ON ROTATION INVARIANCE IN COPY-MOVE FORGERY DETECTION

Vincent Christlein, Christian Riess and Elli Angelopoulou

Pattern Recognition Lab
University of Erlangen-Nuremberg
sivichri@stud.informatik.uni-erlangen.de, {riess, elli}@i5.informatik.uni-erlangen.de

ABSTRACT

The goal of copy-move forgery detection is to find duplicated regions within the same image. Copy-move detection algorithms operate roughly as follows: extract blockwise feature vectors, find similar feature vectors, and select feature pairs that share highly similar shift vectors. This selection plays an important role in the suppression of false matches. However, when the copied region is additionally rotated or scaled, shift vectors are no longer the most appropriate selection technique.

In this paper, we present a rotation-invariant selection method, which we call *Same Affine Transformation Selection (SATS)*. It shares the benefits of the shift vectors at an only slightly increased computational cost. As a byproduct, the proposed method explicitly recovers the parameters of the affine transformation applied to the copied region. We evaluate our approach on three recently proposed feature sets. Our experiments on ground truth data show that SATS outperforms shift vectors when the copied region is rotated, independent of the size of the image.

Index Terms— Blind image forensics, copy-move forgery detection, rotation invariance, shift vectors

1. INTRODUCTION

The goal of blind image forensics is to assess the authenticity and origin of images from an unknown and uncontrolled source. For an overview see e.g. [1, 2, 3]. In general, there are three main approaches which involve the examination of either a) the presence of expected artifacts of the imaging process, b) the statistical properties of the output image, or c) the consistency of scene properties.

Representative examples of exploited imaging-process artifacts are camera identification from sensor noise [4], the analysis of lateral chromatic aberration [5, 6] or the Bayer pattern [7]. There is also a considerable body of work on the analysis of statistical properties of images as evidence of image tampering. Typical instances include methods based on JPEG artifacts [8, 9, 10] or traces of resampling [11, 12]. In comparison, the methods that analyze scene-properties investigate higher-level features like lighting direction [13], or the color of the illuminants [14].

There are also techniques that have been developed for the investigation of specific types of image tampering, like the detection of copy-move forgeries [15, 16, 17, 18, 19, 20, 21]. In this very popular field, the underlying assumption is that an image region has been copied and pasted within the same image. Common applications are either the concealment of information (e.g. to hide persons), or the emphasis of particular image content (for instance, to enlarge a crowd of people in a demonstration).

Most copy-move forgery detection (CMFD) algorithms can be cast in a common CMFD pipeline (see Sec. 2). There exist very robust methods for the detection of plain copy-move forgeries. However, the detection of rotated or scaled copy-move forgery (RS-CMFD) is still considered a challenging problem [22]. One approach to RS-CMFD is to use keypoint-based features like SIFT or SURF [23, 24]. One drawback of these methods is that they require a sufficiently distinctive image structure, such that a dense set of keypoints can be extracted [21].

Rotation or scale invariant features that can be more easily integrated in the CMFD pipeline have been proposed. However, they have their limitations, too. For example, the method by Bayram *et al.* [20] involves a large search space for rotation-invariant matching, which often renders the technique impractical. Bravo-Solorio *et al.* [16] propose rotation-invariant features, but lack a rotation-invariant post-processing. Therefore, their method does not reach its full potential. Finally, the techniques by Ryu *et al.* [21] and Wang *et al.* [25] work well in many cases, but lack a global interpretation of the results, as well as a principled way to rank and further evaluate groups of coherent pairs of blocks.

We believe that these shortcomings in RS-CMFD algorithms are mainly due to the lack of a sophisticated block selection in the verification stage (since shift vectors are not applicable under general affine transformations). In this paper, we offer a straightforward rotation- and translation-invariant replacement for the popular shift vectors. We call this processing step *Same Affine Transformation Selection (SATS)*. We show that any set of rotation-invariant features benefits from the inclusion of this processing step in the pipeline. More specifically, the contributions of this paper are:

- An efficient post-processing method, *SATS*. By construction, it can detect arbitrary variations in rotation and scaling in the copied part. The SATS-detected sets of copy-moved blocks form meaningful connected groups that can be further analyzed.
- A processing scheme where any suitable feature can be used in combination with the SATS post-processing algorithm.
- The evaluation of twelve previously proposed feature sets for their applicability in the detection of RS-CMFD.

An overview about the common CMFD pipeline is presented in Sect. 2. In Sect. 3, we examine different feature sets for their suitability for rotational copy-move forgery detection. In Sect. 4, we present the proposed rotation and scale invariant post-processing method. Experiments on ground truth data are finally presented in Sect. 5. The Code and associated data are available at `http://www5.informatik.uni-erlangen.de/software` and `http://www5.informatik.uni-erlangen.de/data`, respectively.

2. OVERVIEW ON CMFD METHODS

Most CMFD algorithms adhere to a common pipeline, as shown in Fig. 1. First, the image is optionally preprocessed (for instance converted to grayscale). It is then subdivided in overlapping blocks of pixels. On each of these blocks, a feature vector is extracted. Highly similar feature vectors are matched as pairs. Known methods for matching are lexicographic ordering on the feature vectors and nearest neighbor determination in a kd-tree. The similarity of two features can be determined by different similarity criteria, e. g. Euclidian distance.

In the verification step, outliers are removed by filtering the pairs of feature vectors ("verification" in Fig. 1). To the best of our knowledge, three approaches for verification have been proposed so far. In the first one, only basic filtering is applied, e. g. morphologic operations on a map of matched pairs [26]. According to the second approach, a pair of blocks is only considered forged when the neighborhoods of both blocks are also similar [18]. Lastly, the third approach handles outliers by imposing a minimum number of similar shift vectors between block-pairs. A shift vector contains the translation (in image coordinates) between two matched blocks. Consider a case where a number of blocks is copied (without rotation or scaling). Then, the histogram of shift vectors exhibits a peak at the translation parameters of the copy operation. This verification step is the most commonly used one. Note, that some methods don't use any of these methods, but just rely on a similarity criterion.

These verification methods are inherently unable to handle rotation and scaling. Thus, we propose a novel alternative verification step, the *Same Affine Transformation Selection (SATS)*. The core idea behind SATS is to explicitly estimate the affine transformation parameters of a copy-moved area. Knowing these parameters, we can transfer the key properties of shift vectors to arbitrary affine transformations of the image (as long as proper features exist). A number of rotation-invariant features for CMFD have already been proposed. Thus, in demonstrating the effectiveness of our approach we use some of the previously established rotation-invariant features.

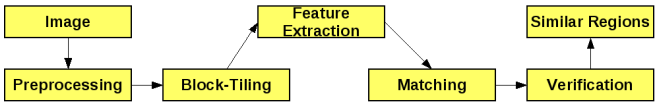

Fig. 1: Common CMFD algorithm pipeline.

3. ROTATION INVARIANT FEATURES

The selection of a suitable feature set is the core of most copy-move forgery detection methods. We evaluated the performance of existing feature sets to match similar blocks when they have undergone rotation. For this purpose, we considered 14 different features. Twelve of them can be divided in four groups: moment-based, dimensionality reduction-based, intensity-based, and frequency domain-based features (see Table 1). Additionally, two keypoint based feature vectors using SIFT and SURF features were also assessed.

For testing the feature performance for rotational CMFD, we picked two images, *tree* and *cattle* (see Fig. 2 and Fig. 3). We inserted the copied parts with rotations of 0°, 60°, 120° and 180°. Then, we subdivided the resulting images into blocks, computed the respective feature vectors per block and matched every feature vector to its nearest neighbor in feature space. Each such match constitutes a *block pair*. Note that, for this particular experiment, no additional noise has been added, since we are only interested in the performance of the features under pure rotation.

Group	Methods
Moments	MOM1 [18], MOM2 [27], MOM3 [21]
Dim.-red.	PCA [28], SVD [19]
Intensity	INT1 [15], INT2 [16], INT3 [17], INT4 [25]
Frequency	DCT [29], DWT [30], FMT [20]

Table 1: Grouping of existing copy-move forgery detection methods

As a first straightforward measure of the suitability of the features, we counted the block pairs, where one block stems from the source and one from the target region of the copied part. Features with good discriminating power which are also rotational invariant will exhibit a low number of false nearest neighbors in feature space. This will result in a high number of correctly matched block pairs. Without rotation, the best-performing feature set found 2588 correct matches in *tree*, and 4755 correct matches in *cattle*. The columns of Table 2 show the correctly matched pairs under rotations of 60°, 120° and 180°. Note that this evaluation is not completely fair towards SIFT and SURF features. These methods inherently examine fewer keypoints. Thus, they contain by definition fewer blocks that can be matched.

	cattle, max 0°: 2588			*tree*, max 0°: 4755		
Feat.	60°	120°	180°	60°	120°	180°
INT2	**2108**	**2154**	**1875**	**2628**	**2625**	**2512**
INT4	**774**	**738**	**541**	**1650**	**1663**	**1762**
MOM3	**609**	**544**	**363**	**1725**	**1686**	**1698**
INT1	736	310	184	853	539	506
MOM2	294	296	389	1172	1210	1308
FMT	54	60	766	191	199	1633
SVD	198	221	232	1084	982	913
MOM1	91	148	112	691	667	715
INT3	130	71	64	853	803	662
DWT	127	73	64	44	49	76
SURF	4	3	4	24	17	23
DCT	9	0	1	20	25	16
SIFT	1	3	1	7	8	8
PCA	0	0	0	7	1	2

Table 2: Number of correct block pair matches, for 60°, 120° and 180° rotations. For comparison, under no rotation, the best performing features found 2588 and 4755, respectively, true closest matches.

Based on the results of these experiments, we chose INT2, INT4 and MOM3 for the demonstration of our proposed method. Our findings support the claim that these three methods are rotation invariant. INT2 [16] uses the average color information of a circular block as the first three features, and the area's entropy as its fourth component. INT4 [25] uses the mean intensities of circles with different radii around the block center. Finally, the feature vector of MOM3 [21] is based on the Zernike moments of circular blocks. Bayram *et al.* also use rotational invariant features, but this property can only be exploited by an exhaustive search over all cyclic shifts of the feature vector [20], which is a prohibitively expensive computation.

Fig. 2 shows a visualization of this test. White pixels belong to block pairs, where both blocks truly belong to a copied region.

Gray pixels denote matches where at least one block is outside the copied area (and thus a false match). Finally, as a copied region has a minimum size, two blocks are not allowed to lie too close to each other. Thus, black pixels belong to matches where two blocks are located within a certain distance.

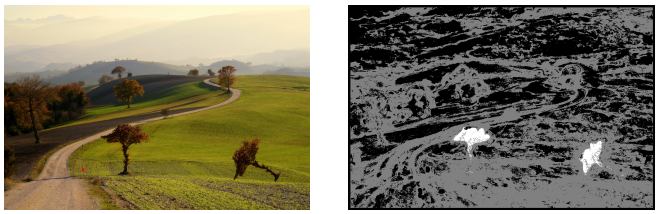

Fig. 2: Visualization of the performance check of the CMFD features under rotation. White denotes matched feature pairs where both blocks came from copy-moved regions. Gray denotes matched pairs where at least one block is in a non-copied region.

4. SAME AFFINE TRANSFORMATION SELECTION

We propose a straightforward yet effective replacement for the shift vectors, that can expressly handle affine transformations. The core idea is to explicitly estimate the rotation and scaling parameters from a few blocks, expressed as an affine transformation matrix. Starting from an initial estimate, we apply region growing on block pairs with similar transformation parameters.

Consider the i-th matched pair $\vec{f_i}$ of feature vectors $\vec{f}_{i1}$, $\vec{f}_{i2}$, $\vec{f_i} = (\vec{f}_{i1}, \vec{f}_{i2})$. In order to determine the rotation and translation between block pairs, we need to examine the coordinates of the block centers. Let $C(\vec{f}_{ij})$ denote the coordinates (in row vector form) of the block center from where $\vec{f}_{ij}$ was extracted. Further, let

$$\vec{p}_i = C(\vec{f}_{i1}), \qquad \vec{q}_i = C(\vec{f}_{i2}). \tag{1}$$

If $\vec{f_i}$ stems from a copy-move operation with rotation and scaling, then $\vec{q}_i$ is related to $\vec{p}_i$ via an affine transformation:

$$\vec{q}_i = \vec{p}_i \cdot A + \vec{b} \ , \tag{2}$$

where A is a 2×2 matrix containing rotational and scaling parameters, and $\vec{b}$ is a translation vector. The six unknowns in A and $\vec{b}$ can be found if at least three matched pairs $\vec{f_1}$, $\vec{f_2}$, $\vec{f_3}$ are available.

Equation 2 can be satisfied by searching for matched block pairs that are spatially close to each other, i.e. within a distance t_1. We recover the transformation and treat it as an initial solution to an RS-CMFD region. Then, we search for further matched block pairs that fit this hypothesis, which is iteratively refined. If the number of block pairs that satisfy the hypothesis exceeds a certain limit t_2, we consider the transformation a candidate for a copy-moved region. We report the involved blocks as well as the transformation parameters as an RS-CMFD result. SATS follows the same principles as shift vectors for robustness to outliers: clustering of similar results, and required minimum number of similar transformations. Thus, it is expected that SATS be equally robust to this type of noise. The details of the proposed verification method is shown in Algorithm 1.

SATS naturally extends the known shift vector selection. It preserves the outlier filtering properties of the shift vector approach. Furthermore, given a rotation-invariant feature set, it can handle arbitrary rotations. The incorporation of different rotation-invariant features is smoothly integrated in the RS-CMFD pipeline. In the experiments, we show how three different rotation-invariant features were used within the RS-CMFD scheme. Within this scheme, one could equally seamlessly use rotation-and-scale-invariant features. Conceptually, it is also straight-forward to extend SATS to serve as a post-processing step for keypoint-based methods.

Algorithm 1 SATS: Rotation and scale invariant verification.

for every matched pair $\vec{f_1} = (\vec{f}_{11}, \vec{f}_{12})$ **do**
  Let the hypothesis-set $H = \{\vec{f_1}\}$;
  for matches $\vec{f_i}$ **do**
    if $\mathrm{d}(C(\vec{f}_{11}), C(\vec{f}_{i1})) < t_1$ and $\mathrm{d}(C(\vec{f}_{12}), C(\vec{f}_{i2})) < t_1$ **then**
      $H = H \cup \vec{f_i}$;
    end if
  end for
  if $|H| < 3$ **then**
    continue; // At least three spatially close block pairs
  end if
  From H, compute A and $\vec{b}$ as described in the text
  for every f_i where $C(f_{i1})$ is close to matched blocks in H **do**
    compute $\vec{q_i} = \vec{p_i} \cdot A + \vec{b}$ as in Eqn. 2
    if $\mathrm{d}(C(\vec{f}_{i2}), \vec{q_i}) < t_1$ **then**
      $H = H \cup \vec{f_i}$
      if $|H| \mod 10 \equiv 0$ **then**
        recompute A and $\vec{b}$ to increase stability of the estimate
      end if
    end if
  end for
  if $|H| > t_2$ **then**
    store A, $\vec{b}$ and mark the blocks in H as copy-moved.
  end if
end for

The runtime complexity is in practice affordable, despite of the two nested loops in Algorithm 1. This is is due to the use of a greedy strategy in the selection of suitable neighbors for the initial hypothesis. Within a local region, we pick two candidates that have been mapped into the same region. Though this might be questionable from a theoretical viewpoint, we found the results to be sufficiently good in practice. Thus, the complexity mainly consists of: a) an iteration over all blocks and b) a per-block neighborhood search for suitable pairs. More precisely, let N_B be the total number of blocks in the image, N_{CB} the number of copied blocks and N the neighborhood size. Then the worst-case runtime is $O(N_B N_{CB} N)$. In practice, it is reasonable to assume that $N_{CB} \ll N_B$. Thus, the complexity is mainly influenced by the number of blocks in the image. When timing our code, we noticed that our unoptimized implementation of SATS takes at most as long as feature extraction and matching. Thus, it at most doubles the processing time for a particular image.

5. EXPERIMENTS

5.1. Dataset

We selected 10 original test images that would help us create a diverse, challenging dataset of small, as well as comparably large pictures. In each of the images one or more regions were selected for copying. The size of the regions varies among the 10 original images

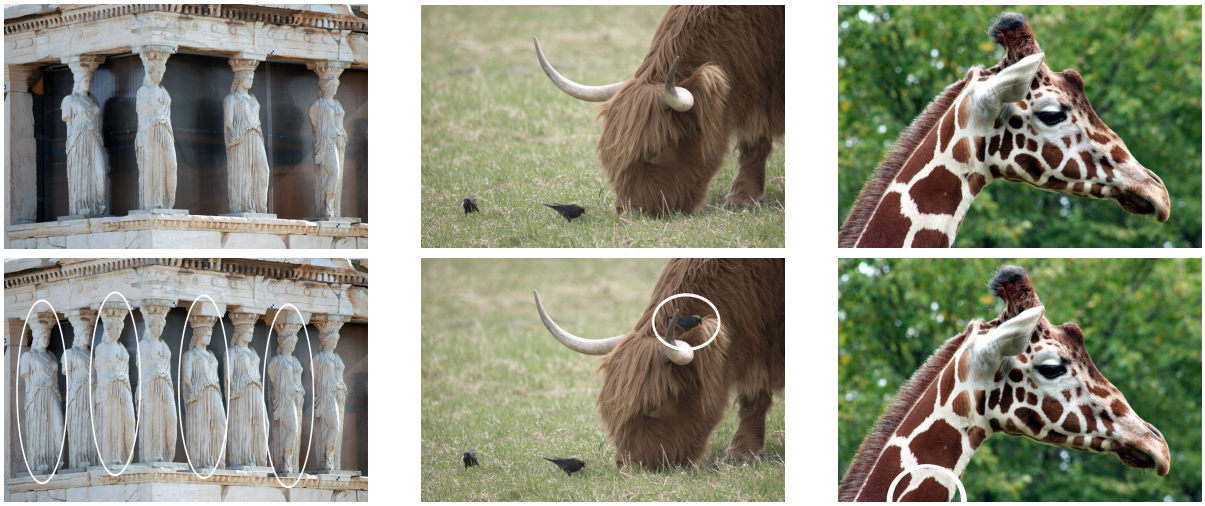

Fig. 3: Example images from our test set. Original images are in the top row, forgeries are in the bottom row. From left to right: *acropolis2* (large copied region), *cattle* (small region), *giraffe* (small region).

(see the examples in Fig. 3). Small copied regions are more difficult to detect, while larger copied regions are computationally more demanding for SATS. The copied regions were rotated by 0° to 180° angles, in steps of 15°. Thus, our dataset consisted of $10 \cdot 13 = 130$ images. Ground truth labels were created for every image. In the ground truth image, the postprocessed boundaries of the duplicated regions are marked differently than the exactly copied pixels (see Fig. 4). In our evaluation, we only use exactly copied pixels. Though many methods claim that they can detect such postprocessed (but still similar) parts, we chose to exclude them for two reasons:

1. Most CMFD methods mark blocks of copied pixels. Thus, a block on the boundary between copied and non-copied pixels is not well defined as "copied" or "original".
2. One can not clearly define a set of permissible boundary manipulations in such a way that these pixels can still be considered as copy-moved.

Hence, we believe that the cleanest solution is to exclude boundary pixels from the evaluation. Note that this definition does not hinder us to apply post-processing on the image like JPEG compression.

5.2. Detection Error Measures

We employed two basic error measurements, following the ideas of [16] and [15]. These are the percentage of erroneously matched blocks F_{P} (false positives) and erroneously missed blocks F_{N} (false negatives). More precisely, let R_1 be the copied region, $R_i, i > 1$, be the i copy-pasted regions and B the unchanged background. Then,

$$F_{\mathrm{P}} = \frac{|\text{matches in } B|}{|B|} \tag{3}$$

and

$$F_{\mathrm{N}} = \frac{|\text{missed matches in } (\bigcup_i R_i)|}{|\bigcup_i R_i|} \quad , \tag{4}$$

so that lower rates of F_{P} and F_{N} indicate higher accuracy.

Note that, as long as a copied region is detected, a high F_{P} rate is considered to be worse than a high F_{N} rate. High F_{P} rates can lead to a highly confusing overdetection result, which: a) requires a man-in-the-loop to examine every result and b) might even conceal the truly tampered regions.

Fig. 4: The image *beachwood* (upper left) is forged with a green patch (bottom left) to conceal a building (upper right). A ground truth map (bottom right) is generated where copy-moved pixels are white, unaltered pixels are black and boundary pixels are gray.

5.3. Impact of the Image Size

The ultimate goal of CMFD methods is to detect copy-move forgeries independent of the type of transformation applied to the copied region, or the type (including size) of the manipulated image. Most existing CMFD methods have been evaluated on relatively small images. Thus, before evaluating SATS itself on our dataset, which includes images of diverse sizes (see Table 3), we examined how the efficiency of the original methods may be affected by size. Once again, we focus our analysis on the best performing RS-CMFD methods: INT2, INT4 and MOM3 as presented in [16, 25, 21]. Example results on these experiments are shown in Fig. 5.

Among the three methods, MOM3 performed best since it yields a lower F_{P} rate. For all three methods the detection accuracy dropped as the image size increased. This effect was more prominent in the false positive rate F_{P}, which is often considered more

Image	x dim.	y dim.
soldiers	420	300
concrete	640	480
camen	640	480
helicopter	640	480
giraffe	800	533
tree	1024	683
cattle	1280	854
beachwood	3264	2448
acropolis2	3872	2592
swan	3888	2592

Table 3: Sizes of original images in the dataset.

critical in forgery detection. This is not surprising, since the probability of collisions in feature space between unrelated image regions increases with image size.

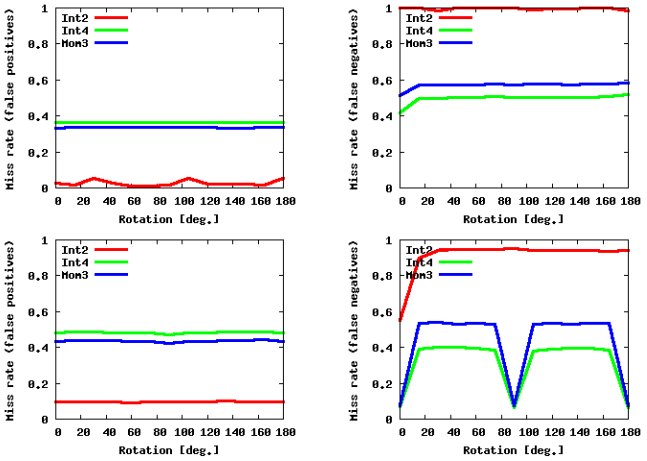

Fig. 5: CMFD performance of the original algorithms on a medium-sized image, *cattle*, (top) versus a larger image, *beachwood*, (bottom). False positives rates are on the left, false negatives rates on the right.

5.4. SATS Evaluation

The three best performing rotation invariant features (INT2, INT4 and MOM3) were also evaluated under the proposed SATS method. In our implementation of SATS, we set the neighborhood size N to 16 (i. e. we search a 4×4 grid). The distance of matched block pairs t_1 was set to 7 and the minimum number of connected matches t_2 was set to 30. For computational efficiency, the underlying feature extraction was performed on every second block. When the spatial offset between two true positive blocks is odd, a pixelwise exact match is not possible anymore. Thus, there is a trade-off between computational efficiency and feature performance. Furthermore feature extraction was only computed on those blocks with a minimum entropy of 4.0, following the idea of [20]. This drastically decreases the runtime and prevents false matches due to too uniform blocks. For the matching step we used a kd-tree as it gives fewer false positives than lexicographic sorting [22] (these steps were of course also included in the evaluation of the original methods).

Table 4 summarizes the performance of SATS in comparison to the original methods. The results show the average performance over the entire dataset. For all three features, SATS drastically reduces the false positive rate F_P, making false alarms very unlikely. This is mainly due to the clustering of transformation hypotheses controlled by t_2.

A further drawback of the original methods of INT4 and MOM3 is the proper adjustment of the Euclidian distance threhold (used as similarity criterion). This threshold depends on the image size while, when using SATS, we have a threshold, which is independent of the image size but dependent of the patch size we want to detect.

	Original method		SATS method	
Feat.	F_P	F_N	F_P	F_N
INT2	4 ± 3	96 ± 9	0 ± 0.4	22 ± 2
INT4	24 ± 19	66 ± 30	0 ± 0.0	41 ± 32
MOM3	0.4 ± 1	88 ± 24	0 ± 0.0	23 ± 1

Table 4: Comparison of the original CFMD method and the proposed SATS approach. The average F_P and F_N rates and the standard deviation are computed over the entire dataset and are given in percent.

Table 5 shows the detailed results for one of the most successful features, MOM3. The average and standard deviation over all rotation angles is depicted. Note that, consistently over all image sizes, about 75% of the copied block pairs are found. A common convention of most copy-move authors is to *mark a copy-moved region as detected, if at least one block pair is correctly matched*. Under this definition, our proposed method exhibits a 100% detection rate of CM forgeries. However, the use of a stricter evaluation measure, like F_P and F_N rates, provides a better insight on the performance of a method.

SATS with MOM3		
Image	F_P	F_N
soldiers	$0.00\% \pm 0.0\%$	$23.0\% \pm 0.9\%$
concrete	$0.00\% \pm 0.0\%$	$23.6\% \pm 0.4\%$
camen	$0.00\% \pm 0.0\%$	$23.3\% \pm 0.3\%$
helicopter	$0.00\% \pm 0.0\%$	$21.8\% \pm 0.9\%$
giraffe	$0.00\% \pm 0.0\%$	$22.1\% \pm 0.4\%$
tree	$0.00\% \pm 0.0\%$	$21.8\% \pm 0.6\%$
cattle	$0.00\% \pm 0.0\%$	$22.7\% \pm 0.8\%$
beachwood	$0.00\% \pm 0.0\%$	$23.1\% \pm 1.2\%$
acropolis2	$0.00\% \pm 0.0\%$	$22.8\% \pm 0.9\%$
swan	$0.00\% \pm 0.0\%$	$22.0\% \pm 1.2\%$

Table 5: SATS performance of the MOM3 features. The columns show the average and the standard deviation over all rotation angles.

We also tested our approach with different degrees of JPEG compression, ranging from JPEG quality 50% to 100% in steps of 10%. Since the performance of SATS-MOM3 and INT2 did not significantly vary with the rotation angle, we only tested a 90° rotation. The results over the different images were highly stable, with a F_P rate of 0% and F_N rates that were comparable to those of the uncompressed images.

6. CONCLUSIONS

We presented an extension to CMFD algorithms, which can handle general affine transformations between the copied and pasted

regions. As such, it can directly handle rotation. The proposed SATS method can smoothly replace the widely-used shift-vectors. Since the operating principle behind shift vectors and SATS is the same (clustering blocks of similar transformations), the method is expected to be similarly robust to outliers. Although not yet tested, it is reasonable to assume that keypoint-based methods can also be successfully extended using SATS. These are preliminary results. Further assessment, involving a rigorous evaluation on the impact of noise and the scale invariance of SATS, is already underway.

7. REFERENCES

[1] T. Ng, S. Chang, C. Lin, and Q. Sun, "Passive-Blind Image Forensics," in *Multimedia Security Technologies for Digital Rights*, chapter 15, pp. 383–412. Academic Press, 2006.

[2] H. Sencar and N. Memon, "Overview of State-of-the-art in Digital Image Forensics," *Algorithms, Architectures and Information Systems Security*, pp. 325–344, 2008.

[3] H. Farid, "A Survey of Image Forgery Detection," *Signal Processing Magazine*, vol. 26, no. 2, pp. 16–25, Mar. 2009.

[4] M. Chen, J. Fridrich, M. Goljan, and J. Lukáš, "Determining Image Origin and Integrity Using Sensor Noise," *Information Forensics and Security*, vol. 3, no. 1, pp. 74–90, Mar. 2008.

[5] M. Johnson and H. Farid, "Exposing Digital Forgeries through Chromatic Aberration," in *Workshop on Multimedia and Security*, New York, NY, USA, Sept. 2006, pp. 48–55, ACM.

[6] T. Gloe, K. Borowka, and A. Winkler, "Efficient Estimation and Large-scale Evaluation of Lateral Chromatic Aberration for Digital Image Forensics," in *SPIE Media Forensics and Security*, 2010, vol. 2, p. 75417547.

[7] H. Cao and A. Kot, "Accurate Detection of Demosaicing Regularity for Digital Image Forensics," *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, Dec. 2009.

[8] Z. Qu, W. Luo, and J. Huang, "A Convolutive Mixing Model for Shifted Double JPEG Compression with Application to Passive Image Authentication," in *International Conference on Acoustics, Speech and Signal Processing*, Mar. 2008, pp. 1661–1664.

[9] J. He, Z. Lin, L. Wang, and X. Tang, "Detecting Doctored JPEG Images Via DCT Coefficient Analysis," in *European Conference on Computer Vision*, May 2006, vol. 3, pp. 423–435.

[10] H. Farid, "Exposing Digital Forgeries from JPEG Ghosts," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 4, pp. 154–160, 2009.

[11] A. Popescu and H. Farid, "Exposing Digital Forgeries by Detecting Traces of Resampling," *Signal Processing*, vol. 53, no. 2, pp. 758–767, Feb. 2005.

[12] A. Gallagher and T. Chen, "Image Authentication by Detecting Traces of Demosaicing," in *Computer Vision and Pattern Recognition Workshops*, June 2008, pp. 1–8.

[13] M. Johnson and H. Farid, "Exposing Digital Forgeries in Complex Lighting Environments," *Information Forensics and Security*, vol. 2, no. 3, pp. 450–461, Sept. 2007.

[14] C. Riess and E. Angelopoulou, "Scene Illumination as an Indicator of Image Manipulation," in *Information Hiding Conference*, 2010.

[15] W. Luo, J. Huang, and G. Qiu, "Robust Detection of Region-Duplication Forgery in Digital Images," in *18th International Conference on Pattern Recognition*, Aug. 2006, vol. 4, pp. 746–749.

[16] S. Bravo-Solorio and A.K. Nandi, "Passive Forensic Method for Detecting Duplicated Regions Affected by Reflection, Rotation and Scaling," *17th European Signal Processing Conference*, Aug. 2009.

[17] H. Lin, C. Wang, and Y. Kao, "Fast Copy-Move Forgery Detection," *WSEAS Transactions on Signal Processing*, vol. 5, no. 5, pp. 188–197, 2009.

[18] B. Mahdian and S. Saic, "Detection of Copy-Move Forgery using a Method Based on Blur Moment Invariants," *Forensic Science International*, vol. 171, no. 2, pp. 180–189, Dec. 2007.

[19] X. Kang and S. Wei, "Identifying Tampered Regions Using Singular Value Decomposition in Digital Image Forensics," in *International Conference on Computer Science and Software Engineering*, Dec. 2008, vol. 3, pp. 926–930.

[20] S. Bayram, H. Sencar, and N. Memon, "An efficient and robust method for detecting copy-move forgery," *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, pp. 1053–1056, Apr. 2009.

[21] S. Ryu, M. Lee, and H. Lee, "Detection of Copy-Rotate-Move Forgery using Zernike Moments," in *Information Hiding Conference*, 2010.

[22] V. Christlein, C. Riess, and E. Angelopoulou, "A Study on Features for the Detection of Copy-Move Forgeries," in *GI SICHERHEIT*, 2010.

[23] Hailing Huang, Weiqiang Guo, and Yu Zhang, "Detection of Copy-Move Forgery in Digital Images Using SIFT Algorithm," in *Computational Intelligence and Industrial Application. Pacific-Asia Workshop on*, Dec. 2008, vol. 2, pp. 272–276.

[24] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool, "Surf: Speeded up robust features," *Computer Vision and Image Understanding*, vol. 110, pp. 346–359, 2008.

[25] J. Wang, G. Liu, H. Li, Y. Dai, and Z. Wang, "Detection of Image Region Duplication Forgery Using Model with Circle Block," in *2009 International Conference on Multimedia Information Networking and Security*, Jun 2009, pp. 25–29.

[26] A. Langille and M. Gong, "An Efficient Match-based Duplication Detection Algorithm," in *Canadian Conference on Computer and Robot Vision*, June 2006, pp. 64–71.

[27] J. Wang, G. Liu, Z. Zhang, Y. Dai, and Z. Wang, "Fast and Robust Forensics for Image Region-Duplication Forgery," *Acta Automatica Sinica*, 2009.

[28] A.C. Popescu and H. Farid, "Exposing digital forgeries by detecting duplicated image regions," Tech. Rep. TR2004-515, Department of Computer Science, Dartmouth College, 2004.

[29] Jessica Fridrich, David Soukal, and Jan Lukáš, "Detection of copy-move forgery in digital images," in *Proceedings of Digital Forensic Research Workshop*, Aug. 2003.

[30] M. K. Bashar, K. Noda, N. Ohnishi, H. Kudo, T. Matsumoto, and Y. Takeuchi, "Wavelet-based multiresolution features for detecting duplications in images," in *Conference on Machine Vision Application*, May 2007, pp. 264–267.