



HAL
open science

Dual JPEG Compatibility: a Reliable and Explainable Tool for Image Forensics

Etienne Levecque, Jan Butora, Patrick Bas

► To cite this version:

Etienne Levecque, Jan Butora, Patrick Bas. Dual JPEG Compatibility: a Reliable and Explainable Tool for Image Forensics. 2025. <hal-04681110v2>

HAL Id: hal-04681110

<https://hal.science/hal-04681110v2>

Preprint submitted on 4 Apr 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Dual JPEG Compatibility: a Reliable and Explainable Tool for Image Forensics

Etienne Levecque, Jan Butora, Patrick Bas

Abstract—Given a JPEG pipeline (compression or decompression), this paper demonstrates how to find the antecedent of an 8×8 block. If it exists, the block is considered *compatible* with the pipeline. For unaltered images, all blocks remain compatible with the original pipeline; however, for manipulated images, this is not necessarily true. This article provides a first demonstration of the potential of compatibility-based approaches for JPEG image forensics. It introduces a method to address the key challenge of finding a block antecedent in a high-dimensional space, relying on a local search algorithm with restrictions on the search space. We show that inpainting, copy-move, and splicing, when applied after JPEG compression, result in three distinct mismatch problems that can be detected. In particular, if the image is re-compressed after modification, the manipulation can be detected when the quality factor of the second compression is higher than that of the first. Through extensive experiments, we highlight the potential of this compatibility attack under varying degrees of assumptions. While our approach shows promising results—outperforming three state-of-the-art deep learning models in an idealized setting—it remains a proof of concept rather than an off-the-shelf forensic tool. Notably, with a perfect knowledge of the JPEG pipeline, our method guarantees zero false alarms in block-by-block localization, given sufficient computational power.

Index Terms—JPEG Forensics, JPEG Compatibility, Forgery localization.

The code related to this paper can be found on GitHub: <https://github.com/EtienneLevecque/jpeg-antecedent>.

I. INTRODUCTION

Nowadays, image editing has become incredibly easy, making reliable image forensic methods more crucial than ever. Almost anyone can alter an image using licensed software or free smartphone apps, and the results are often convincing. While most modifications are benign, some can become malicious. For instance, misinformation can involve deleting, adding, or altering important data in an image. Such images containing fake information can sometimes have serious consequences, including influencing public opinion in politics, falsifying scientific findings, or triggering harassment on social networks.

In this paper, we leverage the fact that digital images are usually compressed. This is a common practice on the Internet, where reducing data traffic is a top priority. Of all the image compression methods, the most widespread is undoubtedly JPEG compression (Joint Photographic Experts Group). It was one of the first methods to be adapted to many applications and to retain good visual rendering. This popularity has been maintained over time, and many cameras still save images in JPEG by default. However, as we will see later, many papers

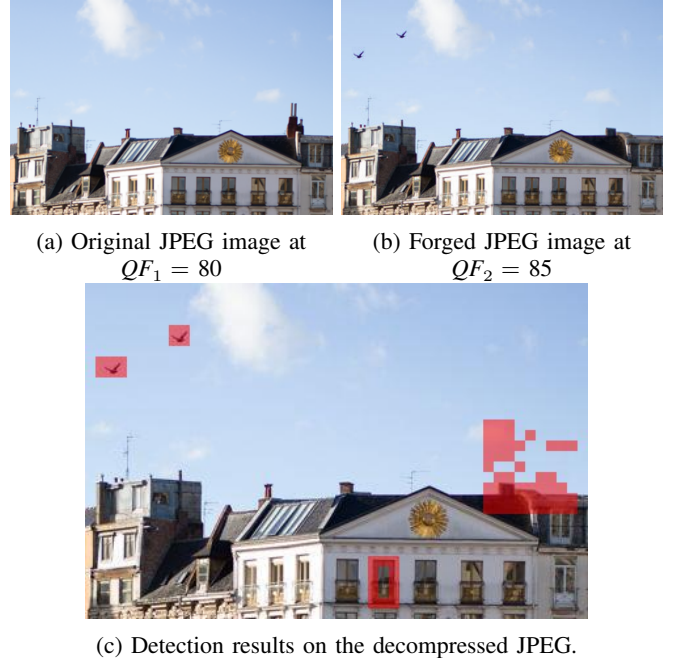


Fig. 1: Demonstration of finding incompatible JPEG blocks. The chimney was removed using in-painting (*i.e.* uncompressed data), and birds were spliced with different Quality Factors (QF s), one of them is fully aligned with the JPEG grid but not the other. The window was copy-moved aligned on the JPEG grid but not fully aligned, *i.e.* portions of blocks belong to the original image.

in the literature show that JPEG compression artifacts can be used to gather information about media authenticity.

There are many different ways of altering an image. In our case, we are focusing on images that have already undergone JPEG compression, and were then modified in the pixel domain. The first modification is the *copy-move* which involves selecting a portion of the image, duplicating it, and pasting it elsewhere. If the portion comes from another image, this is known as *splicing*. Additionally, it is possible to locally change the pixel values to add or remove content, a process referred to as *inpainting*.

Our approach fits into this context of verifying the authenticity of JPEG images and, specifically, detecting local forgeries. The proposed method does not allow for the detection of global image forgeries or the generation of entire images. Our verification mechanism is based on the concept of block compatibility: if we observe an 8×8 block \mathbf{B} , is there a block \mathbf{A} which, once passed through the JPEG pipeline, would

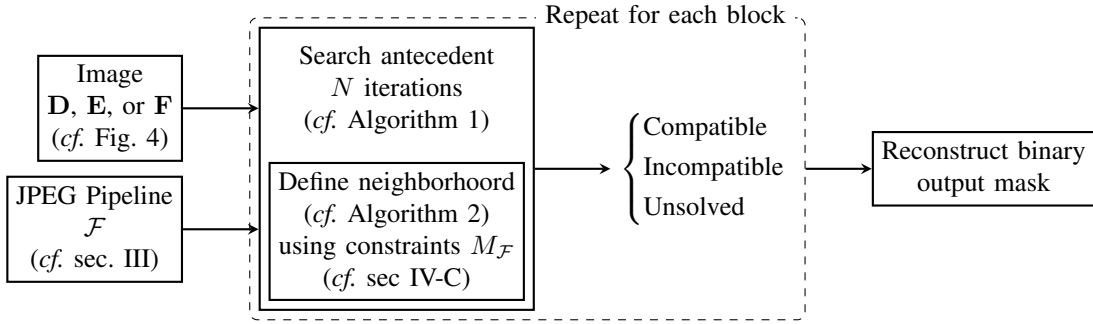


Fig. 2: General flowchart of our method. For each block of the image obtained with a pipeline \mathcal{F} , we search an antecedent with a local search constrained with upper bound $M_{\mathcal{F}}$. If the block is incompatible, we can conclude that it was altered. If the block is unsolved, the decision is left to the agent depending on the number of iterations N .

give **B**? If so, the block is compatible with the pipeline; if not, it is incompatible and has been altered. By construction, all unmodified images are compatible with their development pipeline. But when modifications are made to the result of a JPEG pipeline, it can create a block that is impossible to obtain through the same pipeline: an incompatible block, and therefore a modified block. While the idea is simple, proving whether or not a block is compatible is a complex problem in very high dimensions. However, this approach has a significant advantage: unmodified images are inherently compatible, so there are no false positives.

JPEG block compatibility has previously been applied in steganalysis [17] to detect whether compressed images at $QF100$ contained a hidden message. In that earlier work, the task of finding an antecedent was already addressed using a local search algorithm. However, this work differs in at least two directions from this previous work: it studied a different application of a similar algorithm but also improved this algorithm.

Regarding use cases, the steganalysis paper [17] mainly addresses two questions: how to find antecedents of only DCT (Discrete Cosine Transform) blocks at a fixed $QF100$ for a simple compression? And, what is the relation between manipulations in the DCT space and incompatible blocks at $QF100$? In contrast, this paper tackles broader questions: how to find antecedents of pixel and DCT blocks for any QF across combined pipelines? And, what is the relation between manipulations in the pixel space, considering any QF and multiple compressions/decompressions, and incompatible blocks?

When it comes to algorithm improvement, there are three main differences. First, it now leverages an upper bound to reduce the search space. While this upper bound also existed in the steganalysis work, it was not useful to constrain the search space due to the specific quality factor. Second, after benchmarking different norms, we found that the infinity norm was more effective than the ℓ_1 norm in the objective function. Third, the algorithm has been enhanced to handle any combination of compression and decompression, making it fully 'dual'—a capability that was not present in its initial version.

Our contributions are the following:

- We propose a dual extension of the antecedent search used in steganalysis [17]. The previous method looks for a DCT antecedent of a pixel block, whereas the proposed algorithm can fit any pipeline with any number of compressions or decompressions. This search can find antecedents in the pixel or DCT domain and can find antecedents for integer DCT pipelines.
- A phylogenetic approach of all JPEG forgery scenarios (see Section V) shows that all pixel modification can fall into three types of compatibility mismatch: the grid mismatch, the quantization mismatch, and the pipeline mismatch. Solving each type of mismatch ensures solving the vast majority of JPEG forgery.
- We analyze the impact of the Quality Factor (QF) over the compatibility result in each mismatch case and show that for some combinations of QF the search is fast on compatible blocks. Additionally, the detection is deterministic, *i.e.* contrary to statistical methods, there are no false positives: a block detected as incompatible has been tampered with.
- Finally, we conduct experiments to show that compatibility can be used to build a JPEG forgery detector that achieves state-of-the-art performance. This method is very self-explainable, can gather information about the falsified areas, and localize forgery at the JPEG block level.

Fig. 2 illustrates the global strategy of the method presented in this paper. Section II presents works related to this paper, in particular about JPEG forgery detection and JPEG compatibility. The section III defines the main notations for JPEG compression, decompression, and composed pipelines. The compatibility is presented in section IV along with the algorithm to prove that a block is compatible or not. Section V formulates all JPEG forgery scenarios into 3 mismatches that can create incompatible blocks and therefore could be detected with compatibility. Finally, section VI compares our JPEG forgery detector based on incompatibility with other JPEG-oriented detectors.

II. RELATED WORKS

To tackle the detection of falsified JPEG images, literature can be split into two categories: on one hand, there are

the statistical methods built over an underlying constrained model: if the observed features do not follow the model, the image is classified as forged. On the other hand, there are statistical methods built with data that encompass supervised deep learning approaches that extract complex features, often lacking explainability, to classify or segment the image into falsified or non-falsified regions.

A. Forgery detection using statistical model

In the first category, Lukas *et al.* [21] estimates the primary quantization table of a double-compressed image using the peaks in the DCT histograms. The main limitation is that it is hard to estimate the high-frequency quantization step, especially for low QF s, because nearly all or all the coefficients are quantized to 0.

Lin *et al.* [19] analyses the periodic pattern in the histograms of a DCT coefficient in the presence of double quantization. They show that in forged images with different regions (at least one being singly compressed and one being double compressed), it is possible to separate the blocks responsible for the periodic pattern from those that do not. This classification at the block level exhibits the localization of the forgery at the image level.

Farid *et al.* [10] propose a forgery detection based on a JPEG ghost. It is related to the convergence of multiple compression at the same quality factor: if a block was double compressed, compressing it at the same QF will return a very close value, but for an uncompressed block or block compressed with a different QF , the difference will be more important. This difference reveals a darker region called "ghost". This method was automated for forgery detection in Zach's *et al.* [27] paper. Performances are good except when the modified area has been compressed with a lower QF than the QF of the image.

Luo *et al.* [22] used the Laplacian model of AC DCT coefficients to build a statistic that can be used to identify if an image is JPEG compressed or not. It can also estimate the quantization steps or detect the quantization table in a dictionary. There is no result on forgery localization but it could be derived from the output of their method.

In Chen *et al.* [7], the block artifact periodicity is used to determine if a block has been double compressed even after a grid shift.

Finally, Bianchi *et al.* [5] presents a more detailed model of the double quantized DCT coefficients in case of aligned or non-aligned double JPEG compression. Using an estimation of the first quantization table and the grid shift, a likelihood ratio test is used to classify each block as forged or authentic.

Another category of methods relies on different versions of Benford's law [12, 23, 24] applied on DCT coefficients to derive statistical tests used to detect single/double compression or estimate the different QF used in the pipeline.

B. Forgery detection using deep learning

In the second category, we find deep learning detectors that specifically target modified JPEG images such as the convolutional neural networks proposed by Barni *et al.* [2] to detect

aligned or non-aligned double JPEG compression. Detecting double compression when the first quality factor is larger than the second one is a challenging task, but their results using the noise residuals and the pixel values are promising in this direction for non-aligned double compression. However, the generalization over multiple QF s is not achieved since they need different datasets to train multiple models.

Cat-Net, the model proposed in Kwon's *et al.* [16] paper takes these performances a step further with a good generalization over multiple forgery scenarios and QF s. This model uses DCT volume representation instead of histograms and combines it with pixel values to detect JPEG compression artifacts at a pixel level.

Also in the second category, there are deep learning models not specific to JPEG images such as ManTra-Net, the model proposed by Wu *et al.* [25]. This model targets each type of image forgeries and the JPEG compression or double compression are only two classes among many others. If their visual segmentation is very satisfying, the misclassification error between single-compressed and double-compressed is close to 85%.

Another general model is Noiseprint proposed by Cozzolino *et al.* [8] which relies on camera fingerprint estimation to detect that an image has been forged. Using contrastive learning, the authors train a model to estimate the so-called "noiseprint" of an image that can generalize this operation on unseen cameras. However, the main drawback of this scheme is its dependence on the QF since they need to train a specific detector for each quantization matrix.

TruFor is an improved version of Noiseprint proposed in Guillaro's *et al.* paper [13] which also relies on an estimation of camera noiseprint to localize forgeries. This model aims at detecting splicing but generalizes quite well on any forgery type when there is sufficient pattern in the fingerprint, which is the case with JPEG compression.

Recently, SAFIRE has been presented in Kwon's *et al.* paper [15]. The main novelty of this model is the ability to do a multi-source segmentation of forged images. Moreover, the paper show very promising localization performances in comparison with others deep learning model, in particular, outperforming Cat-Net and TruFor.

C. JPEG compatibility

The compatibility notion in JPEG blocks is not recent, the first paper to present and use it, is the one of Fridrich *et al.* [11]. To verify the compatibility of a block, they filter potential antecedent candidates with the ℓ_2 -distance from those candidates to the recompressed DCT block. They show that this distance can be upper-bounded by the norm of the worst-case scenario of rounding error. This constraint highly reduces the number of candidates for small QF s but not enough for QF s higher than 95. The method is very powerful but limited to grayscale images. The second limitation is that the guarantee for the compatibility is built with the mathematical Discrete Cosine Transform (DCT). Although the mathematical definition of the DCT is linear and bijective, most of the implementations do not respect those properties as detailed in section IV.

Compatibility has been extended to the reverse case (verifying the compatibility of a DCT block) in Butora's *et al.* [6] work. The authors propose a statistical model for the variance of the rounding error due to decompression. It is used to train a neural network to detect steganographic messages embedded in the DCT domain of JPEG images. A similar model was used in Dworetzky's *et al.* [9] work to detect steganographic messages embedded in the spatial domain.

The work of Lewis *et al.* [18] is not about detecting modification in JPEG images but doing exact recompression of a decompressed JPEG image. To do so, they try to find the exact DCT block of the observed pixel block, in other words, they search for an antecedent of a pixel block. The search is built with a very interesting set refining method: each step of the compressor tries to inverse a step of the decompressor by refining a set of possible values. At the end of the compression, the exact DCT antecedent should be the only one in the set. This method can deal with color images but the complexity at QF s higher than 90 makes it intractable.

An antecedent search was already used in [17] to find pixel antecedents of a DCT block. The methodology is similar but was applied to steganalysis at $QF = 100$. In the dual formulation formulated in the present paper (see section IV), it appears that blocks are more likely to become incompatible when the modifications are done in the pixel domain which makes the JPEG compatibility much more general and allows forgery detection for a wide range of QF s.

III. JPEG PIPELINE

A. Notations

We consider that an image is divided into non-overlapping blocks of size $(c, 8, 8)$ pixels with c the number of channels (1 for grayscale, 3 for color images). The analysis of chrominance sub-sampling is beyond the scope of this study but should be addressed in a future work. We suppose for simplicity that the image size can be divided by 8. All operations in this paper are block independent, therefore we will use the same bold letter to refer to the image or to a block of this image. We use the letter \mathbf{Q} to denote the quantization table of the same size $(c, 8, 8)$ as the block. The standard multiplication $x \times y$ and division $\frac{x}{y}$ symbols are used to define element-wise operations. The dot product $x \cdot y$ is used to refer to the matrix product. The rounding operator is defined as $[x]$. The notation tilde $\tilde{\cdot}$ is used to indicate a floating-point vector.

We only present some mathematical formulation of the JPEG pipeline that will be used in the rest of the paper.

Color transformation: Only RGB pixels undergo color transformation to become YC_bC_r pixels (denoted YCC). The mathematical definition of this transformation is linear and bijective and can be applied independently for every pixel $\mathbf{x}_{\text{RGB}} \in [0; 255]^3$ with the matrix \mathbf{T} as follows:

$$\mathbf{x}_{\text{YCC}} = \mathbf{T} \cdot \mathbf{x}_{\text{RGB}} + \begin{pmatrix} 0 \\ 128 \\ 128 \end{pmatrix}. \quad (1)$$

$$\mathbf{x}_{\text{RGB}} = \mathbf{T}^{-1} \cdot \left(\mathbf{x}_{\text{YCC}} - \begin{pmatrix} 0 \\ 128 \\ 128 \end{pmatrix} \right). \quad (2)$$

with $\mathbf{T} \cdot \mathbf{T}^{-1} = \mathbf{I}_3$ and,

$$\mathbf{T} = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ -0.168736 & -0.331264 & 0.5 \\ 0.5 & -0.418688 & -0.081312 \end{pmatrix}. \quad (3)$$

However, most implementations of this color transformation round the result to the nearest integer and use only integer and bit shift (`libjpeg` [14] for instance) making them neither linear nor invertible.

JPEG compression: The JPEG compression is applied independently on every channel of every block. Let \mathbf{x} be a pixel block. Its Discrete Cosine Transform (DCT) coefficients \mathbf{c} can be defined as follows:

$$\tilde{\mathbf{c}} = \frac{\text{DCT}(\mathbf{x})}{\mathbf{Q}}, \quad (4)$$

$$\mathbf{c} = [\tilde{\mathbf{c}}]. \quad (5)$$

where DCT refers to the forward 2D Discrete Cosine Transform function and \mathbf{Q} is the quantization table of the same size as a block, which depends on the QF of the compression. A high QF means small quantization steps (at $QF = 100$, the quantization steps are all equal to 1). DCT algorithm and rounding operation are not unique, there exist several types of implementation depending on the application. Therefore we use the notations $\mathbf{c} = \mathcal{C}(\mathbf{x}; \mathbf{Q})$ to define a specific compressor.

JPEG decompression: The decompression process is almost symmetrical. Given a DCT block \mathbf{c} , the decompressed pixel block is:

$$\tilde{\mathbf{y}} = \text{IDCT}(\mathbf{Q} \times \mathbf{c}), \quad (6)$$

$$\mathbf{y} = [\tilde{\mathbf{y}}]^{[0; 255]}, \quad (7)$$

where the notation $[\cdot]^{[0; 255]}$ means that the results of the rounding operation are clipped to the set $[0; 255]$. For the same reason as for the compressor, we use the notation $\mathbf{y} = \mathcal{D}(\mathbf{c}; \mathbf{Q})$ to define a specific implementation of a decompressor.

As said for the compressor and decompressor, there exist several implementations of the DCT and IDCT algorithms to find a good trade-off between speed and precision for each application [20, 1]. The same applies to the color transformation. In this paper, we focus on the default algorithm from `libjpeg` library [14] called `islow`. We also did all our experiments with the `libjpeg` color transformation that uses integers and thus is not lossless.

B. Composed Pipelines

To deal with double compression or more means that compression and decompression must be composed multiple times with potentially different functions and parameters. We define those notations in this subsection.

Let f_1, \dots, f_n be $n > 0$ compression or decompression functions alternated. This means that for each consecutive pair (f_i, f_{i+1}) there is a compression and a decompression. If $n = 1$, the pipeline is only composed of a single function that can be either a compression or a decompression. Each function is parameterized by its quantization tables $\mathcal{Q} = (\mathbf{Q}_1, \dots, \mathbf{Q}_n)$. We denote the composed pipeline as $\mathcal{F}(\cdot; \mathcal{Q})$.

For example, if we have f_1 a decompression with quantization matrix \mathbf{Q}_1 and f_2 a compression function with the quantization matrix \mathbf{Q}_2 , then the composed pipeline of f_1 and f_2 is \mathcal{F} parameterized with $\mathcal{Q} = (\mathbf{Q}_1, \mathbf{Q}_2)$ and for any block \mathbf{x} ,

$$\mathcal{F}(\mathbf{x}; \mathcal{Q}) = f_2(f_1(\mathbf{x}; \mathbf{Q}_1); \mathbf{Q}_2). \quad (8)$$

We also define the backward pipeline as f_b (\mathcal{F}_b for composed pipelines) that takes as argument a block from the same domain as \mathbf{y} and return a block in the same domain as \mathbf{x} . For example, if \mathcal{F} is a decompression using \mathbf{Q}_1 followed by a compression using \mathbf{Q}_2 , \mathcal{F}_b is a compression using \mathbf{Q}_2 followed by a decompression using \mathbf{Q}_1 .

Regardless of implementations used in the forward pipeline \mathcal{F} , the backward pipeline \mathcal{F}_b can be defined with any DCT, IDCT, or color transformation algorithms. Therefore, we always define it with the floating-point bijective algorithms to avoid any error due to integer operation. Moreover, we do not round or clip any result in this backward pipeline, making it perfectly reversible.

However, note that due to rounding operations, f is not an invertible function, and thus f and f_b are not inverse functions of each other: $f(f_b(\mathbf{y}; \mathbf{Q}); \mathbf{Q}) \neq \mathbf{y}$ (\mathbf{y} is an arbitrary block in the domain of f_b) and $f_b(f(\mathbf{x}; \mathbf{Q}); \mathbf{Q}) \neq \mathbf{x}$ most of the time. The same applies to \mathcal{F} and \mathcal{F}_b .

IV. JPEG COMPATIBILITY

A. Definitions and challenges

In this subsection, we define JPEG compatibility as the existence of a solution to an inverse problem.

Let \mathbb{S} be the set of pixel blocks or the set of DCT blocks. Let $\mathbf{y} \in \mathbb{S}$ be a block and let f be a function of compression or decompression such that the codomain (set of destination) of f is \mathbb{S} . Let \mathbf{Q} be the quantization table that parameterizes f . We say that \mathbf{y} is *compatible* with f if there exists a block antecedent \mathbf{x} such that:

$$f(\mathbf{x}; \mathbf{Q}) = \mathbf{y}. \quad (9)$$

On the contrary, if no antecedent \mathbf{x} exists, we say that \mathbf{y} is *incompatible* with f .

We can easily generalize this definition to any composed pipeline \mathcal{F} with the set of quantization tables \mathcal{Q} and codomain \mathbb{S} :

$$\mathcal{F}(\mathbf{x}; \mathcal{Q}) = \mathbf{y}. \quad (10)$$

This inverse problem is complex because the pipeline is not an invertible function and thus applying the backward pipeline to \mathbf{y} will return a close candidate but not always an antecedent. Even if we suppose that the uniform norm between this close

Algorithm 1 Local search to find antecedent

Require: \mathbf{y}, \mathcal{Q} # Observed block and quantization tables
Require: \mathcal{F} # Pipeline
Require: $N > 0$ # Max iteration
 $\mathbf{x}_s \leftarrow \mathcal{F}_b(\mathbf{y})$ # Starting block of the search
add \mathbf{x}_s to P with cost 0 # Priority queue initialization
while P not empty **and** $k \leq N$ **do**
 $k \leftarrow k + 1$
 $\mathbf{x} \leftarrow$ remove first element of P
for \mathbf{x}_n in $neighbors(\mathbf{x})$ **do**
if \mathbf{x}_n has not been visited **then**
if $\mathcal{F}(\mathbf{x}_n; \mathcal{Q}) = \mathbf{y}$ **then**
return \mathbf{x}_n # \mathbf{x}_n is an antecedent of \mathbf{y}
end if
 $c_n \leftarrow \|\mathbf{y} - \mathcal{F}(\mathbf{x}_n; \mathcal{Q})\|_1$
add \mathbf{x}_n to P with cost c_n # \mathbf{x}_n is a new candidate
end if
end for
end while

candidate and the closest antecedent is 1, this means that for each element in the blocks, the error can be one of -1, 0, or 1. This sums to 3^{64} possibilities for grayscale blocks and 3^{192} possibilities for color blocks without colors sub-sampling. Brute forcing all solutions can not be considered, instead, we will rely on a local search to find antecedents combined with theoretical constraints to reduce the search space.

B. Local search to find antecedents

The local search is detailed in algorithm 1 and shares some similarities with the one presented in [17]. However, this one can be used with any composed or simple pipeline and uses the ℓ_1 norm instead of the infinity norm.

The search begins by applying the backward pipeline to the observed block to obtain a starting candidate. Then, until the number of iterations is reached, we select the best candidate from a priority queue (queue data structure but every element has a priority; highest priority is always first), explore all its neighbors, and feed the priority queue again depending on the ℓ_1 distance to the observation. If the distance is 0, then an antecedent has been found and the search returns the result. A hash table is used to ensure that each block is visited only once.

The neighborhood of a candidate is defined in Algorithm 2 with $M_{\mathcal{F}}$ an upper bound depending on the pipeline \mathcal{F} that will be defined in the next subsection. However, when the block \mathbf{x} is clipped, we ignore this upper bound and always add the new candidate to the neighborhood. Note that we only apply changes of ± 1 to one single value in the block but this process is repeated indefinitely until the maximum number of iterations is reached. Hence, every block in the constrained search space can be explored.

Note that to find an antecedent, this local search algorithm can be considered as fast. Indeed, the worst case is for $QF = 100$ because lots of candidates are very close to the solution and can be solved in a matter of minutes for most blocks.

Algorithm 2 Neighbors function

Require: \mathbf{y} # Observed block
Require: \mathbf{x} # Block for which a neighborhood is required
Require: $\tilde{\mathbf{x}}_b$ # Floating point starting point
Require: $M_{\mathcal{F}}$ # Upper bound of the pipeline

if \mathbf{y} is clipped **then**
 $M_{\mathcal{F}} = \infty$
end if

for i in channels **do**
 for j in rows **do**
 for k in columns **do**
 $\mathbf{x}_n \leftarrow \text{copy}(\mathbf{x})$
 $x_{n,i,j,k} \leftarrow x_{n,i,j,k} + 1$ # Positive change
 if $\|\mathbf{x}_n - \tilde{\mathbf{x}}_b\|_2 \leq M_{\mathcal{F}}$ **then**
 add \mathbf{x}_n to the neighborhood
 end if
 $\mathbf{x}_n \leftarrow \text{copy}(\mathbf{x})$
 $x_{n,i,j,k} \leftarrow x_{n,i,j,k} - 1$ # Negative change
 if $\|\mathbf{x}_n - \tilde{\mathbf{x}}_b\|_2 \leq M_{\mathcal{F}}$ **then**
 add \mathbf{x}_n to the neighborhood
 end if
 end for
 end for
end for

return neighborhood

However, the problem is for incompatible blocks, the search could explore the whole intractable space before emptying the priority queue. We can avoid that by constraining the searching space and thus the neighborhood.

C. Additional theoretical constraints

In order to speed up the search of incompatible blocks we also decided to see to which extent we can use the criterion proposed in the paper of Fridrich *et al.* [11] where the authors defined a theoretical constraint for grayscale candidate blocks that drastically reduce the search space for $QF < 95$. This constraint is presented in this subsection as well as the generalization to color images and double compression. All computations are done under two assumptions. First, we are using the mathematical definition of the DCT, the IDCT, and the color transformation such that all of them are unitary invertible functions. Second, we suppose that there is no clipping in the sense that all pixel values are always in the interval $[0;255]$. This second assumption is necessary otherwise, the rounding error can be bigger than $1/2$ and the constraints will not hold.

Let's start with intermediate results that will simplify the equations. Let $\tilde{\mathbf{y}}$ be a floating point block of size 8×8 and $\mathbf{y} = \lfloor \tilde{\mathbf{y}} \rfloor$ its rounding value. We define the rounding error as $\mathbf{e} = \mathbf{y} - \tilde{\mathbf{y}}$. We have the following upper bound on the rounding

error norm:

$$\begin{aligned} \|\mathbf{y} - \tilde{\mathbf{y}}\|_2 &= \sqrt{\sum_{i=1}^8 \sum_{j=1}^8 |y_{i,j} - \tilde{y}_{i,j}|^2} \\ &\leq \sqrt{\sum_{i=1}^8 \sum_{j=1}^8 \frac{1}{2}} \\ &\leq 4. \end{aligned} \quad (11)$$

Now if we have a color transformation, we use the under-script RGB to denote the $3 \times 8 \times 8$ tensor with the three channels, and Y , R , G , and B for each individual channel. The same result applies to other channels:

$$(T \cdot \mathbf{e}_{RGB})_Y = 0.299\mathbf{e}_R + 0.587\mathbf{e}_G + 0.114\mathbf{e}_B. \quad (12)$$

Since each row of T sums to one, we can invoke the triangle inequality together with equation (11) to upper bound each channel by 4 and obtain:

$$\|T \cdot (\lfloor \tilde{\mathbf{y}}_{RGB} \rfloor - \mathbf{y}_{RGB})_Y\|_2 \leq 4. \quad (13)$$

Note that we have the same upper bound when doing the inverse color transform.

We now derive two upper bounds, one that is used when searching for an antecedent of a compression, and one used when searching antecedent of a decompression.

Let us consider a simple compression setup with notations \mathbf{x} , \mathbf{Q} , $\tilde{\mathbf{c}}$ and \mathbf{c} defined in equations (4) and (5). This compression can be modeled by a pipeline f and its backward equivalent f_b . We suppose that we observe \mathbf{c} and would like to find the antecedent \mathbf{x} from the starting point $\mathbf{x}_b = f_b(\mathbf{c})$. To speed up the search we will derive an upper bound to the distance between the true antecedent and the starting point:

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}_b\|_2 &= \|\mathbf{x} - \text{IDCT}(\mathbf{Q} \times \mathbf{c})\|_2 \\ &= \|\mathbf{x} - \text{IDCT}(\mathbf{Q} \times \mathbf{c} - \mathbf{Q} \times \tilde{\mathbf{c}} + \mathbf{Q} \times \tilde{\mathbf{c}})\|_2 \\ &= \|\mathbf{x} - \mathbf{x} - \text{IDCT}(\mathbf{Q} \times \mathbf{c} - \mathbf{Q} \times \tilde{\mathbf{c}})\|_2 \\ &= \|\mathbf{Q}(\mathbf{c} - \tilde{\mathbf{c}})\|_2 \\ &\leq M_f = \left\| \frac{\mathbf{Q}}{2} \right\|_2 \end{aligned} \quad (14)$$

Indeed, the ℓ_2 -norm is invariant by unitary transformation (the IDCT) and the inequality comes from (11) with a quantization table.

Now, let's consider a decompression setup with notations \mathbf{Q} , \mathbf{c} , $\tilde{\mathbf{y}}$ and \mathbf{y} defined in equations (5), (6) and (7). This decompression can also be modeled by a pipeline f and its backward equivalent f_b . We suppose that we observe \mathbf{y} and would like to find the antecedent \mathbf{c} from the starting point $\mathbf{c}_b = f_b(\mathbf{y})$. Using a very similar calculus, we derive an upper bound to the distance between the true antecedent and the starting point:

$$\begin{aligned} \|\mathbf{Q}\mathbf{c} - \mathbf{Q}\mathbf{c}_b\|_2 &= \|\mathbf{Q}\mathbf{c} - \text{DCT}(\mathbf{y})\|_2 \\ &= \|\mathbf{Q}\mathbf{c} - \text{DCT}(\mathbf{y} - \tilde{\mathbf{y}} + \tilde{\mathbf{y}})\|_2 \\ &= \|\mathbf{Q}\mathbf{c} - \mathbf{Q}\mathbf{c} - \text{DCT}(\mathbf{y} - \tilde{\mathbf{y}})\|_2 \\ &= \|\mathbf{y} - \tilde{\mathbf{y}}\|_2 \\ &\leq M_f = 4. \end{aligned} \quad (15)$$

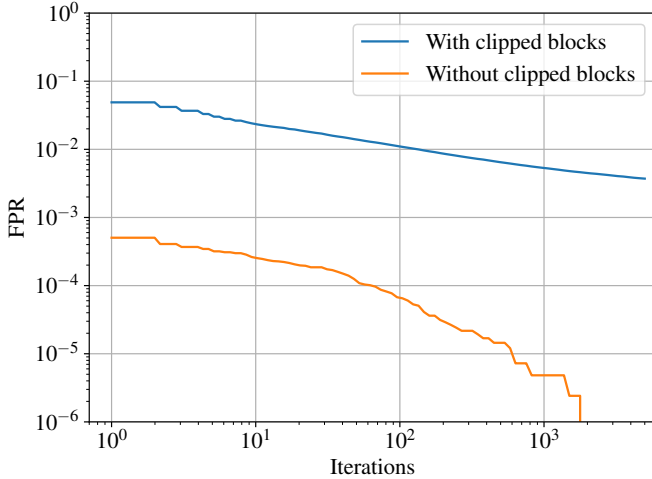


Fig. 3: Probability for compatible blocks to be unsolved after a given number of iterations (FPR). This empirical curves were obtained by averaging the results of all single and double compression experiments ($\sim 1\text{M}$ blocks) from section VI.

Using the triangle inequality of the norm, this constraint can be generalized to any pipeline by adding the upper bounds together. For example, if we want to search for an antecedent through a pipeline of decompression (upper bound is 4) with color transform (upper bound is also 4), the distance between the starting point and the true antecedent should not exceed $M_{\mathcal{F}} = 4 + 4$. In the double compression scenario, the pipeline is composed of a decompression ($M_f = 4$), a color transform ($M_f = 4$), a compression with quantization table \mathbf{Q} ($M_f = \|\mathbf{Q}/2\|_2$), another decompression ($M_f = 4$) and a final color transform ($M_f = 4$), then the distance should not exceed $M_{\mathcal{F}} = 4 + 4 + \|\mathbf{Q}/2\|_2 + 4 + 4$.

Finally, each upper bound obtained defines a sphere around the starting point in which the true antecedent should be. If we explore all candidates inside this sphere and do not find a solution to our inverse problem (9), then we can conclude that the "true antecedent" does not respect this property and the block is incompatible.

D. Analysis of the number of iterations

The main parameter of the local search algorithm is the number of iterations. This subsection will present a short study of this parameter to help the reader select the correct value and see the influence of clipped blocks. Each block can be either compatible if an antecedent has been found, incompatible if no antecedent has been found among all possible candidates or unsolved if the algorithm reaches the maximum number of iteration before the two other cases.

Fig. 3 highlights the impact of clipped blocks on the number of iterations. If we ignore them, the algorithm is able to find antecedent for all of the $\sim 1\text{M}$ compatible blocks in less than 2000 iterations and only 1 block in 5000 (equivalent to 1 block in an image of 560×560) is still unsolved after 1 iterations. However, clipped blocks are an issue because we can not use the reduction of the search space and we can see that even

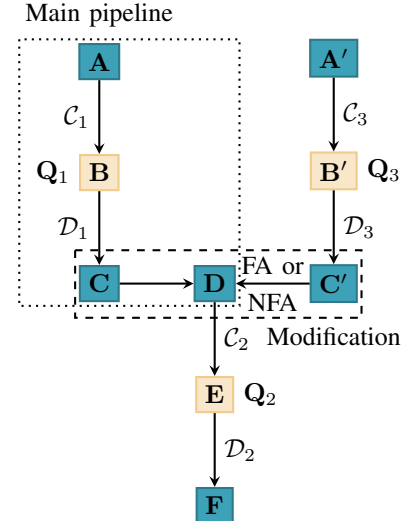


Fig. 4: General scenario for modifying a JPEG image in the spatial domain. Blues boxes represent blocks in the spatial domain (pixels) and orange boxes represent blocks in the frequency domain (DCT coefficients). \mathcal{C} and \mathcal{D} are the compressors and decompressors used. \mathbf{Q} are the quantization tables. Modifications applied to \mathcal{C} can be fully aligned (FA) or not fully aligned (NFA).

after 2000 iterations, there is still 1 block in 500 which is clipped and which is still unsolved.

V. COMPATIBILITY OF FORGED JPEGS

We describe in this section the different types of forgery scenarios, their associated incompatibilities, and the results associated with each scenario.

A. Addressed forensics scenarios

There are several ways of forging a JPEG image and this section presents the underlying assumptions regarding the possible JPEG compression operations and the different scenarios that follow.

In order to verify the JPEG compatibility of an image, we need at least a part of it to come from a decompressed JPEG image. We assume that there exists a pixel image \mathbf{A} , that has been compressed with a compressor \mathcal{C}_1 and quantization table \mathbf{Q}_1 to obtain the DCT coefficients \mathbf{B} . Once decompressed with a decompressor \mathcal{D}_1 using the same quantization table \mathbf{Q}_1 , we obtain \mathbf{C} . This pixel image \mathbf{C} is the one that will be falsified.

The general scenario, which includes all the others, is illustrated in Fig. 4. In this graph, the double arrows towards \mathbf{D} mean that the image \mathbf{D} is built by using the values of the pixels from \mathbf{C} or \mathbf{C}' and placing them at any position in \mathbf{D} . The other common part of each scenario relies on a second compression which can be applied on \mathbf{D} .

Note that the notion of fully aligned or not fully aligned altering is inspired by the definitions of aligned and non-aligned recompression formalized by Bianci and Piva [4], and is illustrated in Fig. 5. A modification is *fully aligned* (FA) if the JPEG grid of \mathbf{C} and the JPEG grid of \mathbf{C}' are the same

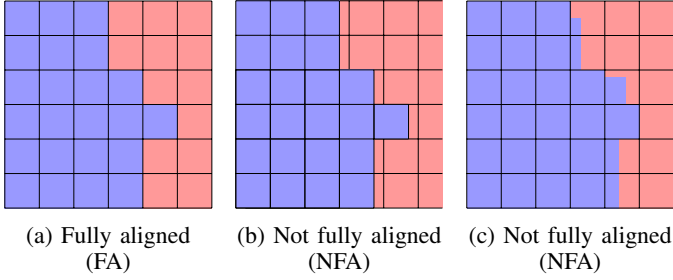


Fig. 5: Examples of fully aligned and not fully aligned modifications. Black lines represent blocks. Example 5b illustrates block translations from one image to another, on the contrary example 5c illustrates a finer segmentation.

and if all modifications are done block-wise. For example, if a block of \mathbf{D} has some pixel values from \mathbf{C} and some other pixel values from \mathbf{C}' , then the modification is not full and therefore *not fully aligned* (NFA) (cf. Fig. 5c). Note that the fully aligned situation is very rare when doing JPEG forgery. Indeed, if the grids can be aligned with a probability of $\frac{1}{64}$, blocks that form the boundary of the modification will most likely be cut, especially if a blur filter is applied to hide the contour artifacts.

Three specific forensics scenarios are described as follows, each encompassing fully aligned or not fully aligned modifications.

Inpainting. In this scenario, the modified pixel comes from a natural image \mathbf{C}' that has not been compressed. Therefore, \mathcal{C}_2 and \mathcal{D}_2 do not exist and the modifications are necessarily non-aligned since \mathbf{C}' is not associated with a JPEG grid.

Copy-Move. In this scenario, $\mathbf{A} = \mathbf{A}'$, $\mathbf{B} = \mathbf{B}'$ and $\mathbf{C} = \mathbf{C}'$ (compressors and decompressors are also equal) but the image \mathbf{D} is built by taking some pixels from \mathbf{C}' and placing them somewhere else. The modifications can be aligned on the JPEG grid or non-aligned.

Splicing. This is the most general scenario. Two images are compressed with potentially different compressors with some quantization tables \mathbf{Q}_1 and \mathbf{Q}_3 that can be equal or different. The decompressed versions of those images are used to build \mathbf{D} . The modifications can be aligned on the JPEG grid or non-aligned.

B. Possible sources of incompatibility

This subsection translates the forgery scenarios described in the previous one into three distinct mismatches that can create incompatible blocks. We denote them as the *grid* mismatch, the *quantization table* mismatch, and the *pipeline* mismatch.

Grid mismatch appears when the dependency between the 64 values of a block (decompressed pixels or DCT coefficients) is broken. It can be broken either because the block was partially modified and therefore there is no dependency between the two parts of the block (NFA copy-move or NFA splicing) or because the modified pixels are simply not JPEG decompressed and thus not dependent (inpainting).

Quantization table mismatch can only be observed when the forgery is fully aligned and when the quantization table

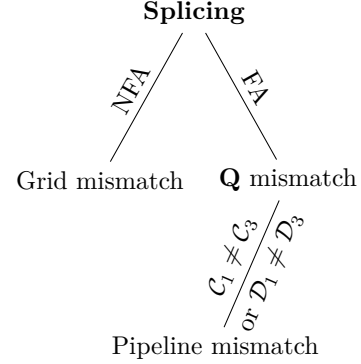


Fig. 6: In the splicing scenario, if the modifications are not fully aligned (NFA), it is a *grid* mismatch, in the fully aligned case (FA), there can be a *quantization table* mismatch if tables are not the same. *Pipeline* mismatch can appear if one of the compression or decompression (or both) differs. Note that a FA splicing with the same quantization tables and functions is a *perfect* splicing in the sense that it does not create any mismatch.

used is not the same as the one used in the main pipeline. In this case, the 64 decompressed pixel values are dependent and are not consistent with the original pipeline. This mismatch can for example be obtained when doing fully aligned splicing.

Pipeline mismatch can also only be observed when doing fully aligned splicing. It comes from the fact that the functions used in the compression and decompression process can be different. For example, if the DCT transform does not have the same implementation or when the rounding function is not the same. It comes from any implementation difference that influences the actual result. In particular, Fig. 6 illustrates the possible mismatches for the splicing scenario and shows that one can observe a single pipeline mismatch if compressors are not the same but decompressors are the same, or a double pipeline mismatch if both compressors and decompressors are different.

In this classification, we assume there can only exist a single mismatch. Indeed, if there is a grid mismatch, it is impossible to define the quantization table, either because it does not exist at all or because it would mean that there are multiple quantization tables used in the image. For the same reason, the pipeline can not be correctly defined, therefore, the quantization table mismatch and the pipeline mismatch do not exist. Using the same reasoning, we assume that if the manipulation is fully aligned (there is no grid mismatch) then if a quantization mismatch exists we can not define the pipeline mismatch.

In our forgery scenario, the agent performing the analysis can observe either \mathbf{D} , \mathbf{E} , or \mathbf{F} . Depending on the observed image and the type of mismatch, it may be necessary to search for an antecedent further up the pipeline. For example, to verify a grid mismatch while observing \mathbf{E} , we need to find a DCT antecedent \mathbf{X} through the combination of \mathcal{D}_1 and \mathcal{C}_2 such that $\mathcal{C}_2(\mathcal{D}_1(\mathbf{X})) = \mathbf{E}$.

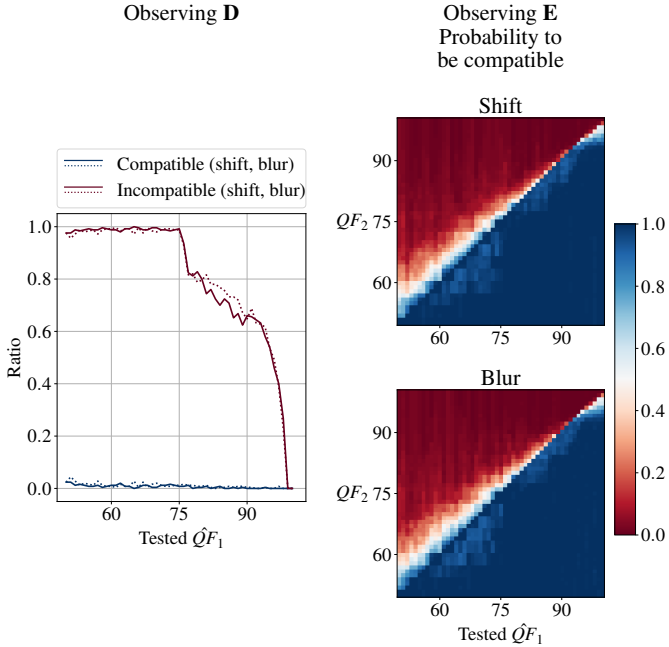


Fig. 7: Statistics in the presence of a grid mismatch due to a grid shift or a blurring kernel to remove the grid. The left plot gives ratio of compatible (blue) and incompatible (red) blocks when observing **D** and the right plots when observing **E** (almost equal to plots when observing **F**). Tested \hat{QF}_1 is the QF used to search an antecedent. Note that, all blocks have been modified in this experiment.

C. Statistics of incompatible blocks

In the previous section, three possible mismatches and three observations yielded nine different experiments. This section presents some experiments to explore the statistics of incompatible blocks and their relation to the QF . We used images from the UCID dataset composed of 1338 RGB color images of size 384×512 . These pixel images are equivalent to **A** in our general scenario (Fig. 4). We use the `islow` DCT method from the `libjpeg` library using all QF between 50 and 100 to create the compressed image **B**. To obtain the decompressed pixel image **C** we use the decompressor of the same library. In this subsection, QF_1 and QF_2 refer to the true quality factors of a block and we also use them to denote the quantization table associated with this QF . However, the tested \hat{QF}_1 refers to the QF used in the pipeline of the search.

Statistics on grid mismatch: We have seen in the last subsection that grid mismatch can be caused by a grid shift (such as non-aligned splicing or copy-move) or by inpainting a region without a grid at all. To explore both cases, we try to find antecedents to blocks modified with a random misalignment (called shift) and blocks modified with a blurring kernel (called blur). We can draw the following conclusions by looking at the results in Fig. 7:

- When observing **D** (*i.e.* forged blocks in the pixel domain), blocks with a grid mismatch are very unlikely to be compatible with any tested \hat{QF}_1 (probability of being compatible almost equals to 0). For small enough tested




Fig. 8: Statistics in the presence of a quantization mismatch. Left plots give statistics when observing **D** and right plots when observing **E** (almost equal to plots when observing **F**). Tested \hat{QF}_1 is the QF used to search an antecedent.

\hat{QF}_1 we can even prove that blocks are incompatible (probability of being incompatible almost equals to 1). The search space becomes bigger with \hat{QF}_1 , so it requires more iterations in the search Algorithm 1 to explore all possible candidates. For this experiment the number of iterations was fixed to 100.

- When observing **E** (*i.e.* recompressed forged blocks), blocks with a grid mismatch are very unlikely to be compatible for tested \hat{QF}_1 small enough compared to QF_2 (above diagonal tested $\hat{QF}_1 = QF_2$).
- But when observing **E**, if the tested \hat{QF}_1 is too big compared to QF_2 , blocks with a grid mismatch are very likely to be compatible to this tested \hat{QF}_1 (below diagonal tested $\hat{QF}_1 = QF_2$).
- Observing **E** or **F** gives very close results.
- Shift and blur yield the same probability, so an inpainted block has the same probability of being incompatible as a non-fully aligned spliced or copy-move block.

Statistics on quantization mismatch: This time, observed blocks are compatible to a true QF_1 but we want to evaluate how likely they are to be compatible to another tested \hat{QF}_1 . Fig. 8 depicts the results of our experiments and again we can draw the following conclusions:




Fig. 9: Result of the algorithm in the presence of a pipeline mismatch for a single compression after 5000 iterations. The original image was compressed with `islow` from `libjpeg` at $QF_1 = 75$ but the algorithm used in the antecedent search is `float` from `libjpeg`. In red, are incompatible blocks, in pink unsolved blocks, and blue compatible blocks. The blue blocks in all the predictions are constant blocks that are always compatible with any DCT algorithm.

- When observing **D**, blocks with a quantization mismatch are very unlikely to be compatible with any other tested $Q\hat{F}_1$ different than the true QF_1 . For small enough true QF_1 we can even prove that blocks are incompatible with any other pipeline using a different $Q\hat{F}_1$. We observe the same separation as for simple compression around true $QF_1 = 75$, which is due to the size of the search space compared to the maximum number of iterations in the search Algorithm 1.
- When observing **E**, blocks with a quantization mismatch are very unlikely to be compatible with any other tested $Q\hat{F}_1$ different than the true QF_1 if QF_1 is small enough compared to QF_2 (below the line true $QF_1 = QF_2$).
- However, when observing **E**, if QF_1 is bigger than QF_2 , blocks with a quantization mismatch are very likely to be compatible with any tested $Q\hat{F}_1$ (above the line true $QF_1 = QF_2$).
- Observing **E** or **F** gives very close results.

Notice that in both Fig 7 and Fig. 8, when QF_2 is smaller than QF_1 , all blocks become compatible. This occurs because a compression with a smaller QF acts as a projection in a smaller space (both in size and in dimension since some coefficient at set to 0). So if a block is incompatible because of a manipulation after the first compression, this second compression acts as a projection on a set of compatible blocks, and therefore, the block after the second compression is compatible.

Pipeline mismatch: Localizing a manipulation because of the pipeline mismatch is very unlikely to happen because this would imply that the manipulation is fully aligned and that there is no quantization mismatch, two rare scenarios. Nonetheless, it's useful to explore this concept, especially when the pipeline is not fully known. For instance, consider the situation where we attempt to find the antecedent of compatible blocks using an incorrect DCT algorithm. We conduct a small experiment where an image is compressed at $QF_1 = 75$ using the `islow` DCT algorithm from `libjpeg`. After the compression, this image is then manipulated. Next, we apply the antecedent search but with the `float` DCT algorithm also presented in `libjpeg` and for different $Q\hat{F}_1$ used in the search. In Fig. 9 we see that most blocks are classified as either incompatible or unsolved, with only a small fraction becoming compatible when $Q\hat{F}_1 = 75$, which is the correct QF . This experiment highlights how sensitive the algorithm is to variations in the pipeline. If the pipeline is

not known, this experiment implies that it should be possible to do a dictionary attack using multiple pipelines and most of them will not work but the correct one.

To sum up the statistics in this subsection, if we observe the image **D** which has been modified without recompression, there is a very high chance that modified blocks due to inpainting, splicing, or copy-move will be incompatible and detected as such (we assume that fully-aligned copy-move or fully-aligned splicing with the same quantization table are rare events). And if we observe **E** (or **F**, the results should be almost the same) which have been recompressed with QF_2 , we should be able to detect modified blocks due to inpainting, splicing or copy-move if QF_1 is smaller than QF_2 . Otherwise, we will not be able to detect anything.

VI. FORGERY LOCALIZATION

A. Experimental setup

Because of our working assumptions we need to control the JPEG pipeline used to create manipulated images. Thus, we took the first 20 images from the Flickr30k dataset [26] and compressed and decompressed them at $QF_1 \in \{60, 75, 90\}$ using a known JPEG pipeline to obtain 3×20 images. The DCT and IDCT algorithms used were `islow` from `libjpeg` with standard quantization tables. Then, we used a method called Beyond-the-Brush [3] to manipulate all the 60 images (inpainting) without any human supervision. In the experiment, we consider the case where the manipulated images have not been compressed again after modification and the case where the image undergone a second compression at $QF_2 \in \{50, 60, 75, 90, 95\}$ to obtain at the end $3 \times 5 \times 20 = 300$ manipulated images. The flowchart in Fig. 10 illustrates this dataset creation.

The number of iterations of our algorithm is arbitrary fixed at $N = 5000$. For most compatible blocks, a solution will be found in fewer iterations. However, for clipped blocks, with at least one pixel at 0 or one pixel at 255, the equations presented in section IV-C are no longer true because the rounding error can be bigger than $\frac{1}{2}$. Therefore, in presence of a clipped block, the algorithm will not apply the reduction of the search space which drastically increases the number of iterations.

To compare our method, we use the TruFor model [13], the Cat-Net v2 model [16] and the SAFIRE [15] model. All of them are deep neural networks trained mainly on splicing and/or copy-move datasets, however, they are all based on




Fig. 10: Images forgery process for both simple and double compression experiments. C_1, D_1, C_2 are the compressions and decompression defined in Fig.4 such as the quality factors QF_1 and QF_2 .

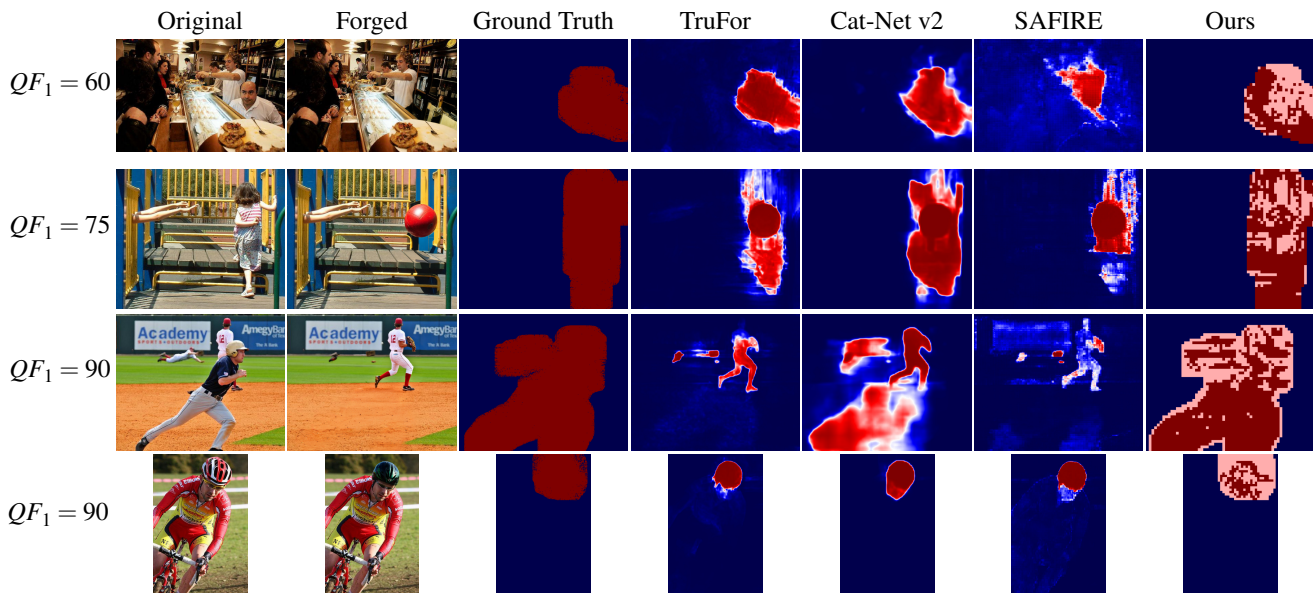


Fig. 11: Visual comparison of inpainting localization. Original images have been compressed using $QF_1 \in \{60, 75, 90, 90\}$ from top to bottom. Then, they have been decompressed, manipulated and stored in PNG. Note that the masks and the predictions are at the pixel level and not at the block level like the metrics. In ours predictions, pink (resp. red) corresponds to unsolved blocks (resp. incompatible) but they are both classified as manipulated.

artifact fingerprints and the JPEG compression applied before the modification should be sufficient to have two different areas in the image: one with the JPEG artifacts and one without. Moreover, Cat-Net v2 is specifically trained to detect JPEG artifacts using an RGB stream combined with a DCT stream. These three models output a probability map, so to obtain a binary map we use a threshold of 0.5 as proposed in Cat-Net v2 paper [16].

The metrics used are all permuted metrics, this means that we take the best metric using the prediction or 1 minus the prediction. This is done to evaluate how well models can distinguish each region. Moreover, every metric is computed by blocks. To get the state of the block we apply the same rule to both the mask and the prediction: a block is classified as manipulated if at least one pixel of this block is classified as such. Finally, the two metrics used are the balanced accuracy, defined as:

$$ACC = \frac{TPR + TNR}{2}, \quad (16)$$

and the False Positive Rate $FPR = FP/(FP + TN)$, where

$TPR = TP/(TP + FN)$ denotes the True Positive Rate, $TNR = TN/(FP + TN)$ the True Negative Rate, TP the True Positives, FN the False Negatives, FP the False Positives and TN the True Negatives.

B. Results

	Balanced Accuracy (ACC)	False Positive Rate (FPR)
TruFor [13]	73.62	0.21
Cat-Net v2 [16]	75.54	0.48
SAFIRE [15]	67.13	6.65
Ours	99.98	0

TABLE I: Localization results on 60 images of size between 496×280 and 496×496 . The balanced accuracy and the false positive rate are defined at the block level. Results has been averaged over $QF_1 \in \{60, 75, 90\}$. There was no visible dependency on the QF .

For the single compression setup, Fig. 11 shows a visualization of the output of our method compared to other forgery

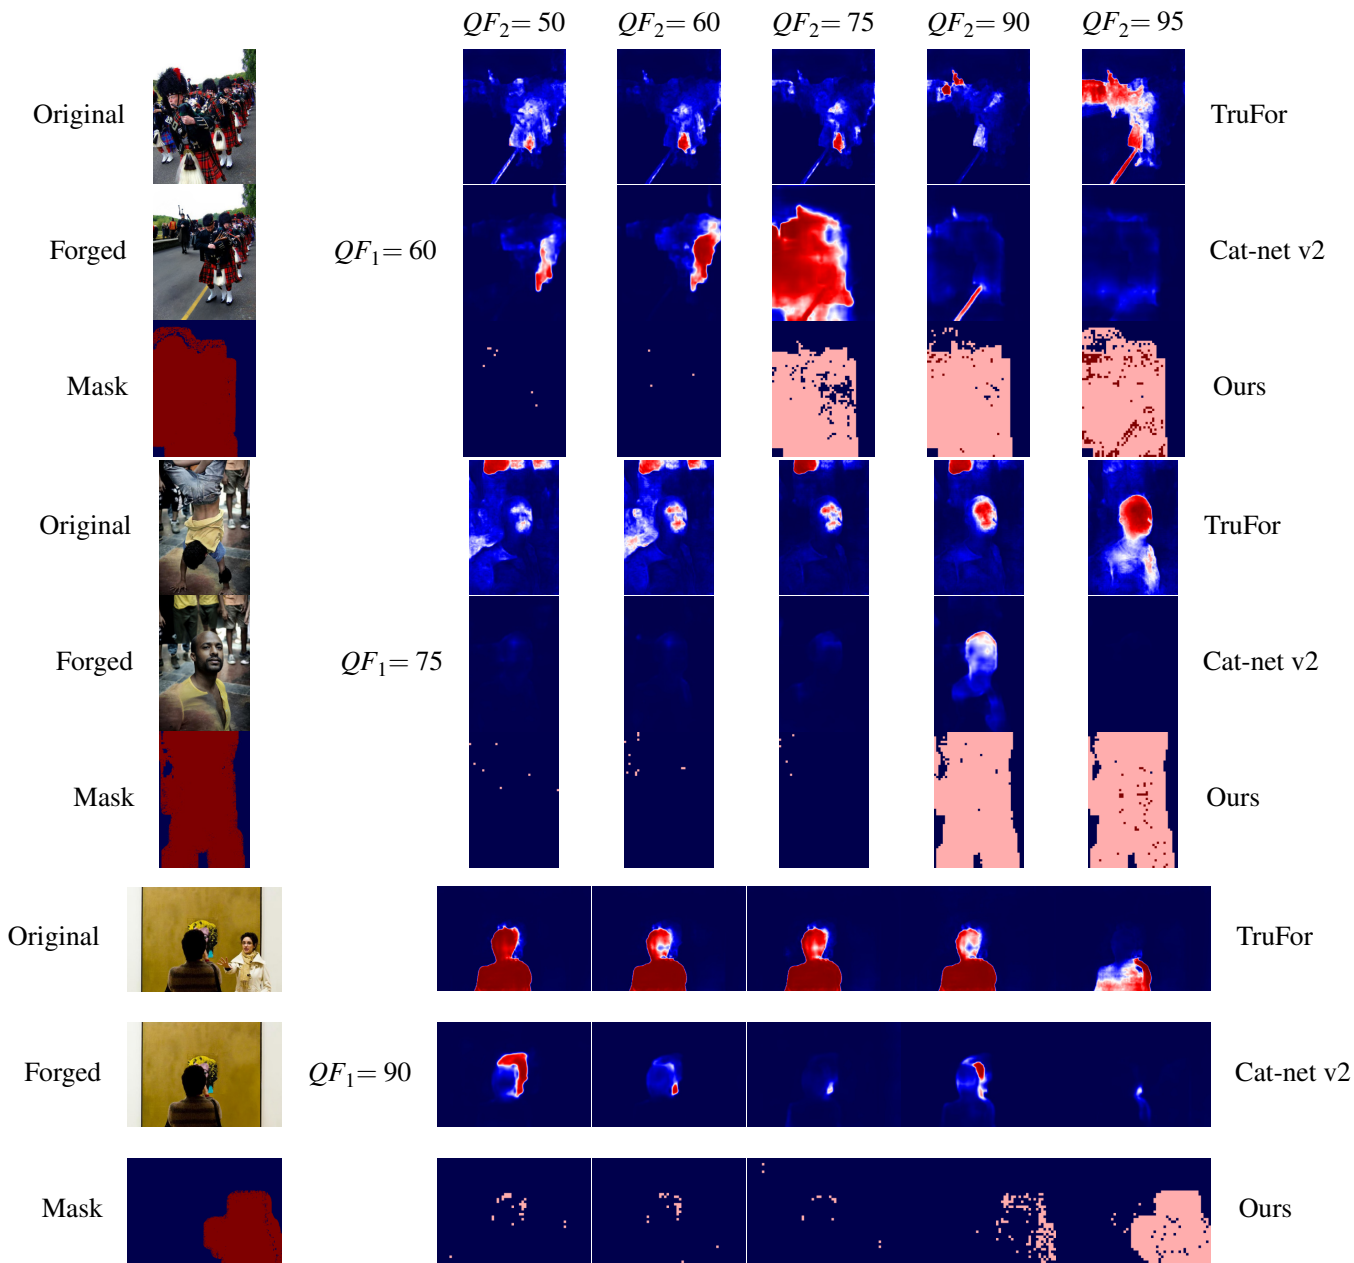


Fig. 12: Localization results in double-compressed inpainted images. Each image has been compressed at QF_1 , modified in the pixel domain, and recompressed at QF_2 . The detection is done on the pixels values \mathbf{F} for all models. In our predictions, pink corresponds to unsolved blocks, and red to incompatible blocks with $N = 5000$ iterations. There are only few incompatible blocks because theoretical constraints in double compression are not tight enough to efficiently reduce the search space.

localization methods, and Table I shows the metrics associated with this experiment. We can see that all models are relatively good for this task but our method is better, especially in terms of false positives. These results match the statistics obtained in the last section: all compatible blocks are detected as is and some incompatible blocks are also found. The pink blocks in our prediction represent unsolved blocks. Those are clipped blocks for which the theoretical upper bound can not be applied so the searching space is not reduced and the algorithm is not able to explore all the space.

For the double compression setup, the visual results are

presented in Fig. 12. As expected, when QF_1 is bigger than QF_2 , almost all blocks are compatible (some clipped blocks are unsolved) and we are not able to detect any forgery. However, in the other case, when QF_1 is smaller than QF_2 , the detection is very good compared to the state-of-the-art methods. Note that when $QF_1 = QF_2 > 85$ we are also able to detect parts of the forgery. Moreover, in a double compression setup, we can see that blocks are mostly either compatible or unsolved and there are almost no incompatible blocks. This is because the upper bound is not tight enough and is not able to reduce the searching space enough to bring it to a reasonable

	QF1	60					75					90				
	QF2	50	60	75	90	95	50	60	75	90	95	50	60	75	90	95
TruFor [13]	ACC	56.3	55.4	64.3	70.4	72.4	62.2	58.6	56.3	71.6	71.6	57.8	58.2	56.8	55.7	68.8
	FPR	3.4	2.6	1.5	0.8	0.1	3.4	1.8	1.6	0.8	0.1	4.0	3.7	2.6	0.9	0.4
Cat-Net v2 [16]	ACC	52.2	53.6	68.3	61.3	57.6	56.8	67.4	52.8	64.6	60.1	53.0	53.5	56.5	50.6	63.5
	FPR	4.0	2.4	1.3	0.5	0.5	13.0	1.9	5.9	0.1	0.1	6.2	3.1	0.5	2.6	0.0
SAFIRE [15]	ACC	59.6	62.8	61.9	64.0	65.7	62.6	62.4	58.9	60.3	63.8	59.7	63.0	60.6	60.2	63.5
	FPR	6.7	19.8	14.3	16.0	13.4	17.0	17.3	12.7	12.2	2.3	22.2	8.7	13.7	8.6	8.3
Ours	ACC	50.0	50.0	92.5	98.6	99.4	50.3	50.1	50.1	97.8	99.2	50.4	50.6	50.4	54.6	99.0
	FPR	0.0	0.0	0.1	0.0	0.0	0.3	0.3	0.0	0.0	0.0	1.4	1.0	0.8	0.1	0.3

TABLE II: Localization results averaged over 20 images of size between 496×280 and 496×496 . The balanced accuracy (ACC) and the false positive rate (FPR) are defined at the block level.

size that could be entirely explored in 5000 iterations.

Table II show the block-wise metrics associated with this experiment. Again the metrics of our method match very closely to the statistics of the last section. In particular, Table II shows a FPR almost null at 5000 iterations (computed over more than 1M predictions). There is also no detection when QF_1 is bigger than QF_2 which was expected based on the last section statistics but note that, other state-of-the-art methods also have poor results in this setup. When QF_1 is smaller than QF_2 , our method outperforms other methods with a very high detection power.

VII. CONCLUSION AND PERSPECTIVES

A. Conclusion

This paper has presented a new approach to solving the problem of verifying the authenticity of an image under the assumption that they were JPEG compressed before a potential forgery. Our method is based on the notion of compatibility of JPEG blocks (existence of an antecedent) or incompatibility (absence of antecedent) which makes it very reliable under some assumptions (no False Positive). We extended this notion of compatibility to any number of compression, decompression or color transforms. We proposed a local search algorithm to look after the antecedent and combined it with a theoretical upper bound to reduce the search space.

We have seen that the type of modification (inpainting, splicing, or copy-move) can create up to 3 types of mismatch (grid, quantization table, and pipeline mismatch). We then studied the link between mismatch type and block compatibility. In particular, we saw that when we look at the pixel image without recompression, almost all mismatched blocks will be incompatible and easy to detect. But when looking at the recompressed (or recompressed then decompressed) image, some modified blocks may remain compatible because of the quantization tables used.

Finally, we conducted a forgery localization experiment to compare with state-of-the-art methods. It appears that without recompression, our method can accurately detect every modified block without error and thus outperforms other methods. If the image is recompressed after modification, our method can be very accurate when the second compression is weaker than the first one and also outperforms other methods. However,

if the second compression is stronger than the first one, our method is unable to detect anything.

B. Perspectives

This work should be seen as a proof of concept of using JPEG compatibility to detect modified images. Indeed, this method requires strong assumptions to work properly. In particular, the exact JPEG pipeline of the original image must be known up to its DCT algorithm implementation. For future work, there are at least two axes to transform this method into an off-the-shelf model.

The first axis is to make it more robust to unknowns. We could start to generalize it to unknown quantization tables by estimating a dictionary of potential tables and verifying the compatibility using each table. But the real challenge is to make it robust to any DCT (and IDCT) algorithm to be able to handle any image without knowing its main pipeline. The second axis is to improve the speed of the algorithm. For high QF , the algorithm requires thousands of iterations for a single block to ensure no False Positive. This means that for big images, it can take hours to analyze them. The speed could be improved either by finding tighter upper bounds to reduce the search space or by finding a new algorithm to find antecedent (or prove incompatibility). In particular, we believe the upper bound could be improved with a probabilistic constraint on the neighborhood. This would create a trade-off between a controlled False Positive Rate and the algorithm efficiency.

Finally, this method only relies on the property of non-surjectivity of compression and decompression processes. Such property could exist in other compressed formats. For example, forgery H26x or HEIC formats could be the source of incompatibilities.

ACKNOWLEDGMENTS

The work presented in this paper received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 101021687 (project “UNCOVER”).

REFERENCES

- [1] Yukihiro Arai, Takeshi Agui, and Masayuki Nakajima. A fast dct-sq scheme for images. *IEICE Transactions (1976-1990)*, 71(11):1095–1097, 1988. 4

- [2] M. Barni, L. Bondi, N. Bonettini, P. Bestagini, A. Costanzo, M. Maggini, B. Tondi, and S. Tubaro. Aligned and non-aligned double JPEG detection using convolutional neural networks. *Journal of Visual Communication and Image Representation*, 49:153–163, 2017. 3
- [3] Giulia Bertazzini, Chiara Albisani, Daniele Baracchi, Dasara Shullani, and Alessandro Piva. Beyond the brush: Fully-automated crafting of realistic inpainted images. In *2024 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2024. 10, 11
- [4] Tiziano Bianchi and Alessandro Piva. Detection of nonaligned double jpeg compression based on integer periodicity maps. *Transactions on Information Forensics and Security*, 7(2):842–848, 2011. 7
- [5] Tiziano Bianchi and Alessandro Piva. Image forgery localization via block-grained analysis of jpeg artifacts. *IEEE Transactions on Information Forensics and Security*, 7(3):1003–1017, 2012. 3
- [6] Jan Butora and Jessica Fridrich. Reverse jpeg compatibility attack. *IEEE Transactions on Information Forensics and Security*, 15:1444–1454, 2019. 4
- [7] Yi-Lei Chen and Chiou-Ting Hsu. Detecting recompression of jpeg images via periodicity analysis of compression artifacts for tampering detection. *IEEE Transactions on Information Forensics and Security*, 6(2):396–406, 2011. 3
- [8] Davide Cozzolino and Luisa Verdoliva. Noiseprint: A CNN-Based Camera Model Fingerprint. *Transactions on Information Forensics and Security*, 15:144–159, 2020. 3
- [9] Eli Dworetzky, Edgar Kaziakhmedov, and Jessica Fridrich. Advancing the jpeg compatibility attack: Theory, performance, robustness, and practice. In *Proceedings of the 2023 ACM Workshop on Information Hiding and Multimedia Security*, pages 67–79, 2023. 4
- [10] Hany Farid. Exposing digital forgeries from JPEG ghosts. *Transactions on Information Forensics and Security*, 4(1):154–160, 2009. 3
- [11] Jessica Fridrich, Miroslav Goljan, and Rui Du. Steganalysis based on JPEG compatibility. In *Multimedia Systems and Applications IV*, volume 4518, pages 275–280. SPIE, 2001. 3, 6
- [12] Dongdong Fu, Yun Q Shi, and Wei Su. A generalized benford’s law for jpeg coefficients and its applications in image forensics. In *Security, Steganography, and Watermarking of Multimedia Contents IX*, volume 6505, pages 574–584. SPIE, 2007. 3
- [13] Fabrizio Guillaro, Davide Cozzolino, Avneesh Sud, Nicholas Dufour, and Luisa Verdoliva. Trufor: Leveraging all-round clues for trustworthy image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20606–20615, 2023. 3, 10, 11, 13
- [14] IJG. libjpeg. <https://libjpeg.sourceforge.net/>, 1990. 4
- [15] Myung-Joon Kwon, Wonjun Lee, Seung-Hun Nam, Minji Son, and Changick Kim. Safire: Segment any forged image region. *arXiv preprint arXiv:2412.08197*, 2024. 3, 10, 11, 13
- [16] Myung-Joon Kwon, Seung-Hun Nam, In-Jae Yu, Heung-Kyu Lee, and Changick Kim. Learning JPEG Compression Artifacts for Image Manipulation Detection and Localization. *International Journal of Computer Vision*, 130(8):1875–1895, 2022. 3, 10, 11, 13
- [17] Etienne Leveque, Jan Butora, and Patrick Bas. Finding incompatible blocks for reliable jpeg steganalysis. *Transactions on Information Forensics and Security*, 19:9467–9479, 2024. 2, 4, 5
- [18] Andrew B. Lewis and Markus G. Kuhn. Exact JPEG recompression. In *Visual Information Processing and Communication*, volume 7543, pages 256–264. SPIE, 2010. 4
- [19] Zhouchen Lin, Junfeng He, Xiaoou Tang, and Chi-Keung Tang. Fast, automatic and fine-grained tampered jpeg image detection via dct coefficient analysis. *Pattern Recognition*, 42(11):2492–2501, 2009. 3
- [20] Christoph Loeffler, Adriaan Ligtenberg, and George S Moschytz. Practical fast 1-d dct algorithms with 11 multiplications. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 988–991. IEEE, 1989. 4
- [21] Jan Lukáš and Jessica Fridrich. Estimation of Primary Quantization Matrix in Double Compressed JPEG Images. *Digital Forensic Research Workshop*, 2003. 3
- [22] Weiqi Luo, Jiwu Huang, and Guoping Qiu. Jpeg error analysis and its applications to digital image forensics. *IEEE Transactions on Information Forensics and Security*, 5(3):480–491, 2010. 3
- [23] Cecilia Pasquini, Giulia Boato, and Fernando Pérez-González. Multiple jpeg compression detection by means of benford-fourier coefficients. In *International Workshop on Information Forensics and Security (WIFS)*, pages 113–118. IEEE, 2014. 3
- [24] Ali Taimori, Farbod Razzazi, Alireza Behrad, Ali Ahmadi, and Massoud Babaie-Zadeh. Quantization-Unaware Double JPEG Compression Detection. *Journal of Mathematical Imaging and Vision*, 54(3):269–286, 2016. 3
- [25] Yue Wu, Wael AbdAlmageed, and Premkumar Natarajan. Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9543–9552, 2019. 3
- [26] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014. 10
- [27] Fabian Zach, Christian Riess, and Elli Angelopoulou. Automated image forgery detection through classification of JPEG ghosts. In *Joint DAGM (German Association for Pattern Recognition) and OAGM Symposium*, pages 185–194. Springer, 2012. 3