

OSCAR-Net: Object-centric Scene Graph Attention for Image Attribution

Eric Nguyen^{1*}, Tu Bui^{2*}, Viswanathan Swaminathan¹, John Collomosse^{1,2}

¹Adobe Research, Creative Intelligence Lab — San Jose, CA.

²CVSSP, University of Surrey, — Guildford, UK

{enguyen | collomos | vishy}@adobe.com, t.v.bui@surrey.ac.uk

Abstract


Images tell powerful stories but cannot always be trusted. Matching images back to trusted sources (attribution) enables users to make a more informed judgment of the images they encounter online. We propose a robust image hashing algorithm to perform such matching. Our hash is sensitive to manipulation of subtle, salient visual details that can substantially change the story told by an image. Yet the hash is invariant to benign transformations (changes in quality, codecs, sizes, shapes, etc.) experienced by images during online redistribution. Our key contribution is OSCAR-Net¹ (Object-centric Scene Graph Attention for Image Attribution Network); a robust image hashing model inspired by recent successes of Transformers in the visual domain. OSCAR-Net constructs a scene graph representation that attends to fine-grained changes of every object’s visual appearance and their spatial relationships. The network is trained via contrastive learning on a dataset of original and manipulated images yielding a state of the art image hash for content fingerprinting that scales to millions of images.

1. Introduction

Fake news and misinformation are centuries-old societal problems, exacerbated today by the ease with which digital images are manipulated and shared. Matching (or ‘attribution’) an image back to a trusted source, improves user awareness of its origins (or ‘provenance’) and so enables more informed trust decisions to be made [17, 45].

Emerging standards for image attribution embed provenance information within metadata [45, 3]. Yet image metadata is commonly stripped by social platforms, and may be replaced to misattribute an image [12]. One solution is to visually match images to a trusted database via ‘fingerprinting’; *i.e.* visual search or content-aware hashing.

This paper contributes a novel object-centric approach for computing a robust hash from an image, to perform visual matching for image attribution. Images often undergo ‘benign transformations’ during online distribution, such as



changes in format, resolution, size, padding, etc. that render cryptographic hashes[16] unsuitable as fingerprints. Thus image fingerprints must be made robust to these ‘benign’ transformations.

We propose a representation learning technique that encourages our hashes to exhibit both invariance to benign transformations and sensitivity to tampering (‘manipulation’) where image content is altered — sometimes quite subtly — but sufficiently to change meaning. For example, the editorial changes to introduce particular objects/motifs, or alterations to salient visual details such as the face, could substantially change the story told by an image (Fig. 1,2). Our method generates substantially different hashes in such cases to avoid corroborating a false story with the provenance information of the original image. Existing perceptual hashes [60, 43, 59, 39, 34]) typically exhibit the opposite desired characteristics (c.f. Sec. 4.5), that is, they are often invariant to minor changes, and sensitive to benign transformations (such as noise, or padding) during content redistribution.

Our core technical contribution is to encode a scene graph representation of the image via a hybrid network architecture, combining a fully-connected Graph Neural Network (GNN) and Transformer to create a robust binary hash of the image. Our *Object-centric Scene Graph Attention for Image Attribution Network* (OSCAR-Net) decomposes

*Equal contribution by co-authors

¹Code, data and model available at <https://exnx.github.io/oscar/>

a scene into salient objects and aggregates their description into a compact visual hash. This creates an object-centric hash that exhibits improved sensitivity to minor manipulation of salient objects, and improved robustness to benign transformations. We leverage contrastive training coupled with benign transformations using data augmentation and manipulations using Adobe PhotoshopTM to learn representations that outperform existing perceptual hashing baselines. Both object appearance and relative object geometry are learned via OSCAR-Net to yield an object-centric hash with improved tamper-sensitivity and scalability. We combine our learned OSCAR-Net representation with geometric verification to create a prototype image attribution tool to assist users in determining the provenance of images encountered online. Such a tool has future potential for fighting fake news by helping users gauge the trustworthiness of content.

2. Related Work

Visual content authenticity methods can be categorized into two camps. The first camp is within digital forensics, and includes techniques such as the ‘blind’ localization of image manipulation [56], detection of ‘deep fake’ (generative) content [14] based either upon its statistical properties [57] or its current limitations (*e.g.* absence of blinking [36]), and anomaly detection [58]. The second camp uses image attribution tools that trace provenance via embedded metadata [45, 3], watermarking [20, 13, 44, 4], or hashing [43, 31, 39, 8, 42, 61, 5]. Our work falls in the second camp – a robust visual hash for image attribution.

Perceptual hashing has been studied extensively for image similarity [55]. Classical approaches sample the spectral domain (*e.g.* DCT [60, 43] or wavelet [26] coefficients), or color-texture features [33, 1]. Compact binary codes may be obtained via random projection [25] or by learning data-dependent similarity-preserving binary projections such as product quantization (PQ) [27, 29]. More recently, convolutional neural networks (CNNs) have been explored for image similarity. Binary codes may also be learned concurrently with feature learning through two methods: continuous relaxation [63, 39, 35, 8] or direct optimization in discrete space [34, 37, 49, 62]. As we later show, these methods are vulnerable to subtle tampering on complex images thus not suitable for image attribution applications.

Scene graph representation is an active research area common in scene understanding [54], image captioning [18], retrieval [40] and synthesis [2]. Graph neural networks (GNN) are either spectral, which distills information via noise filtering in the graph Fourier domain [32, 10, 38], or non-spectral which operates on the graph directly [18, 40, 2, 41]. Non-spectral GNNs can work on graphs with arbitrary structure and various types of nodes and edges, thus can be adapted to complex design elements such as self-attention [53], autoencoder [40] and manifold learnings [41].

To the best of our knowledge, we are the first to leverage scene graph hashing for image attribution. Our contributions are three-fold: (i) we setup new benchmarks for safe-

guarding visual content integrity via image attribution, (ii) we propose a robust tamper-aware binary fingerprinting approach for images “in-the-wild” and (iii) we achieve state-of-art performance compared with existing hash methods.

3. Robust Image Fingerprinting

Our goal is to learn a model for tamper-aware image hashing that is invariant to benign image transformations, but sensitive to tampering due to subtle yet salient manipulation of content (sec. 3.1). Since content manipulation often occur at the object level, we explore an object-centric approach that decomposes images into a graph structure which explicitly models their objects and spatial relationships (Fig. 1). Sec. 3.2 describes how we construct a scene graph from an image, sec. 3.3 describes our scene graph encoder network, sec. 3.4 details our online hashing and losses, and sec. 3.5 outlines the training procedure.

3.1. Manipulation vs. Benign Transformation

We trained and evaluated on a dataset called ‘Photoshopbattles’ collected from Reddit forum discussions [23], which contain manipulated images using Adobe PhotoshopTM. This dataset contains **original images** each with several user-manipulated images that often change the story told by the image. The **manipulated images** include object removal/insertion, face editing/swap, geometry change among other type of manipulations. We also create **benign transformed images** from each original, applying a suite of primary and secondary transformations such as compression, padding, noise, resampling and affine transformation. More details of this dataset and transformations are discussed in sec. 3.5 and 4.1.

3.2. Scene Graph Representation

We represent an image x by a *fully-connected* attributed graph $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, where \mathcal{N} represents node features of the objects in x , and \mathcal{E} represents pairwise relationships between every object. We specifically used fully-connected graphs to model any potential tampering between all objects. To build \mathcal{G} , we first identify the objects in an image using MaskRCNN instance segmentation [21]. This results in a set of object crops $\{c_i | i = 1, 2, \dots, N\}$, where N = number of objects, its corresponding instance segmentation masks $\{m_i\}$, and bounding boxes $\{b_i = [x_i, y_i, h_i, w_i]\}$ characterized by the box center (x_i, y_i) and height/width (h_i, w_i) . We also treat the background (image minus the objects) as its own ‘object’. The output of the scene graph representation is N object and N^2 relation embeddings.

Node features. We integrate 3 visual cues to construct the node features \mathcal{N} . First, for each object, we extract visual features v_i from the object crop c_i . Specifically, we use $v_i = f^v(c_i) \in \mathbb{R}^{256}$, where f^v is a trainable ResNet50 [22] with the softmax layer being replaced by a 256-D fully connected layer (training details in sec. 3.5). Second, we compute shape features $s_i \in \mathbb{R}^7$ from the object mask m_i via 7 affine-invariant Hu moments [24]. Third, we use ge-




Figure 1. **Object-centric Scene Graph Attention Network for Image Attribution (OSCAR-Net)**. In **A**, our object-centric hashing method first decomposes an image into a fully connected scene graph of N detected objects by their appearance (c_i), masks (m_i) and bounding box geometry (b_i). The whole image, N object and N^2 relation embeddings are fed into the 3-stream attention-based encoder, comprised of a global CNN branch \mathcal{F}_G , an object encoder \mathcal{F}_O and a relation encoder \mathcal{F}_R to encode each stream to z_g, z_o, z_r , respectively. A single 64-D embedding z is created, and passed through a sign function to create a 64 bit image hash. In **B**, we indicate how tampered objects and relations are explicitly captured in the scene graph and allow for fine-grained manipulations to produce substantially different hash compared to the original hash.

ometry features $g_i \in \mathbb{R}^5$ directly from bounding boxes b_i :

$$g_i = \left[\frac{x_i}{w}, \frac{y_i}{h}, \frac{w_i}{w}, \frac{h_i}{h}, \frac{A_i}{A} \right], \quad (1)$$

where h, w are the image height and width; and $A = hw, A_i = x_i y_i$ are the areas of the image and the bounding box respectively.

We then aggregates v_i, s_i, g_i to create the node features:

$$n_i = [E_v(v_i), E_s(s_i), E_g(g_i)] \quad (2)$$

$$\mathcal{N} = \{n_i | i = 1, 2, \dots, N\} \in \mathbb{R}^{N \times D_N}, \quad (3)$$

where E_v, E_s, E_g are linear projections of the visual, shape and geometry features; $[\cdot]$ denotes a concatenation; D_N is the resulting feature dimension.

Edge features. We use pairwise geometry relations for edge features in the pixel coordinate space. The geometric connection between the i^{th} and j^{th} objects is defined as:

$$r_{ij} = \left[\frac{\Delta x}{\sqrt{A_i}}, \frac{\Delta y}{\sqrt{A_i}}, \frac{\sqrt{\Delta x^2 + \Delta y^2}}{\sqrt{w^2 + h^2}}, \frac{w_j}{w_i}, \frac{h_j}{h_i}, \theta_{ij}, \gamma_{ij} \right], \quad (4)$$

where transition $\Delta x = x_j - x_i, \Delta y = y_j - y_i$, box angle $\theta_{ij} = \arctan(\Delta y / \Delta x)$ and standard Intersection over Union (IoU) $\gamma_{ij} = \frac{m_i \cap m_j}{m_i \cup m_j}$. r_{ij} is then used to connect the two objects features:

$$e_{ij} = [n_i, E_r(r_{ij}), n_j] \quad (5)$$

$$\mathcal{E} = \{e_{ij} | i, j = 1, 2, \dots, N\} \in \mathbb{R}^{N^2 \times D_E}, \quad (6)$$

where E_r is another linear layer and D_E is the resulting dimension of the edge features.

3.3. Scene Graph Encoder

Existing GNN approaches typically define the graph structure based on relationships between objects, creating salient but sparse adjacency matrices. These matrices are acquired primarily either through manual annotations [10], heuristic assumptions about visual similarity [38] or geometry [30, 18], which may not necessarily hold true for “images in the wild” that have been manipulated. Instead, we are interested in *all* relationships between objects, since tampering can occur anywhere. Therefore, we encode graph \mathcal{G} via an attention mechanism that learns object and connection importance between all objects in a fully-connected manner.

The object encoder \mathcal{F}_O learns an embedding of all objects using node features \mathcal{N} as input. We use a Transformer encoder architecture [52], \mathcal{H}_O , to implicitly learn the importance of each object’s features in an order-invariant manner (we removed the positional encoding so that \mathcal{N} can be treated as an unordered sequence). Here, \mathcal{H}_O has 6 layers each having 8 attention heads. Since the Transformer encoder uses a constant latent vector size for all of its layers, its output sequence has the same length and feature dimension as the input, $H_O(\mathcal{N}) \in \mathbb{R}^{N \times D_N}$. To address this, we add a self-attention module, P_O , that acts as a node-wise pooling layer, aggregating the output sequence into a single embedding, following [50]:

$$w = \sigma(\tanh(\mathcal{N}' K^T) V); \quad P_O(\mathcal{N}') = \sum_i w_i \mathcal{N}'_i, \quad (7)$$

where \mathcal{N}' is the Transformer output, the Key $K \in \mathbb{R}^{N \times D_N}$ and Value $V \in \mathbb{R}^{D_N}$ behave the same way as the attention mechanism in Transformer but now are learnable parameters. Overall, the Transformer and self-attention layer form




Figure 2. Examples of original (left) and manipulated (right, highlighted) images in PSBattles24K. From the top-down: change in gaze, geometry and pose (see coloured box). More in Supp. Mat.

our object encoder and together output an object embedding $z_o = \mathcal{F}_O(\mathcal{N}) = P_O(\mathcal{H}_O(\mathcal{N})) \in \mathbb{R}^{D_N}$.

The relation encoder \mathcal{F}_R uses an identical architecture to \mathcal{F}_O but instead uses edge features as input. \mathcal{F}_R outputs a relation embedding $z_r = \mathcal{F}_R(\mathcal{E}) = P_E(\mathcal{H}_E(\mathcal{E})) \in \mathbb{R}^{D_E}$ where \mathcal{H}_E and P_E are analogous to the object Transformer \mathcal{H}_O and self-attention P_O .

Object and global image level fusion is achieved by fusing our ‘object-centric’ embeddings (z_o, z_r) with a visual appearance feature derived from the entire image, $z_g = \mathcal{F}_G(x)$. \mathcal{F}_G has the same CNN architecture as the object feature extractor f^v , but encodes the ‘global’ appearance instead of object crops. We do not share the weights of \mathcal{F}_G and f^v during training.

Our final embedding is achieved by concatenating the embeddings from the 3 above modules:

$$z = E_d([z_g, z_o, z_r]) \in \mathbb{R}^D, \quad (8)$$

with E_d being a linear projection layer for a desired output dimension D ($D = 64$ in our work). The late-fusion of 3 components - \mathcal{F}_G , \mathcal{F}_O and \mathcal{F}_R - allows our model to be sensitive to object level as well as global manipulations.

The closest work to our design is GAT [53] which also leverages attention mechanism for graph encoding. Our design differs in 3 key aspects: (i) our domain is images instead of text; (ii) we address a retrieval problem, and require a compact embedding; and (iii) OSCAR-Net uniquely uses 3 encoder streams to encode the global- and object-level features. OSCAR-Net assumes all objects and relations have connections with each other; the weights of these connections are learned by our Transformer.

3.4. Hashing and Losses

Hashing is essential for scalable visual search in an image attribution system. We obtain our hash by discretizing our continuous embedding space z with a $\text{sign}(\cdot)$ layer:

$$u = \text{sign}(z) \in \{-1, 1\}^D, \quad (9)$$

The quantization error is approximated via loss:

$$\mathcal{L}_B(z, u) = \|z - u\|^3, \quad (10)$$

where $\|\cdot\|$ denotes entrywise vector norm. Since eq. 9 has an ambiguous gradient at zero, we bridge the gradient of u to z during backprop, that is $\partial\mathcal{L}/\partial z = \partial\mathcal{L}/\partial u$, following Discrete Proximal Linearized Minimization (DPLM) [48, 49] (for more detailed analysis of eq. 10 see Supp. Mat.)

To distinguish manipulation from other benign transformations we propose a supervised version of NTXent-Loss/SimCLR [9] as the content loss. Under this loss, which we call SimCLR+, each image is pulled towards other benign transformed variants of itself while separating from manipulated variants and other images (different identity) in the batch:

$$\mathcal{L}_C(u_i) = -\log \frac{\sum_{i+} \varsigma(u_i, u_{i+})}{\sum_{i-} \varsigma(u_i, u_{i-}) + \sum_{j|\psi_j \neq \psi_i} \varsigma(u_i, u_j)} \quad (11)$$

where $\varsigma(u_i, u_j) = e^{d(E_b(u_i), E_b(u_j))/\tau}$; E_b is a linear layer, $d(\cdot)$ is cosine similarity, and τ is the temperature for contrastive loss. E_b serves two purposes: (i) mapping the binary code back to a continuous space before the loss computation and (ii) providing an intermediate learning buffer as recommended for SimCLR [9]. ψ_i denotes instance class (identity) of image x_i , and u_{i+} (*resp.* u_{i-}) indicates hash code of a benign transformation (*resp.* manipulated) image derived from x_i . Without u_{i-} the loss degenerates to the standard NTXentLoss and becomes self-supervised.

The total loss becomes $\mathcal{L}(\cdot) = \mathcal{L}_C(\cdot) + \alpha\mathcal{L}_B(\cdot)$ where α is the binary loss weight ($\alpha = 10^{-2}$). During inference, u is converted to $\{0, 1\}^D$ as $u := (u + 1)/2$.

3.5. Training Procedure

We first train a single CNN model using SimCLR+ loss (eq. 11), then use it to initialize the weights of both the global CNN (\mathcal{F}_G) and object-level CNN (f^v) modules. This pretraining step helps to learn a robust encoder for visual appearance, and improves the overall performance (sec. 4.4). In the main training phase, the whole network is trained end-to-end. We train our models using PyTorch with the ADAM optimizer and step decay learning rate initializing at 10^{-3} in pretraining and 10^{-4} in the subsequent phase. We use an 8GB GTX2080 GPU with a batch size of 20 unique images (prior to data augmentation), and ensure every original image has at least one manipulated version in a batch.

Benign transformations. To ensure our model is invariant to benign transformations, we apply a set of transformations on our training images for both the originals and manipulated variants. We define two categories of transforms: primary and secondary. The primary transforms (including random JPEG compression and resize) are always applied, while 1-3 secondary transforms (horizontal flip, rotation, sharpness, color enhancement, Gaussian noise, padding) are randomly chosen and applied after the primary transforms. This is to reflect that images are commonly compressed and

Method	Benchmark I		Benchmark II	
	mAP	mmAP	F_{mAP}	F_{R1}
OSCAR-Net	0.8898	0.7411	0.3782	0.3648
GNN	0.8807	0.7111	0.3643	0.3542
CNN	0.7980	0.6086	0.3516	0.3308
GreedyHash [49]	0.6635	0.3456	0.2367	0.2428
HashNet [8]	0.8093	0.4031	0.2047	0.2344
CSQ [59]	0.5785	0.2838	0.2404	0.2356
DFH [37]	0.3207	0.1595	0.2118	0.1842
DBDH [62]	0.6908	0.3339	0.2199	0.2268
DSDH [34]	0.6958	0.3280	0.2078	0.2186
ADSH [28]	0.3339	0.1887	0.1753	0.1289
DPSH [35]	0.8202	0.3917	0.1724	0.2082
DSH [39]	0.2416	0.1358	0.1556	0.1137
DHN [63]	0.1803	0.0898	0.1407	0.1108
wHash [6]	0.5338	0.2274	0.0922	0.1059
aHash [6]	0.5764	0.2668	0.0919	0.1077
pHash [6]	0.6008	0.3260	0.0787	0.0972
ISCC [43]	0.6003	0.3252	0.0787	0.0967
dHash [6]	0.6164	0.2890	0.0604	0.0703
cHash [6]	0.2509	0.1018	0.1896	0.1674

Table 1. Evaluation of proposed method OSCAR-Net versus baselines; all use 64 bit hash codes.

	Method	Bm.I mmAP	Bm.II F_{R1}
Content loss	SimCLR+	0.7411	0.3648
	Triplet+	0.7384	0.3627
	SimCLR [9]	0.7013	0.3500
	Triplet [46]	0.7190	0.3295
	Pairwise Contrastive [19]	0.2199	0.1668
Hash loss	No hashing	0.8722	0.3984
	sign() (proposed)	0.7411	0.3648
	sign() + Bit Balance	0.7217	0.3582
	tanh() + Bit Balance [62]	0.7076	0.3556
	tanh() [8]	0.7096	0.3550
	PQ [27] (offline)	0.6552	0.3572

Table 2. Evaluation of different content (top half) and hashing (bottom half) losses. For content loss experiments, we fix the proposed hash loss \mathcal{L}_B (eq. 9) and vary \mathcal{L}_C . Likewise, we fix the proposed content loss \mathcal{L}_C (eq. 11) and vary \mathcal{L}_B in the hash loss experiments (except for PQ being an offline hashing method).

resized when redistributed online. For details of the parameter ranges for these transformations, see the Supp. Mat.

4. Experiments and Discussion

We evaluate the performance of OSCAR-Net against 6 classical and 10 deep learning baselines for image hashing.

4.1. Datasets

PSBattles24K. We train and evaluate on a set of user-generated manipulated images posted released as PSBattles dataset [23]. The raw PSBattles dataset has 11K original and 90K manipulated variants of those images. To increase the challenge of this dataset, paired images were sorted by ImageNet distance (ResNet50/ImageNet) between the original and manipulated variant, and the lower quartile (most subtle changes) was retained. After duplicate removal, 7K originals and 24K manipulated variants remain (Fig. 2). The

dataset is split into distinct train and test sets, consisting 21K and 3K pairs of original-manipulated images respectively. 5% of the training set is for validation. Further, we create a test set of 150K images (50 random primary and secondary transformations per original) to evaluate invariance to benign transformations (c.f. sec. 3.5).

PSBattles360K-S We study the robustness of our models toward *individual* augmentations using two additional test sets of 10 randomized augmentations applied to each original image within PSBattles24K. The first set contains images from 6 benign transformations seen during the training: compression, rotation, color enhancement, Gaussian noise, padding and sharpness. The second set contain images made from 6 noise sources common in photography and unseen during the training: shot noise, impulse noise, speckle noise, Gaussian blur, de-focus blur and pixelate. Each set has in total 180K ($3K \times 10 \times 6$) images. More details of these transformations can be found in Supp. Mat. **Stock4.7M** is a diverse, unannotated dataset of images used to evaluate retrieval performance in the presence of large-scale distractor images. The dataset comprises 4.76M images from Adobe Stock website ² at VGA resolution. Combined with PSBattles24K this enables us to scale evaluation to a test dataset of $\sim 5M$ images.

4.2. Benchmarks and Evaluation Metrics

To evaluate model performance specifically under the image attribution setting, we defined 2 benchmarks (both are correlated, but differ in terms of the amount of benign vs. manipulated images in the query and database sets).

Benchmark I is used to evaluate embedding separability between manipulated and benign images. It considers a query set of 3K original images and a search database comprised of (i) 3K manipulated images (ii) 150K benign transformed variants of images from PSBattles24K (iii) 100K distractor images from Stock4.7M. Given an original image as the query, we wish to retrieve all benign augmentations and reduce the ranking of any manipulated imagery in the returned results. To this end, we propose a masked Mean Average Precision (mmAP) metric for benchmark I. For a single query q ;

$$mmAP_q = \frac{\sum_k r_q(k) m_q(k) P_q(k)}{\sum_k r_q(k)}, \quad (12)$$

where $P_q(k)$ is precision at k , $r_q(k) = 1$ if the k^{th} retrieval is relevant otherwise 0, $m_q(k) = 1$ if the manipulated image is ranked below k otherwise 0. mmAP thus penalizes early ranking of manipulated images, by computing standard mAP *up to* but excluding the first manipulated result. mmAP achieves an upper bound of 1 if all benign images are returned on top of the rankings. In scalability experiments where all Stock4.7M distractors are included in the search database, computing full mmAP is not feasible. We therefore use mmAP@ R instead; $mmAP_q$ is computed on top R retrieval results only, where R is number of images relevant to query q in the database.

²<https://stock.adobe.com>. Image list at <https://exnx.github.io/oscar>.

Components/Exp. ID	A	B	C	D	E	F (CNN)	G	H (GNN)	I	J (full)
Fully connected graph	✓	✓	✓	✓	✓		✓	✓		✓
Transformer self-attention	✓	✓	✓	✓	✓		✓		✓	✓
Visual appearance features v_i	✓	✓	✓	✓	✓		✓	✓	✓	✓
Vis. appear. feat. w/ pretrain f^v		✓	✓	✓	✓		✓	✓	✓	✓
Geometry feature g_i			✓	✓	✓		✓	✓	✓	✓
Shape feature s_i				✓	✓		✓	✓	✓	✓
Edge feature e_{ij}					✓		✓	✓	✓	✓
Global CNN \mathcal{F}_G						✓	✓	✓	✓	✓
Benchmark I - mmAP	0.2715	0.6282	0.7028	0.7015	0.7148	0.6086	0.6597	0.7111	0.7322	0.7411
Benchmark II - F_{R1}	0.1622	0.3221	0.3547	0.3594	0.3621	0.3308	0.3532	0.3542	0.3599	0.3648

Table 3. Ablation of the proposed (OSCAR-Net) method omitting various architectural stages (see subsec 4.4 for details on Exps. A-J).

Benchmark II is aligned closely with our image attribution use case, where the test dataset contains the original images (*e.g.* a database of images with associated provenance information) plus Stock4.7M distractors, and the query set comprises benign and manipulated variants of those originals. An ideal model would rank an original highly when querying with a benign transformed version, and rank it low for manipulated queries. Recall at top-1 $R@1$ is computed for benign queries; and $\overline{R@1} = 1 - R@1$ for manipulated queries. We propose F_{R1} score to measure the trade-off:

$$F_{R1} = R@1 \times \overline{R@1} / (R@1 + \overline{R@1}), \quad (13)$$

We report F_{mAP} in the same manner, computed over mAP.

4.3. Baseline Comparison

For comparison, we evaluated our OSCAR-Net against the following baselines. We consider OSCAR-Net and GNN as object-centric, and other baseline methods as “global” *i.e.* encoding the whole image vs decomposing into objects.

1. GNN: We replace the Transformer modules $\mathcal{H}_{O,R}$ of OSCAR-Net with MLP layers to learn graph embeddings while keeping everything else the same. This degenerates to a similar approach to the graph CNNs in [18, 40] although we differ in data domain, losses, and training method.

2. CNN: In contrast with our object-centric approach, this uses a single CNN (ResNet50) to encode the whole image. The model is trained using the same losses and data augmentations detailed in sec. 3.4-3.5.

3. Classical Methods: A set of 6 statistical methods for perceptual hashing. Five methods (D-ifference Hash, P-erception Hash, A-verage Hash, W-avelet Hash and C-olor Hash) via public implementations in [6], also ISCC [43] an ISO standard proposal similar to pHash. All methods produce 64-bit hash codes.

4. Deep Hashing Methods: A set of 10 deep supervised hashing approaches: CSQ [59], DBDH [62], DFH [37], HashNet [8], GreedyHash [49], ADSh [28], DPSH [35], DSH [39], DSDH [34] and DHN [63]. We train these models using provided public code using uniform architecture (ResNet50) and code length (64-bit).

Tab. 1 reports the performance of all methods on the PS-Battles24K test set with 100K distractor images for benchmark I and II (sec. 4.2). Overall the classical methods are

at the bottom of the rankings, followed by the global deep hashing methods. The object-centric methods OSCAR-Net and GNN outperform the rest, especially at benchmark I mmAP scores (which takes manipulated images into account as opposed to standard mAP). This indicates the methods’ capability in discriminating manipulated content. The object-centric methods also outperform others on benchmark II which requires the balance between benign-transform robustness and manipulation sensitivity. Our proposed OSCAR-Net outperforms all existing methods and sets a new state-of-the-art in image hash for content fingerprinting.

4.4. Ablation Studies

Design Component Ablation. We study the contribution of each design component in OSCAR-Net in Tab. 3. We begin with a minimal design where \mathcal{N} contains only the object features v_i and OSCAR-Net has only 1 encoder stream \mathcal{F}_O (exp.A). We then add additional components including the pretrained f^v (exp.B), geometry features g_i (exp.C), shape feature s_i (exp.D), the second encoder stream \mathcal{F}_R with edge features e_{ij} (exp.E) to the full model (exp.J). We also remove important components from the full model or test alternative designs (exp.G-I). Pretraining the CNN module f^v is essential as it helps double performance (exp.B vs. A). The geometry features g_i play an important role (exp.C) while the shape s_i and edge e_{ij} features steadily improve performance (exp.D-E). Integrating the global CNN branch \mathcal{F}_G to our object-centric network also boosts performance especially on benchmark I (exp.J), although the single CNN alone does not perform well (exp.F). Exp.G and H show the complementary effects of the transformer $\mathcal{H}_{O,R}$ and self-attention $P_{O,R}$ (eq. 7) modules in our network design. We implemented exp.G by removing the self-attention layer and used the feature of the first element in the output sequence as the embedding (similar to [15]), resulting in a performance drop by 8.1% and 1.2% on the two benchmarks. Exp.H is the GNN architecture described earlier in sec. 4.3, causing 3% and 1.1% drop in performance respectively. Additionally, we experimented with a non-fully connected graph in exp.I where two objects are considered as connected only if their relative distance and intersection are within a given threshold, similar to [30, 18]. Non-connected objects are masked out during




Figure 3. Characterizing our proposed OSCAR-Net content hashing performance vs. baseline methods. (a-b): robustness of the image hash to seen/unseen benign transformations; (c): performance of different hash code lengths for online (sign) and offline (PQ) hashing.

the attention computation in \mathcal{H}_O and their pairwise relations are removed from the edge features \mathcal{E} . Exp.I reports 0.9% performance drop in benchmark I and 0.5% on benchmark II. This demonstrates our OSCAR-Net capability in learning the connection weights within a fully connected graph.

Content Loss Ablation. Tab. 2 (top) shows how alternative content loss functions affect performance on the two benchmarks. Standard pairwise contrastive loss [19] fails to discriminate subtle differences in manipulated images. Standard SimCLR [9] lags behind by 4% and 1.5% on the two benchmarks but still outperforms standard triplet loss [46] despite being a self-supervised method. We modified triplet loss by: (i) randomly sampling negative images from other image instances instead of just manipulated images, (ii) randomly feeding the anchor branch with manipulated images and the negative branch with originals, and (iii) adding another L2 term to explicitly push the negatives from the originals. This improved version, named Triplet+, has the closest performance to SimCLR+ but requires a more complex sampling strategy.

Hash Loss Ablation. Tab. 2 (bottom) reports on alternative hash functions. We experimented with $\tanh()$ as a continuous relaxation method [8]. We also implemented the bit balancing strategy that regularize the number of 0 and 1 bits, as suggested by [62] to improve the hash discrimination. Performance on the 256-D continuous space *i.e.* no hash is reported as an upper-bound reference. Additionally, we compare online hashing with an offline method using Product Quantization (PQ) [27] via FAISS [29]. Offline hash with PQ had the lowest performance (8.6% and 0.8% behind our proposed method on the two benchmarks). For online hashing, optimizing directly on the discrete space ($\text{sign}()$) outperforms the relaxation method ($\tanh()$). Also, we find that bit balancing offers only minor improvement for the online approximation method and hinders unnecessary constraints on the proposed direct optimization method. Overall the use of hashing for image matching enables retrieval speeds of 380ms on average (an indicative speed for 1M images on single-core, i7 CPU) for a 64-bit hash, versus an average of 1.5s for the same using a 256-D real-valued embedding (128x larger footprint). The corresponding difference in accuracy (for our proposed, best case: discrete space optimization) is 13.1% on bench-

Method	1M	2M	3M	4M
OSCAR-Net	0.6680	0.6346	0.6112	0.5930
GNN	0.6295	0.5949	0.5711	0.5522
CNN	0.5226	0.4836	0.4619	0.4444
HashNet [8]	0.4064	0.3937	0.3854	0.3785
pHash [6]	0.3689	0.3689	0.3682	0.3673

Table 4. Scalability of best performers under Benchmark I: mmAP@R with increasing distractors (Stock4.7M). OSCAR-Net is our proposed method. GNN and CNN are ablations of it. HashNet and pHash are the benchmark I highest performing deep and classical baseline methods from Tab.1.

Method	1M	2M	3M	4M
OSCAR-Net	0.3247	0.3068	0.3003	0.2952
GNN	0.3118	0.2944	0.2822	0.2770
CNN	0.3003	0.2925	0.2880	0.2854
GreedyHash [49]	0.2284	0.2204	0.2153	0.2105
cHash [6]	0.1483	0.1464	0.1457	0.1447

Table 5. Scalability of best performers under Benchmark II: F_{R1} with increasing distractors (Stock4.7M). GreedyHash and cHash are the benchmark II highest performing deep and classical baseline methods from Tab.1.

mark I and 3.4% on benchmark II justifying the approach.

4.5. Scalability and Robustness

Scalability. Tab. 4-5 show the performance change when up to Stock4.7M images are added to the search database as distractors. We report mmAP@R for benchmark I and F_{R1} for benchmark II, comparing OSCAR-Net, GNN and CNN with the top classic and deep hash performers according to Tab. 1. Benchmark I shows OSCAR-Net and GNN (ablation) outperforming baselines by a large margin. For Benchmark II, GNN underperforms CNN after 2M distractors, potentially due to lack of attention mechanism. This is addressed by OSCAR-Net (via the Transformer module).

Transformation robustness. We study how the types of benign transformations affect the model performance. We set up an experiment in benchmark II style where the benign query set is from PSBattles360K-S. Fig. 3 (a-b) shows F_{R1} performance against seen and unseen transformations. Scores for a subset of transformations used in PSBattles24K are also reported as reference (labeled as “Mixed”). cHash performs reasonably on blurring/sharpness transformations



Figure 4. Representative examples (4 rows) of nearest neighbors in the learned embedding. For this visualization we mix the original, benign and manipulated images. Originals and benign variants of these are co-proximate, whereas manipulated content (red boxes highlight manipulations) is distant from the original.

but is vulnerable to additive noise or color enhancement. The global deep learning methods GreedyHash [49] and CNN perform better than cHash but consistently underperform the object-centric ones. OSCAR-Net beats other methods at most transformation experiments and generalizes well for unseen transformations.

Code length. Fig. 3(c) shows the performance of OSCAR-Net on different code length settings and hashing methods. For offline hashing, we train a OSCAR-Net model outputting 256-D continuous features using the content loss \mathcal{L}_C only, then apply PQ [27] for a desirable code length. For online hashing, we train separate models varying dimension of z and u in eq. 9. Fig. 3(c) confirms the benefits of online hashing versus offline, since the hash code could be optimized at the same time as its representation is being learned. Also, the optimal code length is shown around 128-bit; we use 64-bits in our experiments for fair comparison with baselines.

4.6. Similarity Visualization

We visually explain the image regions that contribute toward (dis-)similarity, by adapting GradCAM [47] to have objective function:

$$\mathcal{L}(x|x_+, x_-) = d(f(x), f(x_+)) - d(f(x), f(x_-)), \quad (14)$$

where $d(\cdot)$ is cosine similarity. Note that $\mathcal{L}(x|x_+, x_-)$ is similar to a triplet function but with opposite effect. Fig. 1 visualizes a heatmap on several original-manipulated image pairs. Regions of interest that contribute most to distance in our embedding are highlighted, regardless of benign transformations (which are learned to be ignored). Fig. 4 includes further retrieval examples. Manipulations are well separated from originals, but benign images are nearby.

5. Image Attribution using OSCAR-Net

We incorporate OSCAR-Net into a prototype image attribution system where image hashes are used to match images to a database in order to lookup provenance information such as authorship, secure timestamp etc. [45]. Here,

Method	FPR (%) ↓		FNR (%) ↓	
	0.1M	4M	0.1M	4M
OSCAR-Net	3.91	4.06	< 1	1.35
CNN	5.41	5.71	1.01	2.03
GreedyHash [49]	19.36	37.43	75.86	79.51
cHash [6]	52.54	53.43	81.24	87.59

Table 6. Image attribution: False positive/negative rate (FPR/FNR) for proposed (OSCAR-Net) method and baselines (lower is better). Match via L2 distance in search embedding filtered via geometric verification on a shortlist of the top 10 results.

images should be matched irrespective of any benign transformations applied to the image during online re-sharing. However, images that have been manipulated should not be matched to their originals in the database to avoid corroborating provenance to a false story. We conduct retrieval using 10% of the PSBattles24K query test to identify the top-10 ranked results for a given query, followed by a geometric verification (GV) step. GV is a common second stage visual search technique used to identify whether any of the top ranked results are relevant or not, based on registration of sparse feature points between the query and result. We fit a fundamental matrix to homogeneous SIFT point coordinates under MLESAC [51]. We use GV to avoid the need to specify threshold on distance within the hash embedding; instead relying on pass/fail of registration.

Tab. 6 reports the false positive rate (FPR) - the fraction of manipulated queries returning originals as a hit; and the false negative rate (FNR) - the fraction of benign queries failing to return originals as a hit. Our OSCAR-Net model exhibits lower FPR/FNR than all deep and classic baseline methods from Tab. 1.

6. Conclusion

We presented OSCAR-Net, an object-centric image hashing method that leverages a paired dataset of original and manipulated images to learn a search embedding for robust visual matching. Our hash is invariant to benign transformations typically applied to images distributed online, yet sensitive to fine-grained visual manipulations. We proposed a hybrid architecture that combined a GNN with Transformers and self-attention. We showed how our object-centric method significantly improved image search performance and scalability over global (full) image hashing using millions of distractor images. Future work will integrate our learned embeddings within content provenance systems [45, 3] to trace images circulating in the wild. For example, combining our hashes with decentralized immutable storage *e.g.* blockchain and IPFS, in the spirit of recent explorations for digital preservation [11, 7] but applied to new verticals *e.g.* journalism to combat fake news.

Acknowledgement

We thank Andy Parsons, Leonard Rosenthol, Bill Marino and the Adobe Content Authenticity Initiative (CAI) for discussions. This work was supported by an Adobe Research internship and latterly by DECaDE via EP-SRC Grant Ref EP/T022485/1.

References

- [1] B. Amrutkar and L. Singh. Efficient content based image retrieval using combination of dominant-color shape and texture features and k-means clustering. *IJECS*, 4(12), 2015. [2](#)
- [2] O. Ashual and L. Wolf. Specifying object attributes and relations in interactive scene generation. In *Proc. ICCV*, pages 4561–4569, 2019. [2](#)
- [3] J. Aythora, R. Burke-Agüero, A. Chamayou, S. Clebsch, M. Costa, J. Deutscher, N. Earnshaw, L. Ellis, P. England, C. Fournet, et al. Multi-stakeholder media provenance management to counter synthetic media risks in news publishing. In *Proc. Intl. Broadcasting Convention (IBC)*, 2020. [1](#), [2](#), [8](#)
- [4] S. Baba, L. Krekor, T. Arif, and Z. Shaaban. Watermarking scheme for copyright protection of digital images. *IJCSNS*, 9(4), 2019. [2](#)
- [5] A. Bharati, D. Moreira, P.J. Flynn, A. de Rezende Rocha, K.W. Bowyer, and W.J. Scheirer. Transformation-aware embeddings for image provenance. *IEEE Trans. Info. Forensics and Sec.*, 16:2493–2507, 2021. [2](#)
- [6] J. Buchner. Imagehash. <https://pypi.org/project/ImageHash/>, 2020. [5](#), [6](#), [7](#), [8](#)
- [7] T. Bui, D. Cooper, J. Collomosse, M. Bell, A. Green, J. Sheridan, J. Higgins, A. Das, J. Keller, and O. Thereaux. Tamper-proofing video with hierarchical attention auto-encoder hashing on blockchain. *IEEE Trans. Multimedia (TMM)*, 22(11):2858–2872, 2020. [8](#)
- [8] Z. Cao, M. Long, J. Wang, and P. S. Yu. Hashnet: Deep learning to hash by continuation. In *Proc. CVPR*, pages 5608–5617, 2017. [2](#), [5](#), [6](#), [7](#)
- [9] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *Proc. ICML*, pages 1597–1607. PMLR, 2020. [4](#), [5](#), [7](#)
- [10] Z. Chen, X. Wei, P. Wang, and Y. Guo. Multi-label image recognition with graph convolutional networks. In *Proc. CVPR*, pages 5177–5186, 2019. [2](#), [3](#)
- [11] 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. [8](#)
- [12] IPTC Council. Social media sites photo metadata test results. <http://embeddedmetadata.org/social-media-test-results.php>, 2020. [1](#)
- [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*, 5(11):831–837, 2019. [2](#)
- [14] B. Dolhansky, J. Bitton, B. Pflaum, J. Lu, R. Howes, M. Wang, and C. C. Ferrer. The deepfake detection challenge (DFDC) dataset. *CoRR*, abs/2006.07397, 2020. [2](#)
- [15] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2020. [6](#)
- [16] H. Gilbert and H. Handschuh. Security analysis of sha-256 and sisters. In *Int. Ws. Selected Areas in Cryptography*, pages 175–193. Springer, 2003. [1](#)
- [17] S. Gregory. Ticks or it didn't happen. Technical report, Witness.org, 2019. [1](#)
- [18] L. Guo, J. Liu, J. Tang, J. Li, W. Luo, and H. Lu. Aligning linguistic words and visual semantic units for image captioning. In *Proc. ACM ICM*, pages 765–773, 2019. [2](#), [3](#), [6](#)
- [19] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *Proc. CVPR*, pages 1735–1742, 2006. [5](#), [7](#)
- [20] K. Hameed, A. Mumtax, and S. Gilani. Digital image watermarking in the wavelet transform domain. *WASET*, 13:86–89, 2006. [2](#)
- [21] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proc. ICCV*, pages 2961–2969, 2017. [2](#)
- [22] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. [2](#)
- [23] S. Heller, L. Rossetto, and H. Schuldt. The PS-Battles Dataset – an Image Collection for Image Manipulation Detection. *CoRR*, abs/1804.04866, 2018. [2](#), [5](#)
- [24] M.-K. Hu. Visual pattern recognition by moment invariants. *IRE Trans. Information Theory*, 8(2):179–187, 1962. [2](#)
- [25] P. Indyk, R. Motwani, P. Raghavan, and S. Vempala. Locality preserving hashing in multi-dimensional spaces. In *Proc. ACM STOC*, pages 618–625, 1997. [2](#)
- [26] C. Jacobs, A. Finkelstein, and D. Salesin. Fast multiresolution image querying. In *Proc. ACM SIGGRAPH*, pages 277–286. ACM, 1995. [2](#)
- [27] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE TPAMI*, 33(1):117–128, 2010. [2](#), [5](#), [7](#), [8](#)
- [28] Q.-Y. Jiang and W.-J. Li. Asymmetric deep supervised hashing. In *AAAI*, 2018. [5](#), [6](#)
- [29] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with gpus. *IEEE Trans. on Big Data*, 2019. [2](#), [7](#)
- [30] N. Khan, U. Chaudhuri, B. Banerjee, and S. Chaudhuri. Graph convolutional network for multi-label vhr remote sensing scene recognition. *Neurocomputing*, 357:36–46, 2019. [3](#), [6](#)
- [31] F. Khelifi and A. Bouridane. Perceptual video hashing for content identification and authentication. *IEEE TCSVT*, 1(29), 2017. [2](#)
- [32] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *ICLR*, 2017. [2](#)
- [33] I. H. Laradji, L. Ghouti, and E. Khiari. Perceptual hashing of color images using hypercomplex representations. In *Proc. ICIP*, pages 4402–4406, 2013. [2](#)
- [34] Q. Li, Z. Sun, R. He, and T. Tan. Deep supervised discrete hashing. In *Proc. NeurIPS*, pages 2482–2491, 2017. [1](#), [2](#), [5](#), [6](#)
- [35] W. Li, S. Wang, and W.-C. Kang. Feature learning based deep supervised hashing with pairwise labels. In *Proc. IJCAI*, pages 1711–1717, 2016. [2](#), [5](#), [6](#)
- [36] Y. Li, M.-C. Ching, and S. Lyu. In icu oculi: Exposing ai created fake videos by detecting eye blinking. In *Proc. IEEE WIFS*, 2018. [2](#)
- [37] Y. Li, W. Pei, and J. van Gemert. Push for quantization: Deep fisher hashing. *BMVC*, 2019. [2](#), [5](#), [6](#)
- [38] F. Liu, C. Zou, X. Deng, R. Zuo, Y. Lai, C. Ma, Y. Liu, and H. Wang. Scenesketcher: Fine-grained image retrieval with scene sketches. In *Proc. ECCV*, pages 718–734. Springer, 2020. [2](#), [3](#)
- [39] H. Liu, R. Wang, S. Shan, and X. Chen. Deep supervised hashing for fast image retrieval. In *Proc. CVPR*, pages 2064–2072, 2016. [1](#), [2](#), [5](#), [6](#)
- [40] D. Manandhar, D. Ruta, and J. Collomosse. Learning structural similarity of user interface layouts using graph networks. In *Proc. ECCV*, pages 730–746. Springer, 2020. [2](#), [6](#)

- [41] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proc. CVPR*, pages 5115–5124, 2017. 2
- [42] D. Moreira, A. Bharati, J. Brogan, A. Pinto, M. Parowski, K.W. Bowyer, P.J. Flynn, A. Rocha, and W.J. Scheirer. Image provenance analysis at scale. *IEEE Trans. Image Proc.*, 27(12):6109–6122, 2018. 2
- [43] T. Pan. Digital-content-based identification: Similarity hashing for content identification in decentralized environments. In *Proc. Blockchain for Science*, 2019. 1, 2, 5, 6
- [44] D. Proffrock, M. Schlauweg, and E. Muller. Content-based watermarking by geometric wrapping and feature-based image segmentation. In *Proc. SITIS*, pages 572–581, 2006. 2
- [45] L. Rosenthal, A. Parsons, E. Scouten, J. Aythora, B. McCormack, P. England, M. Levallee, J. Dotan, et al. Content authenticity initiative (CAI): Setting the standard for content attribution. Technical report, Adobe Inc., 2020. 1, 2, 8
- [46] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proc. CVPR*, pages 815–823, 2015. 5, 7
- [47] R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. CVPR*, pages 618–626, 2017. 8
- [48] F. Shen, X. Zhou, Y. Yang, J. Song, H. T. Shen, and D. Tao. A fast optimization method for general binary code learning. *IEEE TIP*, 25(12):5610–5621, 2016. 4
- [49] S. Su, C. Zhang, K. Han, and Y. Tian. Greedy hash: Towards fast optimization for accurate hash coding in CNN. In *Proc. NeurIPS*, pages 798–807, 2018. 2, 4, 5, 6, 7, 8
- [50] S. Sukhbaatar, J. Weston, and R. et al. Fergus. End-to-end memory networks. In *Proc. NeurIPS*, pages 2440–2448, 2015. 3
- [51] P. Torr and A. Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *CVIU*, 78(1):138–156, 2000. 8
- [52] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Proc. NeurIPS*, pages 5998–6008, 2017. 3
- [53] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *ICLR*, 2018. 2, 4
- [54] P. Vicol, M. Tapaswi, L. Castrejon, and S. Fidler. Moviegraphs: Towards understanding human-centric situations from videos. In *Proc. CVPR*, 2018. 2
- [55] J. Wang, H. T. Shen, J. Song, and J. Ji. Hashing for similarity search: A survey. *arXiv preprint arXiv:1408.2927*, 2004. 2
- [56] S-Y. Wang, O. Wang, A. Owens, R. Zhang, and A. Efros. Detecting photoshopped faces by scripting photoshop. In *Proc. ICCV*, 2019. 2
- [57] 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. 2
- [58] 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*, pages 9543–9552, 2019. 2
- [59] 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*, pages 3083–3092, 2020. 1, 5, 6
- [60] C. Zauner. Implementation and benchmarking of perceptual image hash functions. Master’s thesis, Upper Austria University of Applied Sciences, Hagenberg, 2010. 1, 2
- [61] X. Zhang, Z. H. Sun, S. Karaman, and S.F. Chang. Discovering image manipulation history by pairwise relation and forensics tools. *IEEE J. Selected Topics in Signal Processing.*, 14(5):1012–1023, 2020. 2
- [62] X. Zheng, Y. Zhang, and X. Lu. Deep balanced discrete hashing for image retrieval. *Neurocomputing*, 2020. 2, 5, 6, 7
- [63] H. Zhu, M. Long, J. Wang, and Y. Cao. Deep hashing network for efficient similarity retrieval. In *Proc. AAAI*, 2016. 2, 5, 6