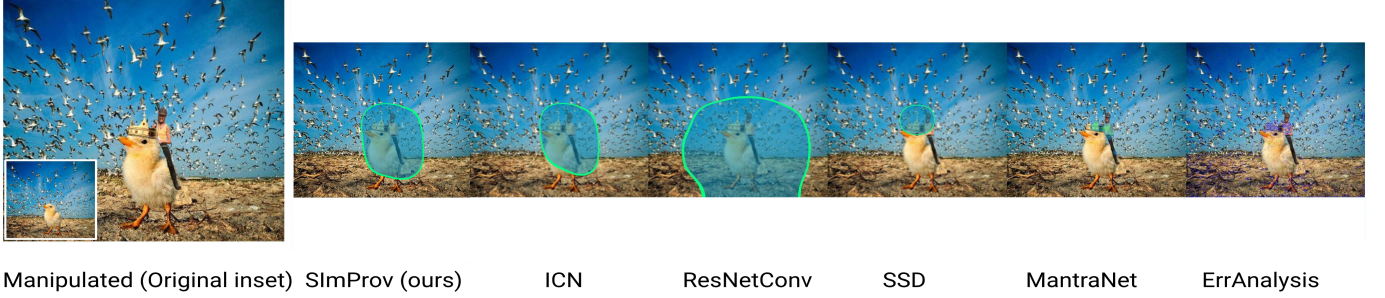
Fig. 7. Comparison of heatmap visualizations (in green) from our method, and baseline methods for thresholded (top) and non-thresholded (bot.) heatmaps. The heatmap visualizes manipulation of an image (crown/rider added on bird).

Fig. 8. Heatmap results. Left col.: Original image. Middle col.: Manipulated image also subjected to benign transformation. Right col.: Heatmap output (green) ignoring benign transformation and highlighting manipulation (ground truth in yellow).

TABLE III

EVALUATING HEATMAP ACCURACY AND INTREPRETABILITY FOR BASELINE METHODS. OUR PROPOSED SIMPROV METHOD IS COMPARED AGAINST BASELINES BOTH OBJECTIVELY FOR ACCURACY (IoU) AND SUBJECTIVELY VIA USERS TO DETERMINE WHICH EXHIBITS BEST INTREPRETABILITY (% METHOD PREFERENCE). $+\mathcal{F}_A$ INDICATES GEOMETRIC ALIGNMENT MODULE APPLIED.

Method	Accuracy (IoU)	Interp. (%)
Ours	**0.565**	**84.6**
ResNetConv+Geo. Align. $\mathcal{F}_A$	0.243	5.50
ResNetConv [36]	0.238	2.20
SSD+Geo. Align. $\mathcal{F}_A$	0.231	3.30
SSD	0.149	0
ErrAnalysis+Geo. Align. $\mathcal{F}_A$	0.143	0
ErrAnalysis [5]	0.109	0
MantraNet+Geo. Align. $\mathcal{F}_A$	0.061	2.75
MantraNet [8]	0.027	1.65

1) Localization Accuracy: We compare the heatmaps generated by our SImProv with baseline methods. Heatmaps are produced by upsampling the 7×7 heatmap output of the SImProv to the size of the image using bicubic interpolation. Heatmaps may be presented on false-colour scale (*e.g.* jet) in this form, or thresholded to produce an outline of the predicted manipulated region. In our experiments, we threshold the normalized heatmaps at 0.35 determined empirically (Fig. 5). Tab. III (first column) reports the IoU metric between the predicted heatmap and the ground truth, both with and without

the thresholding. Whilst most baselines are improved through use of our geometric alignment ($\mathcal{F}_A$) process, our SImProv significantly exceeds baseline performances by at least 0.30. Fig. 7 shows performance SImProv and other baselines on an example image. More change localization examples for SImProv is shown in Fig. 8.

2) Heatmap Interpretability: Heatmap Interpretability is assessed against baseline methods via a crowd-sourced study on Amazon Mechanical Turk (MTurk). Participants see an original image, and an image subjected to both manipulation and benign transformation. The latter is annotated with the ground truth as a guide. The participants are shown a grid of heatmaps generated by 9 methods: ours, 4 baselines SSD, MantraNet, ErrAnalysis and ResNetConv, and 4 warp-corrected baselines pre-applying $\mathcal{F}_A$ for geometric alignment. Participants indicate which of the 9 heatmaps best summarizes the image modification; 200 such tasks are each annotated by 5 unique participants.

Tab. III (final col.) presents the results, which favor our proposed method, even when the image pair are pre-aligned.

F. Large Scale Retrieval

We evaluate the scalability of our method by indexing the BAM-100M database. We compare the $IR@k$ performance of SImProv to its earlier version ICN [2]. Fig. 9 shows $IR@k$ versus database size curves of SImProv and ICN on BAM-100M with BAM-Q-Res as the query set. We demonstrates that SImProv's performance does not degrade nearly as much as ICN with increase in database size. For SImProv, the $IR@k$ remains nearly 1.0 at all image database sizes, dipping to 0.999 for the most challenging case of $IR@1$ for database sizes above 30M. ICN, on the other hand, is much greater affected by database size, with $IR@1$ dropping from 0.997 at 1M images to 0.985 at 100M images. Results for the more challenging query set BAM-Q-Aug are depicted in Fig. 11. SImProv outperforms ICN by a large margin at early k values, on both BAM-100M and a subset of 1M images. The performance drop when increasing the database size from 1M to 100M for SImProv is also lower than ICN. The IR@k curves converge as k value reaches 100, and saturated performance is achieved at IR@100 for both methods regardless of database size, which justifies our design choice

TABLE IV
SIMPROV STAGE 2 - SAME / DIFFERENT CLASSIFIER

Test	Average Precision (AP)	
	ICN [2]	Ours
Original	1.000	1.000
Benign	0.9996	1.000
Manip.	0.9976	0.9922
Benign+Manip.	0.9962	0.9909
Distinct	0.9895	0.9973

TABLE V
SIMPROV STAGE 3 - BENIGN / MANIP CLASSIFIER

Test	Average Precision (AP)	
	ICN [2]	Ours
Original	1.000	1.000
Benign	0.9635	0.9800
Manip.	0.9726	0.9932
Benign+Manip.	0.8807	0.9793

of selecting top-100 images for SImProv subsequent stages. Fig. 10(b) shows examples of retrieval results for SImProv and ICN. SImProv has better retrieval performance in the case of severely distorted queries.
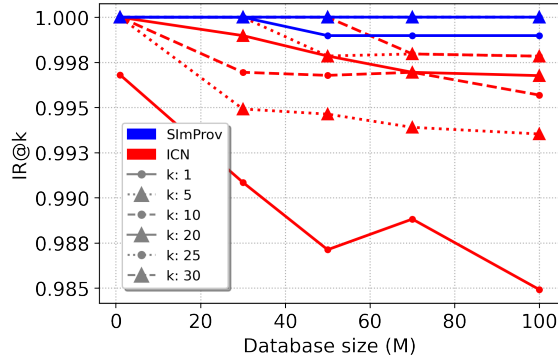

Fig. 9. Top-k retrieval performance comparison of ICN [2] and SImProv versus database size, using BAM-Q-Res queries.

G. Evaluating Classification

We evaluate the classification performance of two classifiers: same/different in SImProv stage 2 and bengin/manipulated in stage 3. Same/different classification is an output of PSEN III-C2, which classifies a pair of images as either being two entirely different images, or the same image, potentially under different transformations. Benign/manipulated classification is an output of the change localization network III-D, which assumes that the input images are not distinct and focuses on classifying whether the differences between them are benign or editorial. We compare the performance of our approach with ICN [2], which has a single 3-way (benign, manipulated, distinct) classifier. In case of same/different evaluation, we combine the confidences of 'benign' and 'manipulated' to count as 'same'.

We evaluate the performance of the 2-way same/different classification by comparing each original image in the test set with: itself, benign transformed version of itself, manipulated

version, manipulated as well as benign transformed and an entirely different image, chosen at random. All of the cases except the last are expected to be classified as 'same' and the last one as 'distinct'. Tab. IV shows the Average Precision (AP) scores achieved for each case. A non-modified original-original pair is always correctly classified as the same image by both both methods. Introduction of benign transformations reduces the accuracy of ICN slightly, but does not affect our approach. The most challenging case is queries that are both manipulated and benign transformed, however both methods maintain AP near 0.99 in all of the cases.

The bigger difference in performance can be seen in Tab. V which shows the AP scores for benign/manipulated classification. Here, our approach outperforms ICN by 2% in the cases where the query image is either just benignly transformed or just manipulated. The difference in performance grows to 9.9% when the query is both manipulated and benign transformed.

H. Re-Ranker Ablations

We perform ablations to determine the impact of each component of the proposed stage 2 re-ranking architecture. Results are shown in Tab. VI. Following the same experimental setup as in IV-C, we report $IR@1$ for the most challenging *manip+benign* dataset split. The result with all three components in use is equivalent to the architecture shown in Fig. 2b and all components turned off indicates only stage 1 retrieval is used and no re-ranking is performed. When using only f_r or PEN, the similarity confidence score from Eq. 10 is predicted by using only r or z features, respectively. In the cases with no re-orderer R, the order of images is sorted by PSEN confidence scores.

We observe that simply passing through the same r features that were initially used to search for the images does not provide any new information, and so the performance remains the same. However, using a pairwise feature from PEN improves the retrieval by nearly 5%. Using both r or z features provides the model with additional context and further improves the performance by 2%. Finally, the re-orderer step improves the performance by additional 4%. This shows that each of the components of stage 2 re-ranking plays a significant role in SImProv retrieval performance.

TABLE VI
ABLATION STUDY OF STAGE 2 RE-RANKING, EXCLUDING ONE OR SEVERAL COMPONENTS OF THE ARCHITECTURE.

f_r	PEN	R	IR@1
✓	✓	✓	**0.9142**
✓	✓	✗	0.8762
✗	✓	✗	0.8544
✓	✗	✗	0.8064
✗	✗	✗	0.8064

V. LIMITATIONS

Fig. 12 illustrates failure cases of SImProv. If the degradation is very severe it will not be ignored by the model's trained invariance, and spurious additional detection (top-left) may occur. The 7×7 heatmap activations are of insufficient

(a)

(b)

Fig. 10. Retrieval examples on PSBat-Ret (a) and BAM-100M (b). For each retrieval example, the left-most image is the query, followed by top-10 returned results for SImProv (top row) and ICN (bottom row). Green bounding box indicates relevant result.
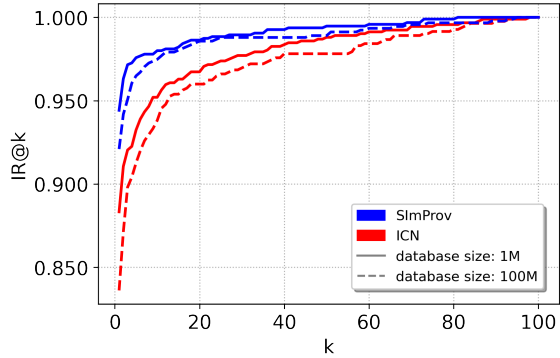

Fig. 11. Retrieval performance comparison of ICN [2] and SImProv on a 1M subset and fullset of BAM-100M, using BAM-Q-Aug queries.

Fig. 12. SImProv Limitations. Top-left: Spurious detections due to benign degradation; Top-right: Mismatch due to annotators missed the removed object in the ground-truth; Bot-left: Heatmap unable to separate many small manipulations. Bot-right: Missed detections due to poor geometric alignment.

resolution to separate many small manipulations (bot-left). Inaccurate ground-truth (all 3 annotators missed the removed object) gives an artificially low IoU (top-right). Geometric alignment may introduce artifacts that the model interprets as manipulations (bot-right).

VI. CONCLUSION

We presented a Scalable Image Provenance (SImProv) framework for large-scale retrieval and visual comparison of a pair of images in order to detect and localize manipulated regions. SImProv enables users to match images circulating 'in the wild' to a trusted database of original images. Given a query and matched original, the SImProv visualizes areas of manipulation as a 'heatmap'. The heatmap ignores artifacts due to benign transformations that commonly occur as images are reshared. We introduced two main architecture changes compared to an earlier version of the work [2]: incorporation of instance-level feature pooling for image retrieval and combination of individual and pairwise descriptors for near-duplicate detection, followed by re-ranking. We show that feature pooling improves retrieval performance by enabling the use of queries of larger resolution. We use a large corpus of 100 million diverse images to demonstrate that these changes improve retrieval performance and make our approach applicable to the web-scale content authenticity problem.

ACKNOWLEDGMENTS

REFERENCES

[1] Coalition for Content Provenance and Authenticity, "Draft technical specification 0.7," C2PA, Tech. Rep., 2021. [Online]. Available: https://c2pa.org/public-draft/

[2] A. Black, T. Bui, H. Jin, V. Swaminathan, and J. Collomosse, "Deep image comparator: Learning to visualize editorial change," in *Proc. CVPR WS*. IEEE, 2021, pp. 972–980.

[3] B. Dolhansky, J. Bitton, B. Pflaum, J. Lu, R. Howes, M. Wang, and C. C. Ferrer, "The deepfake detection challenge (DFDC) dataset," *CoRR*, vol. abs/2006.07397, 2020. [Online]. Available: http://arxiv.org/abs/2006.07397

[4] S.-Y. Wang, O. Wang, A. Owens, R. Zhang, and A. Efros, "Detecting photoshopped faces by scripting photoshop," in *Proc. ICCV*, 2019.

[5] W. Wang, J. Dong, and T. Tan, "Tampered region localization of digital color images based on jpeg compression noise," in *Intl. WS Digital Watermarking*. Springer, 2010, pp. 120–133.

[6] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. Efros, "Cnn-generated images are surprisingly easy to spot... for now," in *Proc. CVPR*, 2020.

[7] N. Yu, L. Davis, and M. Fritz, "Attributing fake images to gans: Learning and analyzing gan fingerprints," in *Proc. ICCV*, 2019.

[8] Y. Wu, W. AbdAlmageed, and P. Natarajan, "Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features," in *Proc. CVPR*, 2019, pp. 9543–9552.

[9] Y. Li, M.-C. Ching, and S. Lyu, "In ictu oculi: Exposing ai created fake videos by detecting eye blinking," in *Proc. IEEE WIFS*, 2018.

[10] C. A. I. (CAI), "Setting the standard for content attribution," Adobe Inc., Tech. Rep., 2020.

[11] J. A. et al., "Multi-stakeholder media provenance management to counter synthetic media risks in news publishing," in *Proc. Intl. Broadcasting Convention (IBC)*, 2020.

[12] S. G. K. Hameed, A. Mumtax, "Digital image watermarking in the wavelet transform domain," *WASET*, vol. 13, pp. 86–89, 2006.

[13] P. Devi, M. Venkatesan, and K. Duraiswamy, "A fragile watermarking scheme for image authentication with tamper localization using integer wavelet transform," *J. Computer Science*, vol. 5, no. 11, pp. 831–837, 2019.

[14] D. Profrock, M. Schlauweg, and E. Muller, "Content-based watermarking by geometric wrapping and feature-based image segmentation," in *Proc. SITIS*, 2006, pp. 572–581.

[15] S. Baba, L. Krekor, T. Arif, and Z. Shaaban, "Watermarking scheme for copyright protection of digital images," *IJCSNS*, vol. 9, no. 4, 2019.

[16] T. Pan, "Digital-content-based identification: Similarity hashing for content identification in decentralized environments," in *Proc. Blockchain for Science*, 2019.

[17] H. Liu, R. Wang, S. Shan, and X. Chen, "Deep supervised hashing for fast image retrieval," in *Proc. CVPR*, 2016, pp. 2064–2072.

[18] Z. Cao, M. Long, J. Wang, and P. S. Yu, "Hashnet: Deep learning to hash by continuation," in *Proc. CVPR*, 2017, pp. 5608–5617.

[19] F. Khelifi and A. Bouridane, "Perceptual video hashing for content identification and authentication," *IEEE TCSVT*, vol. 1, no. 29, 2017.

[20] I. Council, "Social media sites photo metadata test results," http://embeddedmetadata.org/social-media-test-results.php, 2020.

[21] J. Collomosse, T. Bui, A. Brown, J. Sheridan, A. Green, M. Bell, J. Fawcett, J. Higgins, and O. Thereaux, "ARCHANGEL: Trusted archives of digital public documents," in *Proc. ACM Doc.Eng*, 2018.

[22] J. Wang, H. T. Shen, J. Song, and J. Ji, "Hashing for similarity search: A survey," *arXiv preprint arXiv:1408.2927*, 2004.

[23] C. Zauner, "Implementation and benchmarking of perceptual image hash functions," Master's thesis, Upper Austria University of Applied Sciences, Hagenberg, 2010.

[24] H. Zhu, M. Long, J. Wang, and Y. Cao, "Deep hashing network for efficient similarity retrieval," in *Proc. AAAI*, 2016.

[25] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convnets," in *Proc. NIPS*, 2012.

[26] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping," in *Proc. CVPR*, 2006, pp. 1735–1742.

[27] L. Yuan, T. Wang, X. Zhang, F. Tay, Z. Jie, W. Liu, and J. Feng, "Central similarity quantization for efficient image and video retrieval," in *Proc. CVPR*, 2020, pp. 3083–3092.

[28] A. Gordo, J. Almazán, J. Revaud, and D. Larlus, "Deep image retrieval: Learning global representations for image search," in *Proc. ECCV*, 2016, pp. 241–257.

[29] Q. Li, Z. Sun, R. He, and T. Tan, "Deep supervised discrete hashing," in *Proc. NeurIPS*, 2017, pp. 2482–2491.

[30] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. ICML*, 2020, pp. 1597–1607.

[31] M. Douze, G. Tolias, E. Pizzi, Z. Papakipos, L. Chanussot, F. Radenovic, T. Jenicek, M. Maximov, L. Leal-Taixé, I. Elezi, O. Chum, and C. C. Ferrer, "The 2021 image similarity dataset and challenge," 2021. [Online]. Available: https://arxiv.org/abs/2106.09672

[32] M.Huh, A. Liu, A. Owens, and A. Efros, "Fighting fake news: Image splice detection via learned self-consistency," in *Proc. ECCV*, 2018.

[33] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *Proc. ECCV*. Springer, 2020, pp. 402–419.

[34] S. Jenni and P. Favaro, "Self-supervised feature learning by learning to spot artifacts," in *Proc. CVPR*, 2018.

[35] J. Johnson, M. Douze, and H. Jegou, "Billion-scale similarity search with gpus," *IEEE Transactions on Big Data*, 2019.

[36] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.

[37] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. Zhu, S. Parajuli, M. Guo *et al.*, "The many faces of robustness: A critical analysis of out-of-distribution generalization," *arXiv preprint arXiv:2006.16241*, 2020.

[38] F. Radenović, G. Tolias, and O. Chum, "Fine-tuning cnn image retrieval with no human annotation," *IEEE TPAMI*, vol. 41, no. 7, pp. 1655–1668, 2018.

[39] S. Heller, L. Rossetto, and H. Schuldt, "The PS-Battles Dataset – an Image Collection for Image Manipulation Detection," *CoRR*, vol. abs/1804.04866, 2018. [Online]. Available: http://arxiv.org/abs/1804.04866

[40] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," in *Proc. ICLR*, 2019.

[41] Z. Lenyk and J. Park, "Microsoft vision model resnet-50 combines web-scale data and multi-task learning to achieve state of the art," https://pypi.org/project/microsoftvision/, 2021.

[42] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *Proc. CVPR*, 2009.

[43] J. Buchner, "Imagehash," https://pypi.org/project/ImageHash/, 2020.

[44] G. Tolias, R. Sicre, and H. Jégou, "Particular object retrieval with integral max-pooling of cnn activations," in *Proc. ICLR*, 2016, pp. 1–12.
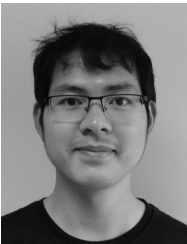
VII. Biography Section

Alexander Black is a PhD candidate at the Centre for Vision Speech and Signal Processing (CVSSP), University of Surrey, United Kingdom. He received BSc in Physics at University of Birmingham in 2019 and MSc in Computer Vision and Machine Learning at University of Surrey in 2020. He is currently a Research Scientist Intern at Adobe Research. His research interests include representation learning, media forensics and multi-modal deep learning.

John Collomosse is a principal scientist leading the Deep Learning Group at Adobe Research, where he works closely with the Content Authenticity Initiative (CAI). John is concurrently a full professor at the University of Surrey, where he founded and co-directs the DECaDE centre exploring the intersection of AI and Blockchain including applications to content provenance and media authenticity.

Tu Bui is currently a Research Fellow at the Centre for Vision Speech and Signal Processing (CVSSP), University of Surrey, United Kingdom. He received BEng and PhD degrees in electrical and electronic engineering at University of Surrey in 2014 and 2019, respectively. His research interests include representation learning, content provenance and cross-domain retrieval.

Simon Jenni received the BSc, MSc, and Ph.D. degrees from the University of Bern in 2015, 2017, and 2021. He is a Research Scientist at Adobe Research. He joined Adobe in May 2021. His research interests are in computer vision and deep learning, focusing on self-supervised learning methods for images and videos.
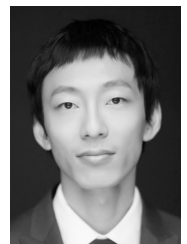
Zhifei Zhang is currently a Research Engineer with Adobe Research. His research interests are in representation learning and image editing/synthesis. He has worked extensively on internal frameworks for scalable visual search systems.

Viswanathan (Vishy) Swaminanthan is a senior principal scientist in Adobe Research focused on next generation video and digital experience technologies. His areas of research include video streaming, processing, coding, protection, digital rights management, and analytics. Vishy received his MS and Ph.D. in electrical engineering from Utah State University.