# Digital Image Manipulation Forensics

*Anthony Sutardja*
*Yan Zhao*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2015-125
http://www.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-125.html

May 15, 2015

# DIMF

## DIGITAL IMAGE MANIPULATION FORENSICS

UC BERKELEY MASTER OF ENGINEERING

EECS CAPSTONE PROJECT

Omar Ramadan

Anthony Sutardja

Yan Zhao

Advisor: Professor James F. O'Brien

# Table of Contents

*For the UC Berkeley Master of Engineering Capstone Project*

# Introduction

## Problem Statement

For the past century, photographs have served as reliable primary sources of evidence, but that is quickly changing. Photo manipulation tools have become widespread and it is easy to manipulate images. Photo manipulations tools such as Adobe PhotoShop afford greater artistic expression, and enable users to create manipulations that challenge the limits of our natural perception. The difference between authentic and manipulated photos has become harder to distinguish, and can only be detected by digital forensic experts. The image forensics capstone project aims to create an online software service that performs the work of forensic analysts, and visualizes and analyzes the possible manipulations that may have been performed on an image. Our goal is to empower anybody to perform image forensic analysis, and help increase our faith in digital content.


Screenshots of our web detection service in action. *Left*: the landing page where a user can upload a suspected image. *Middle*: the progress page where users await detection results. *Right*: a section of the results page with our image manipulation analysis.

Our capstone has completed a prototype of this image manipulation detection web service that helps a user identify fraudulent features within an image. Although fully functional, our service is not yet ready for commercialization as many of the techniques are still under research. In the following sections, we perform a market, industry, and intellectual property analysis on how we would hypothetically take our online software service to market.

# Trends, Market, and Strategy

## Trends

There are strong technological and regulatory trends that reveal a growing need for effective image forensics solutions. Smartphones and the cameras embedded within them have rapidly become the most popular method of photo capture. According to the Pew Research Center, the percentage of U.S. adults who own smartphones jumped from 35% in 2011 to 56% in 2013 (Smith, "Smartphone Ownership 2013"). This relentless growth in smartphone adoption will continue, especially in regions outside of the United States (Dulaney 5). More people have the ability to take digital photos, as seventy-six percent of smartphone users say they use the smartphone's cameras to take photos (Smith, "Mobile Access 2010"). This instant and convenient photo-snapping ability is becoming recognized by many institutions as a way of documenting incidents and filing claims ("How to Document Auto Accident Damage.") (Thomson 3). However, as we increase our reliance on digital photography, we become susceptible to fraud. Digital image editing software is easy to obtain and is growing in utilization as well (Kahn 9). With the growing trend in smartphone adoption and the rise in access to image editing software, digitally manipulated images will become commonplace.

Due to the growing trend of digital image documentation, the regulatory landscape is also changing with new court precedents that are beginning to require verifying the authenticity of digital images. In Israel, a new law makes it illegal "to use images in advertisements that have been retouched to make models look thinner without printing a disclosure on the picture ("Picture Imperfect")". In the United States, the Federal Rules of Evidence now require the authentication of digital images before they are used as evidence in court (Thomson 3). The growing trend in regulatory measures for verifying digital images will only continue to push the need for image forensics service.

Hence, identifying these trends has helped us realize the customers of the image forensics industry. In journalism, we see the need to verify digital images for journalistic integrity. In the courtroom, we see the need to verify digital images to establish evidentiary authenticity. In insurance, we see the need to identify fraudulent digital insurance claims. As more services begin to rely on digital documentation, the need for such media verification services will increase.

# Market

## Potential Markets

There are several markets in which an image forensics solution is greatly needed.

Our technology has serious implications for the press image market. The exploding growth of the Internet and digital content has empowered citizen journalism and transformed the press industry with exclusive rich media collections. With this growth, press companies now have the additional challenge of maintaining their credibility by validating media from new unverified sources. Photo editors at such institutions work to ensure that only genuine photos are published, but with advanced editing tools, even an experienced analyst can fail to identify image manipulations. The $32 billion dollar newspaper industry takes photo manipulation as a serious threat because of the damage it brings to their credibility (McKenna). Their zero tolerance policy towards photo manipulation frequently results in hundreds of thousands of dollars in losses; each manipulated photo is issued a "mandatory kill" that alerts all affected customers to control damages (Lum). The involved photographers are usually terminated such as in the case of Getty freelancer Marc Feldman (Lum), L.A Times staff photographer Colin Crawford (Irby), or Reuters freelancer Adnan Hajj ("Altered Images Prompt Photographer's Firing").

Another potential market resides in the court of law. Admitting manipulated photographic evidence in court is both a growing trend and a growing concern. Perjury is not uncommon, and police officers are estimated to commit perjury twenty to fifty percent of the time on fourth amendment issues (Slobogin). Lawyers also commit perjury, and even good lawyers and are untruthful in their profession. As it stands, the authenticity of photographs is contingent on the honesty of the witnesses. An image forensics expert is only called upon when there are no witnesses to testify in a lawsuit or when the photographic evidence has its integrity challenged (Kashi). Otherwise, there is no standard procedure to have the photo evidence thoroughly vetted. The $400 billion dollar legal industry needs a better way to admit photographic evidence in a court of law that is more robust than a single testament of the truth (Novet). We see our image forensic software as an opportunity to change the way photo evidence is processed.

## Segmented Market Entry

The last market that we have identified is with insurance companies. Our capstone team has decided to focus on catering our image forensics software service towards the insurance segment due to the great value we will provide to insurance companies. This value stems from the major losses insurance companies face due to fraudulent claims. While digitally documented photography and documentation has made it more cost effective for the insurance industry to survey damages, it has exposed them to significant losses. This shift towards digital evidence has made it easier to tamper with photographs to exaggerate or fabricate insurance claims. One study found that one in seven hundred general insurance claims has been digitally altered, and one in seventy-five property damage insurance claims has been digitally altered ("Picture Imperfect"). These digitally altered insurance claims account for a significant portion of the $40 billion dollars lost by the insurance industry every year in the United States (FBI). These monumental losses indicate that insurance companies have strong needs for innovative ways to combat fraudulent claims. Our image forensics software seeks to address this problem by helping insurance companies identify fraudulent claims accurately, and in a cost effective manner.

## Marketing Strategy

### Product

The detection of manipulated images is an important problem that requires the consultation of highly specialized forensic experts. Given the growing nature of this problem, we propose a service that puts the analytical power of image forensic methods in the hands of the common user. We generate detailed reports that can be interpreted by non-experts, effectively allowing lower level analysts to create more value while relieving the workloads of busy forensic experts.

Our product increases the capacity of analysts; it provides them with forensic information that helps them make informed decisions when evaluating claims. While images submitted with a claim are only a portion of the evidence that analysts use to base their decisions, the information we provide helps them identify and flag suspicious claims for further review, ultimately enabling them to prevent more fraud.

## Value Guarantee

The largest risk for our customers is a negative return on investment, in that the cost of our platform is not exceeded by the value of fraud that we stop. To reduce this risk, and promote the sale of our service, our strategy is to minimize the software acquisition risks for our customers. We can reduce the switching costs by supplying our own forward deployed engineers to integrate with our customer's software stacks and our own support team to provide training services. Furthermore, with efficiency our service offers and its effectiveness in combating fraud, we can offer a guaranteed return on investment.

## Customer Outreach

Customer acquisition is the largest challenge that we will face. Given the business-to-business nature of our business model and that customers are likely to have a large bureaucracy with several decision makers, our sales timeline will likely take months. To generate leads, we plan to attend digital forensic conferences and insurance fraud conferences to promote our services.

# Industry

The image forensics services industry is sub-industry of the much larger digital fraud detection services industry. By focusing on the insurance market first, we balance our needs for large potential customer volume with our customers' urgent needs for a solution to their fraud problem. With a large number of insurance corporations across the country that face large insurance fraud losses, we can mitigate the bargaining power of buyers and justify big ticket sales. Furthermore, the alternatives to our service are weak in this market, so the barriers to entry are low. Additionally, our software as a service (SaaS) model allows us to replicate and scale our offerings with ease.

## Substitute Offerings

The use of digital evidence has become accepted practice. In the case of automobile insurance claims, agencies actually encourage the behavior. The DMV.org procedures for filing an auto insurance claim suggests using a "smartphone or camera to take pictures at the scene" of the automobile damages ("How Car Insurance Companies Investigate Accident Claims"). These images are easily manipulated to exaggerate damages to increase liability or remove evidence to lessen liability. Insurance industry studies reveal that 10 percent or more of claims filed are fraudulent (GEICO, "GEICO's Special Investigations Unit"), many of which include digitally edited photos (Verifeyed, "Verifeyed – Image, Video and PDF Forensics, Authentication, Manipulation Detection, and and Digital Security"). By analyzing claims for fraudulent images, we can greatly improve the efficiency of the claim review process.

## Current Methods

The details as to how insurance companies handle insurance claims varies based on the nature and severity of the claim, as well as the company's own policies. However, there are certain steps which are common amongst investigations.

Claims process begins when a report of damage is submitted to the insurance company using an online form or by directly calling the company's claim division. The claimant is expected to supply sufficient documentation that allows the company to understand the circumstances of the accident, and estimate the damages associated with it.

The case created is assigned to an employee of the insurance company. These employees have a number of titles across companies including Liability Investigators or Claim Adjusters, Appraisers, and Investigators (U.S. Bureau of Labor Statistics). However, the responsibilities of such personnel are fairly consistent; their role is to determine whether the insurance policy will pay out, and if so, how much it will cover.

The first step of evaluating the claim involves surveying the damage, and checking the case for inconsistencies. This is achieved through in-depth interviews with the involved parties and witnesses followed by a comparison of all accounts (GEICO, "How GEICO Investigates a Claim"). The interviews are also compared with information from the police report, and the investigator reviews auxiliary documents such as photos associated with the incident and public information from social media (GEICO). In the event that a case is flagged by an examiner as being suspicious, the case is reassigned to a Special Investigations Unit (SIU) which examines the case closely for evidence of fraud (GEICO). If the images associated with the claim come into question, the SIU will consult an image forensics expert.

The main problem in identifying fraudulent cases is the intense labor that it requires. Before a case is even flagged as being fraudulent and referred to the SIU, it has already consumed the resources of the claim adjuster. In the case of low-value claims where it is cheaper to pay the claimant rather than verify the damage, image verification can cut costs.

Manipulated images are dangerous to claim investigators. In the case of a fraudulent claim, it is unlikely that a perpetrator would submit photographic evidence that contradicts the claim story. Manipulated photographs can go unnoticed to the naked eye, and since "seeing is believing," convincing fakes can bring investigators to lapses in judgement. While there are forensic tools that can be used to verify photos, such tools require technical analysis skills that are beyond the scope of a case examiner's capabilities.

Another option insurance companies exercise in the face of fraud is to simply be passive. Losses are a common occurrence in the insurance market, and the way they are dealt with is by passing them onto their customers in the form of increased premiums (State Farm®). Not only does it make insurance less affordable, but according to the National Insurance Crime Bureau (NICB), insurance fraud is also correlated with higher taxes and inflated prices for consumer goods (National Insurance Crime Bureau).

### Our solution in comparison

Our capstone solution provides an analytical report for photographs uploaded to our web service. Instead of using photographs as a means to validate a fraudster's narrative, we advocate using the images as one of the primary means for detecting fraud. As soon as a case is uploaded to the insurance company, the photographic contents can be automatically processed to allow case examiners and SIUs to quickly identify fraudulent claims from our reports. With our simple interface that localizes the manipulated regions within a photograph, we empower analysts to perform forensic tasks that were only accessible to experts. Our forensic service allow analysts to detect more fraud, and work through more cases by increasing their efficiency.

Our software-as-a-service platform can be integrated into the existing workflows of analysts, creating synergy. Compared to current substitutes for detecting fraud, the superior accuracy and labor savings that our workflow introduces is likely to yield a near immediate return on investment.

## Competitive Landscape

Our capstone team is not the first group of people who have identified the value to be restored to the insurance companies from detecting manipulated images. We have identified two companies that our product will directly potentially compete with.

### Competitors

System of Methods and Tools of Digital Processing Technology LLC, known as SMTDP is a Russian technology company founded in 2011 that focuses on automated business processes and image manipulation detections ("Company Overview of SMTDP Technology, LLC"). Some of the technologies used by SMTDP are image metadata analysis and image compression analysis. SMTDP's business mainly focuses on the development of technology and relies on partnership agreements with other companies, like Belkasoft and PricewaterhouseCoopers, to distribute and utilize their products ("Company Overview of SMTDP Technology, LLC"). Since SMTDP interacts only with value added resellers, the customers that receive the end user product depends entirely on these partners. Being based in Saint Petersburg, SMTDP is limited in its customer reach. SMTDP instead works with its partners to create third-party products that integrate their image manipulation detection methods

(SMTDP, "SMTDP | Company"). This is a model that leads to customer support and integration difficulties as opposed to working directly with customers to meet their needs.

Besides SMTDP, the second largest competitor in the digital image forensics space is Verifeyed, a privately owned company located at Czech Republic (Verifeyed, "Verifeyed | What Is Verifeyed"). Verifeyed's product is a software suite that can be installed by users who obtain a license; then end users can process suspect images to evaluate if the images are manipulated or not (Verifeyed, "Verifeyed | What Is Verifeyed"). Unlike SMTDP, Verifeyed distributes their software suite by themselves and sells directly to its customers. From the technology perspective, Verifeyed focuses on image metadata analysis and ballistics analysis (a type of image compression analysis).

Both SMTDP and Verifeyed have direct and indirect customers that are from major insurance companies, journalism industries, and credit card companies. The companies belong to markets that highly value image authenticity in order to reduce the potential losses to fraud. Like SMTDP and Verifeyed, our targeted markets are the insurance, journalism, and credit card markets, though we will target the insurance market segment first.

## Comparison

Compared to our competitors' businesses and products, our image detection web service is strategically positioned to have more informative detection results and to be simpler in integration.

Our product is more informative in local manipulation changes than other products on the market. Our algorithms focus on both low level image analyses and high level image features, while our competitors mainly focus on bringing meaning out of a suspected image's metadata and compression analysis. Verifeyed's key differentiating factor is their image ballistics method which can identify the source camera of a photo using its quantization table. While this requires building a database of camera signatures, augmenting publicly available datasets is not an insurmountable challenge. We have matched some features of our competitors by providing the same low-level analysis, but we have also added more advanced forms of analysis. We look for high-level image features like Copy-Move detections and image-splicing detections. These high level features are common amongst most digital image manipulations, and are more difficult to conceal by image forgers. These extra dimensions help

our customers further in interpreting what happened to a particular image. The result of all these feature detections is a more comprehensive and informative image manipulation detection.

Our product has simple and interpretable results. Our competitors' software tools only outputs simple binary classification result indicating if an image has been manipulated or not. In contrast, our web service provides a detailed report showing visualizations of interesting manipulated features, as well as human-readable interpretations of what manipulations may have happened in the image. This detailed report is crucial to our customers since a simple but powerful tool would reduce the need to hire image analysis experts to investigate suspect images.

Screenshots of Verifeyed Professional Edition's user interface on Windows, which is not accessible through different platforms (Verifeyed).

Most importantly, our image detection software is a web service product that can be integrated into any existing workflow. SMTDP distributes its technology through third party resellers; its technology could only benefit users after integrating with third parties products (SMTDP, "SMTDP | Partners"). Verifeyed sells traditional softwares -- such that users need to purchase and install Verifeyed's software and install on their own computers (Verifeyed, "Verifeyed | What Is Verifeyed"). In contrast, our customers are not required to use a particular operating system to run a standalone application. Furthermore, the REST service's flexibility allows our product to be integrated into existing workflows helping to reduce customer training burdens and switching costs.

## Combatting New Entrants

The largest threat that our product faces comes from academia. Academia is constantly working to develop new techniques that can offer incremental accuracy improvements. There is little we can do to mitigate this problem. That being said, we believe that the added utility of new accuracy improvements from another entrant would not be worth the switching costs to another platform.

The very nature of our detection service innately combats new entrants by using machine learning. Our system will become more accurate as we collect more data from our customers. Thus, the longer our system is in use, the better our accuracy becomes. Data acquisition of good quality for this space is not easy to obtain (as we have discovered from working on this capstone project). This first-mover advantage would make it difficult for new entrants to achieve our accuracy levels using the same techniques.

## Scalability

Our software is designed to be deployed on a distributed infrastructure that can scale linearly with the rate at which we process images. We can grow and scale with the requirements of our customers.

Software-as-a-service (SaaS) products only have one supplier: cloud datacenters. The companies that offer cloud datacenters services have been relentlessly cutting prices to the point where all the providers are extremely cheap (State Farm®). If one provider tried to raise prices, we could easily shift our product offerings to a different provider since software service deployment architecture is generally the same. The cheap landscape of cloud service providers enables to host our detection service at minimal costs.

The alternative to deploying our software service on the cloud is to rent or buy our own servers. However, this would have to be at a price point in which the cloud offerings are no longer cheaper, which is possible depending on the usage of the servers (Leong 12). Even in the event where we must purchase servers, hosting and maintaining servers is a rather small fixed cost. Hence, the choice of hosting on cloud providers versus hosting on our own servers is not a decision that will constrain our profitability.

# Intellectual Property

For this capstone, we have developed new algorithms for detecting and identifying manipulated features in a digital image. Our team has taken careful consideration of the intellectual property laws of the United States to protect the future prospects of our business. Copyright and trade secret laws in the United States can help us enforce our business model and keep our competitive advantage.

## (No) Patent Strategy

Our team recognizes the difficulty in receiving a software patent claim. There are several criteria needed to acquire a patent, and the two criteria that challenge most software patents, including our own are nonobviousness and novelty. After analyzing each of these criterion, we believe that our web service, overarching detection process, and feature-specific sub-processes are not patentable.

### Non-obviousness

US patent law requires a patent claim to be non-obvious "to a person having ordinary skill in the art" (35 USC 103). Software patents can be invalidated due to the ideas being of "common sense" or "obvious to try" as a next step (Perfect Web Technologies, Inc).

We hold concerns that our detection web service as a whole is "obvious to try." Although our detection service uses a combination of different methodologies, these techniques that we've created are heavily based on academic papers and resources that are in the public domain. Using all of these different features in a machine learning classifier could be seen as the next step for performing a general image manipulation detection classification. Hence, the United States Patent and Trademark Office may invalidate our claim to a patent due to nonobviousness.

### Novelty

The other main challenge in obtaining software patents is establishing novelty. US patent laws requires that a patent claim to be filed before "the claimed invention was patented, described in a printed publication, or in public use, on sale, or otherwise available to the public" (35 USC 102).

## Existing Patents

Although we believe our processes to be non-patentable, we must make sure that we are not infringing on any patent claims. We performed a thorough search of patents relevant to our image manipulation detection service. Although we were unable to find any patents that captured the essence of our end-to-end classification system, we were able to find a few more-specific patents that pertain to the features that our classification engine uses.

Looking into patents filed into other countries, we found that Chinese patent CN102609948A had many of the similar keywords and the same ultimate objective to one of our subprocesses related to Copy-Move detection (张华熊). The United States Patent and Trademark Office will look to foreign patents to determine if there is already prior art (United States. Dept. of Commerce). However, upon a closer reading of the technical specifications of the Chinese patent, we have determined that the Chinese patent would not be seen as prior art to our processes. Even though we are using some of the same image-level characteristics to deliver an analysis like Scale-Invariant Feature Transform (abbreviated as SIFT), our process for Copy-Move detection involves technical processes that are unrelated and different to the technical processes detailed in the Chinese patent. We believe this differentiation is enough such that this claim cannot be applied to our Copy-Move detection process in the United States courts, as extreme specificity results in narrow applicability of a patent's claims (LizardTech, Inc). Hence, we do not believe we are infringing on any prior claims. In addition to the differentiated procedures of the Chinese patent, we can also disregard this foreign patent claim since it is not enforceable in the United States.

We were also able to find existing patents in the United States such as US8200034B2 by New Jersey Institute of Technology, US7439989B2 by Microsoft, and US7720288B2 by Eastman Kodak that target double JPEG compression. What is unique to these methods is not the artifacts that they are detecting, but the classification methods being used. All of the patents listed describe unique processes for detecting double compression, while the method we employ uses distinct workflows for aligned and nonaligned JPEG (Bianchi). That being said, the specifics of how our technique detects JPEG double compression are distinct from the claims made in the patent. Hence, like in the Chinese patent case, we believe that our implementation does not infringe on the patents and their ways of extracting this feature.

Even though there are several relevant patents to how we extract various types of image manipulation features, we believe our techniques are differentiated enough such that we are not infringing on any patent claims.

## Existing Publications

Having processes that do not infringe on existing patent claims does not mean that we necessarily have a novel and patentable claim. In addition to an existing patent search, the United States Patent and Trademark Office will also look to publications to determine the existence of prior art when determining whether or not to grant a patent (United States. Dept. of Commerce).

Many of our components are based on academic papers from external research groups that have heavily guided our implementations of detecting interesting image manipulation features. These features include higher order statistics, JPEG error level analysis, and double JPEG compression (Farid and Lyu) (Krawetz) (Bianchi and Piva). Although our copy-move detection technique was developed independently from knowledge that is documented by the academic publications of Irene Amerini at the University of Florence in Italy, the United States Patent and Trademark Office would still use these publications as evidence of prior art due to the substantial similarity of processes and its publication date being in 2011 (Amerini).

In addition to the prior art within the feature extraction context, prior art in publications also exists in the context of our full image detection service. A dissertation from Columbia University by Yu-Feng Hsu and Shih-Fu Chang captures the essence of using a variety of image features to perform a classification for image manipulation detection (Hsu and Chang 1). Although the dissertation use different features, Hsu and Chang describe a full system that is very similar to image detection service.

## No Patents

The prior art found in existing publications further reduce our possible claims to a patent. Furthermore, our use of our classification system with these specific features can be seen as obvious to try. We conclude that both the overall classification system and the individual feature components for our classifier are to be non-patentable.

## Trademark Law

We anticipate that other new entrants could mimic our image detection processes since the processes are not patentable. Our advantage over new entrants is the time that we have already spent in finding and tuning the optimal parameters. In order to minimize the threat of another entrant copying our systems, we will use trade secret protections to protect our "methods, techniques, processes, procedures, programs, or codes" (18 USC 1839). In other words, we will be able to protect the parameters in our machine learning process, the manipulated image datasets that test our methods, as well as the detailed lists on our customers. Protecting these assets of our business will be crucial to maintaining our competitive advantage.

## Copyright Law

The last, and most important part, of our intellectual property strategy is to enforce the copyright protections offered by United States law.

Primarily, we will use copyright to enforce our business model of selling the license to view the image forensic reports produced by our web service. United States copyright law gives us the sole right to reproduce, distribute, and display our image analysis reports (17 USC 106). Successful enforcement of copyright protections would prevent the scenario in which a middleman takes advantage of our online software service by reselling or redistributing our reports indirectly to our potential customers. Rather than selling our software or our reports, we will sell the license to view the image analysis reports that our web service produces.

United States copyright will also protect our right to create any derivations of our product (17 USC 106). This would prevent any scenario in which a third party used our image analysis reports to make a slightly more comprehensive report or redressed the report by plagiarizing our metrics and analyses. Courts in the United States would enforce our copyright based on the substantially similarity of ideas that are contained in a potentially infringing work, as well as the "look and feel" of the potentially infringing work (Data East USA, Inc).

Hence, copyright is the strongest portion of our intellectual property strategies. The enforceability of copyright protections will help our business model flourish.

## Overall Intellectual Property Strategy

All in all, intellectual property laws will be used by our capstone team to minimize threats from competing businesses and to enforce our business model. Although we are unlikely to receive any patent claims to our product, other intellectual property laws can help our business flourish. Trade secret protections will help keep our head-start in training and tuning our machine learning models. Lastly, copyright protections will help enforce that no third-party can redistribute or reproduce our image analysis reports that we will sell to our customers.

# Works Cited


17 USC. Sec. 106. 2008. *Cornell University Law School*. Web. 26 Feb. 2015.

18 USC. Sec. 1839. 1996. *Cornell University Law School*. Web. 26 Feb. 2015.

35 USC. Sec. 102. 2012. *Cornell University Law School*. Web. 2 Mar. 2015.

35 USC. Sec. 103. 2011. *Cornell University Law School*. Web. 26 Feb. 2015.

"Altered Images Prompt Photographer's Firing." *Msnbc.com*. Accessed April 15, 2015. http://www.nbcnews.com/id/13165165/ns/world_news-mideast_n_africa/t/altered-images-prompt-photographers-firing/.

Bianchi, T., and A. Piva. "Image Forgery Localization via Block-Grained Analysis of JPEG Artifacts." *IEEE Transactions on Information Forensics & Security* 7, no. 3 (June 2012): 1003–17.

"Company Overview of SMTDP Technology, LLC." *bloomberg.com*. N.p., n.d. Web. 12 Apr. 2015.

Data East USA, Inc. v. Epyx, Inc. 862 F. 2d 204. Court of Appeals, 9th Circuit. 1988. *Google Scholar*. Web. 26 Feb. 2015.

Dulaney, Ken, et al. "Predicts 2015: Mobile and Wireless" Gartner, Inc. Stamford, Connecticut. November 5, 2014. Print.

Farid, Hany, and Siwei Lyu. "Higher-Order Wavelet Statistics and Their Application to Digital Forensics." Accessed December 19, 2014.

FBI, "Insurance Fraud." *FBI*. N.p., n.d. Web. 12 Apr. 2015.

GEICO, "GEICO's Special Investigations Unit." *geico.com*. N.p., n.d. Web. 12 Apr. 2015.

GEICO, "How GEICO Investigates a Claim." *geico.com*. N.p., n.d. Web. 12 Apr. 2015.

GEICO, "Join GEICO in the Fight against Insurance Fraud by Reporting Suspicious Activity." *geico.com*. N.p., n.d. Web. 12 Apr. 2015.

"How Car Insurance Companies Investigate Accident Claims." *DMV.org*. N.p., n.d. Web. 12 Apr. 2015.

"How to Document Auto Accident Damage." *DMV.org*. N.p., n.d. Web. 12 Apr. 2015.

Hsu, Yu-Feng, and Shih-Fu Chang. "Image Tampering Detection For Forensics Applications." Columbia University, 2009. Print.

I. Amerini, L. Ballan, R. Caldelli, A. D. Bimbo, G. Serra. A SIFT-Based Forensic Method for Copy–Move Attack Detection and Transformation Recovery. Information Forensics and Security, IEEE Transactions. Volume 6, Issue 3. IEEE. 17 March 2011.

Irby, Kenneth. "L.A. Times Photographer Fired Over Altered Image," April 2, 2003.

Kahn, Sarah. "Design, Editing & Rendering Software Publishing in the US" IBISWorld. September 2014. Print. Accessed on February 17, 2015.

Kashi, Joe. "Authenticating Digital Photographs as Evidence: A Practice Approach Using JPEG Metadata." *americanbar.org*. N.p., n.d. Web. 12 Apr. 2015.

Krawetz, Neal. "A Picture's Worth… Digital Image Analysis and Forensics". Hacker Factor Solutions. Black Hat Briefings USA 2007.

Leong, Lydia. "Technology Overview for Cloud Infrastructure as a Service" Gartner, Inc. Stamford, Connecticut. June 30, 2014. Print.

LizardTech, Inc. v. Earth Resource Mapping, Inc. 424 F. 3d 1336. Court of Appeals, Federal Circuit. 2005. *Google Scholar*. Web. 26 Feb. 2015.

Lum, Jessica. "Getty Photographer Terminated Over Altered Golf Photo." *PetaPixel*. N.p., n.d. Web. 12 Apr. 2015.

McKenna, Farrell. *Newspaper Publishing in the US*. Melbourne, Australia: IBISWorld Services, 2014. *IBISWorld*. Web. 12 Apr. 2015.

National Insurance Crime Bureau, "INSURANCE FRAUD: UNDERSTANDING THE BASICS." *National Insurance Crime Bureau* n. pag. Web. 12 Apr. 2015.

Novet, Jordan. "All Rise: The Era of Legal Startups Is Now in Session | VentureBeat | Entrepreneur | by Jordan Novet." *venturebeat.com*. N.p., n.d. Web. 12 Apr. 2015.

Perfect Web Technologies, Inc. v. InfoUSA, Inc. 587 F. 3d 1324. Court of Appeals, Federal Circuit. 2009. *Google Scholar*. Web. 26 Feb. 2015.

"Picture Imperfect." *The Economist* 9 Mar. 2013. *The Economist*. Web. 12 Apr. 2015.

Porter, Michael E. "The Five Competitive Forces that Shape Strategy" Harvard Business School Publishing Corporation. 2008. Print.

State Farm®, "Reporting Fraud – State Farm®." *State Farm*. N.p., n.d. Web. 12 Apr. 2015.

"Silver Lining." *The Economist*. *The Economist*. Web. 12 Apr. 2015.

Slobogin, Christopher. "TESTILYING: POLICE PERJURY AND WHAT TO DO ABOUT IT." *University of Colorado Law Review, Inc.* 67.1037 (1996): n. pag. Print.

Smith, Aaron. "Mobile Access 2010." *Pew Research Center's Internet & American Life Project*. N.p., n.d. Web. 12 Apr. 2015.

Smith, Aaron. "Smartphone Ownership 2013." *Pew Research Center's Internet & American Life Project*. N.p., n.d. Web. 12 Apr. 2015.

SMTDP, "SMTDP | Company." N.p., n.d. Web. 12 Apr. 2015.

SMTDP, "SMTDP | Partners." N.p., n.d. Web. 12 Apr. 2015.

Thomson, Lucy L. "Mobile Devices: New Challenges for Admissibility of Electronic Evidence" The SciTech Lawyer, Volume 9, Number 3, Winter/Spring 2013. Print.

U.S. Bureau of Labor Statistics, “Claims Adjusters, Appraisers, Examiners, and Investigators : Occupational Outlook Handbook: : U.S. Bureau of Labor Statistics.” N.p., n.d. Web. 12 Apr. 2015.

United States. Dept. of Commerce. The United States Patent and Trademark Office. “Prior Art”. Chapter 0900. Section 901. *The United States Patent and Trademark Office.* Web. 26 Feb. 2015.

Verifeyed, “Verifeyed – Image, Video and PDF Forensics, Authentication, Manipulation Detection, and and Digital Security.” N.p., n.d. Web. 12 Apr. 2015.

Verifeyed, “Verifeyed – Image, Video and PDF Forensics, Authentication, Manipulation Detection, and and Digital Security | Technology & Documentation.” N.p., n.d. Web. 12 Apr. 2015.

Verifeyed, “Verifeyed – Image, Video and PDF Forensics, Authentication, Manipulation Detection, and and Digital Security | What Is Verifeyed.” N.p., n.d. Web. 12 Apr. 2015.

张华熊, 胡洁, 薛福冰, 黄海. “Manipulation detection method for copy-paste distorted photo digital photos.” Patent CN 102609948 A. 16 Apr. 2014.

# Appendix

Figure 1: The landing page for our prototype image detection service.

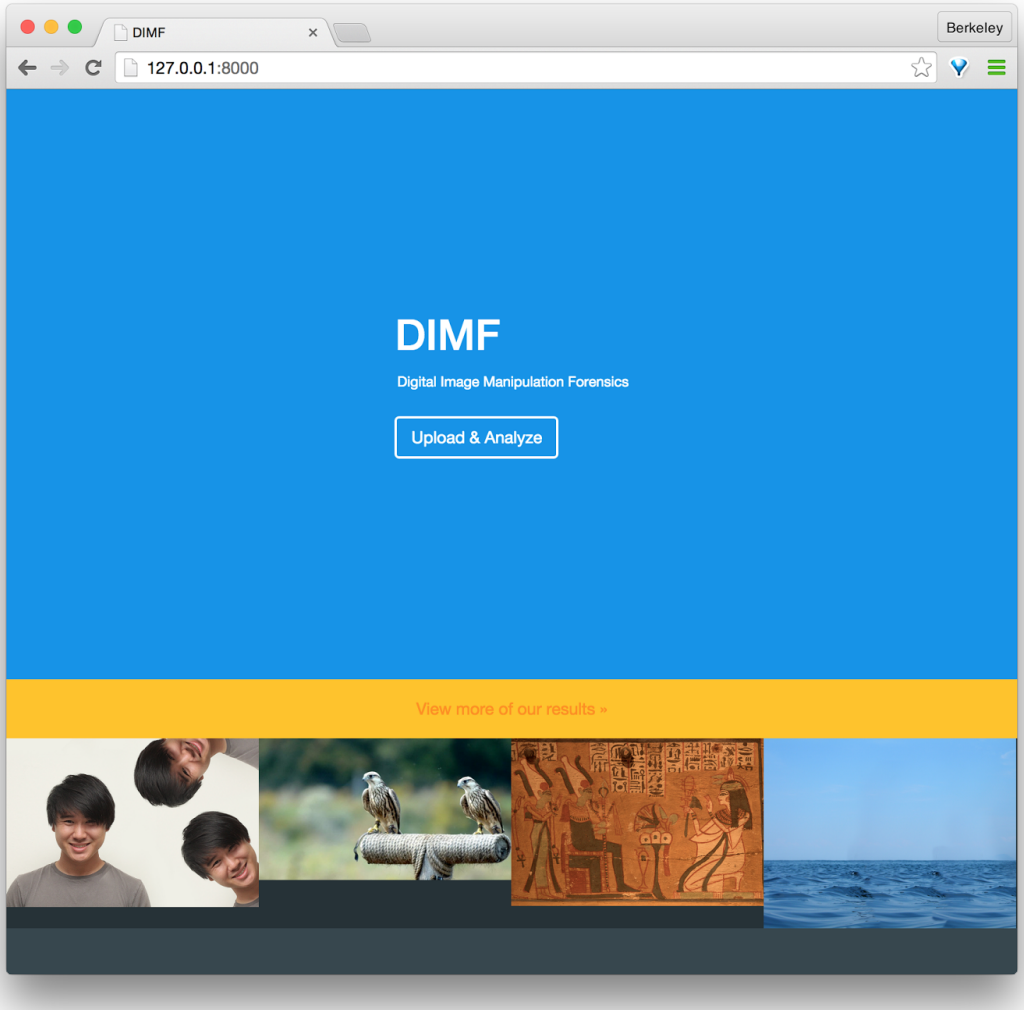

Figure 2: The progress page while a photo is being processed.

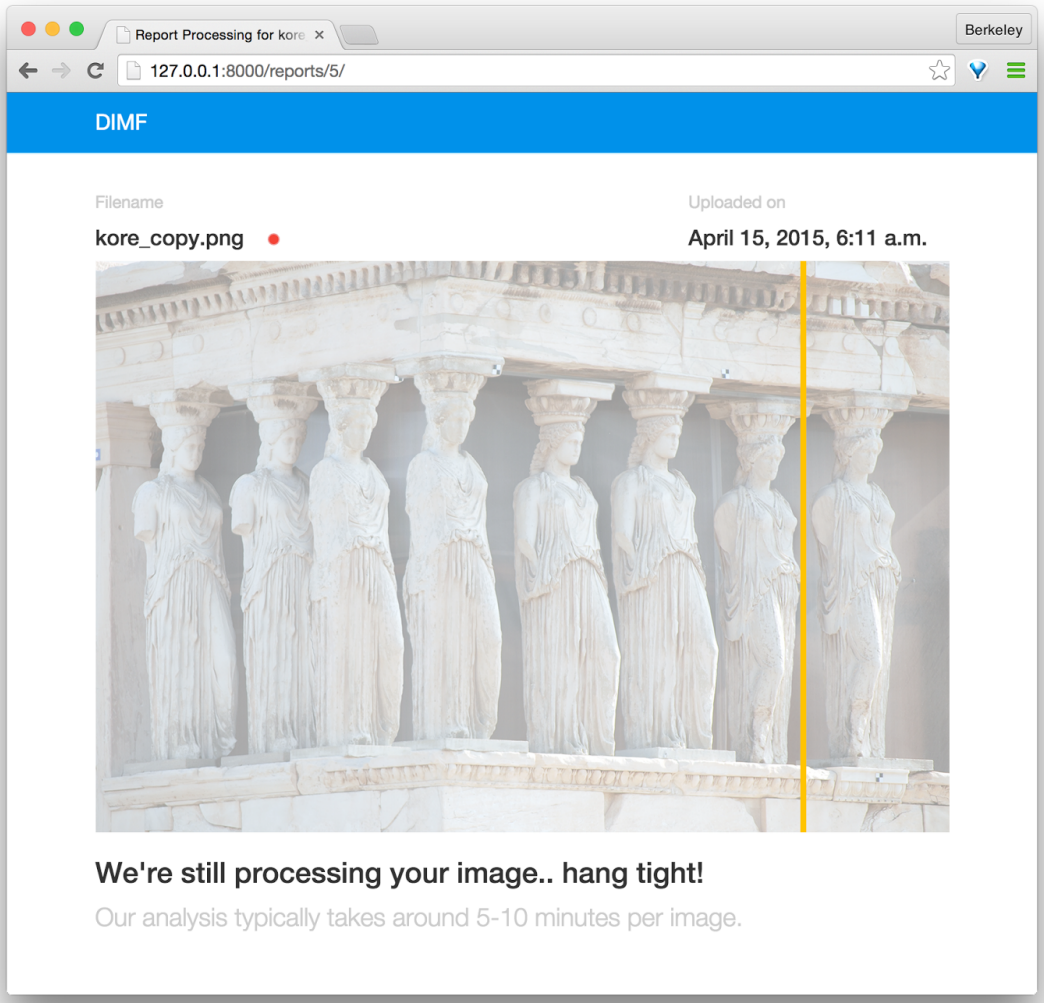

Figure 3: The results page detailing a copy-move detection and higher order statistic analysis.

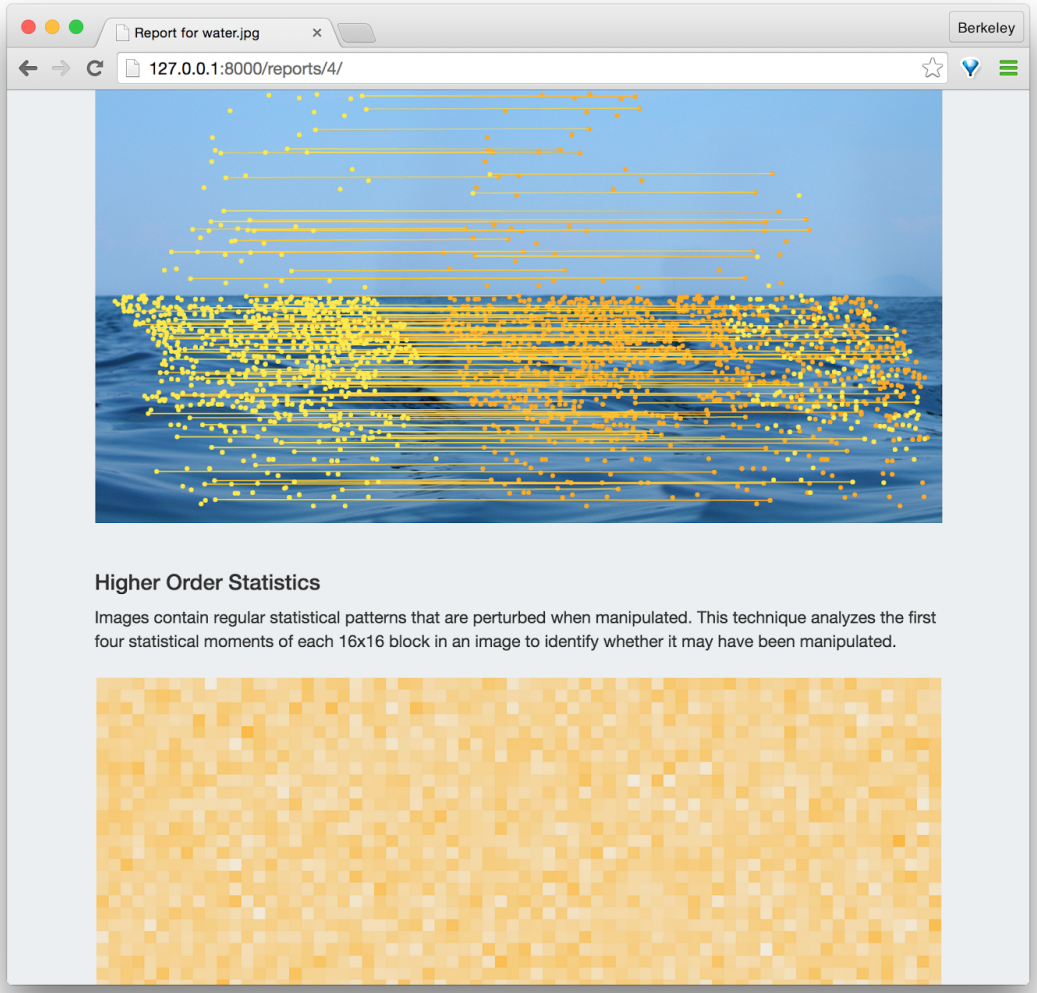

# Detection of Photomontage using Higher Order Wavelet Statistics

Omar Ramadan
Department of Computer Science
University of California, Berkeley
omar.ramadan@berkeley.edu

## I. Introduction

It is relatively simple to generate images that are indistinguishable from authentic photographs. Recent advancements in photo editing tools have made manipulated photos common occurrences in our lives, undermining the credibility of digital photographs. There is an increased demand for forensic systems that can verify the authenticity of digital photographs. For such a classification system, it is desirable that it be both passive and blind, meaning it requires no interaction to generate a response so that it can be used in batch processing, and require no prior information about the photograph such as the original. It is possible to design such systems by leveraging the artifacts that are introduced into images through photo editing.

The task of image forensics first begins with searching for evidence of tampering. Manipulations frequently leave behind artifacts that differentiate them for authentic photographs. An inconsistency in the time that a digital file was created and the time it was last modified is a particularly obvious indicator of a manipulation, though it is not completely reliable. These irregularities can exist at the metadata level such as in this case. They can also occur at a visual level with inconsistencies in shadows or lighting [1], or within the encoding and compression of the file itself [2] [3]. Forensic analysts can also detect forgeries by modeling irregularities that occur in natural images such chromatic aberrations caused by the camera lens [4] and correlations introduced in sensor arrays [5].

Any combination of these pieces of evidence can exist in the case of a tampered photograph, and the more a system is capable of detecting, the better its performance. For this capstone project, one of the core tasks is to design and implement multiple methods for detecting forgeries. In this paper, we describe a method to detect an image splicing forgery.

## II. Prior Work

### A. Image Statistics

Image splicing is defined as the construction of an image from multiple image sources. Spliced images can be some of the more misleading manipulations because they fuse together different contexts to create deception as in Figure 1.

To detect these forgeries, we exploit the use of statistical correlations contained within the wavelet representation of images. The decomposition of images into basis functions (i.e wavelets) that are localized in position, orientation and scale has proven to be extremely effective in texture synthesis, image compression, and noise removal. They have also been used for a number of forensic detection applications such as stenography, rebroadcasting detection, and classifying between computer graphics and natural images [3]. This paper serves to explore whether this method can also be used to detect forgeries in a particular image.

We utilize features derived from a recursive pyramid decomposition of images based on Quadrature Mirror Filters (QMF's) [6]. This decomposes the frequency space into multiple scales and orientations as illustrated in Figure 2. The decomposition is achieved by applying separable high and low pass filters across the axes yielding low-pass, vertical, horizontal and diagonal sub-bands. This is applied recursively to generate the pyramid.

Using this decomposition, the original work extracts features for each image the image statistical features by extracting the first four statistical moments (mean, variance, skewness, and kurtosis) of each of each orientation, at each level of the image pyramid, for each color channel for a total of 108 features.

An interesting property of the sub-bands of a QMF is that the coefficients are correlated to their spatial, orientation, and scale neighbors [7]. The magnitude of each pixel in each sub-band at level $i$ can be approximated as a linear combination of *(1-4)* four immediate spatial neighbors, *(5)* its down sampled "child" at level $i + 1$, *(6)* its counterpart in the orthogonal sub-band at level $i$, and *(7)* its down sampled orthogonal counterpart at level $i + 1$ [3]. For each sub-band, we can estimate the set of coefficients $w_1, ..., w_7$ that correspond to how much each of these spatial neighbors contributes to the magnitude of the pixel. Using this linear relation, we can find the of coefficients $w_1, ..., w_7$ for each sub-band that fits best to the data with a least squares loss.

[3] uses this interpolation to compute a regression estimate for each sub-band which we will call $Qw$, the using the actual sub-band $V$, defines a measure of deviation from this linear estimate:

$$E = log_2(V) - log_2(Qw) \tag{1}$$

From the interpolation error of each of the sub-band, we also take the mean, variance, skewness and kurtosis for an additional 108 features.

This statistical correlation amongst pixels that is captured in the sub-bands of a QMF is reminiscent of other correlations that occur naturally in authentic photographs, due to

Fig. 1: Image-splicing can be used to create convincing fakes: (a) A star ship hovering over the valley. (b) A spliced picture of the Singapore skyline.

the response of the image sensor, and applied filters in the image formation process. If such correlations exist between the sub-band, they would be perturbed with the introduction of a manipulation, resulting in higher reconstruction errors, especially along the manipulation boundaries. In this paper we explore whether such statistical signatures can be used to detect inconsistencies within a photograph.

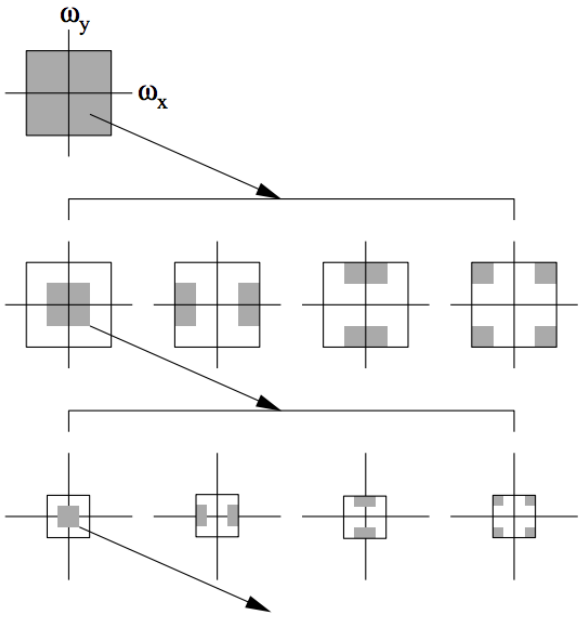


Fig. 2: From [3]: "An idealized multi-scale and orientation decomposition of frequency space. Shown, from top to bottom, are levels 0, 1, and 2, and from left to right, are the low-pass, vertical, horizontal, and diagonal sub-band."

## III. Methods

In this paper we present a method for the detection of photomontage. We reduce the problem of identifying a spliced image to the problem of detecting whether an image contains segments that are from a different source. These individual segment measurements can then be aggregated to give a classification on the image level.

### A. Image Segmentation

We segment images for a number of reasons. For such a system, the detection and classification of manipulated photos is the most important function it provides. However, localizing the area in question we provide an increased level of confidence to the classification result, and make it easier for forensic experts to validate the result or provide context to how the image was generated.

Performing analysis of segments also improves performance as opposed to image level analysis since localized analysis helps amplify manipulation signatures in the area of analysis. While the manual segmentation of the image to isolate the questionable region would yield improved results, to make the classifier passive, we need to be able to perform the segmentation automatically. We employed two methods of automatic segmentation for this analysis, using a lattice grid, and using Normalized Cuts.

### B. Normalized Cuts

One of the methods analyzed for the purpose of image segmentation is the state-of-the-art Normalized Cuts (NCuts) algorithm [8]. The NCuts algorithm is appropriate for this task in that it groups segments in the image and takes into account the total dissimilarity between different groups, as well as the similarity between groups which is very much in line with the task of identifying anomalous regions in an image. NCuts performs well in finding segments that are aligned along boundaries, which is frequently where manipulations take place. The number of cuts to perform is a hyper-parameter that must be carefully selected. It must be large enough to be able to isolate the manipulated region while too large of a number can result in the over-segmentation of consistent regions. We use 20 cuts in this particular analysis.

For each of the image segments, we compute a segment mask. We then compute the pyramid decomposition associated

with this mask. Because of the irregular shape of the segments in this pyramid, the pixels on the edges may sample points from neighboring segments.

### C. Lattice Grid

The other method of segmentation that we explore is a lattice grid of non-overlapping square windows whose size depends on the number of pyramid layers we decompose into. The smallest window size that can be selected is $2^k$ where $k$ is the number of layers in the image pyramid. The image is center cropped so that its dimensions are divisible by $8 * 2^k$, making each segment and its features only computed from its decomposition.

### D. Feature Extraction

For the $k$ levels of the QMF pyramid decomposition of each segment, we compute the first four moments corresponding to each of the three orientations. The bottom two layers from the pyramid contain 1 and 4 pixels which are too small of samples to extract statistics from, so we exclude them. This gives $(k-2)*3*3*4$ features per block. We also compute the interpolation error as given by (1) which doubles the number of features we collect.

Since at each level of the pyramid, we are leveraging the spatial relations of neighboring pixels, we are forced to discard pixels on the edges, which translates to an analysis region of the image that is $2^k$ smaller on each side.

With the lattice segmentation, our segments are closely aligned to the manipulated region, with the downside that there are more segments to analyze.

## IV. Results

The data set used is the image splicing data set from Columbia [9] consisting of 180 spliced images. The images are from four different camera sources and contains 6 sets of 30 spliced images created by splicing two different camera sources.
We divide the segments into three categories, *(1)* authentic: this is a detected segment that lies fully within the larger of the edge-masked region, *(2)* spliced: a segment that lies fully within the minority portion of the edge-mask meaning it is spliced from a secondary source, and *(3)* unaligned: a segmented region that spans across both edge-masks regions. Since the unaligned segments do not have a clearly defined ground truth, we do not use them in training.

This results in some issues with the availability of training data for the images segmented by NCuts. Since the number of strongly separable segments varies significantly across images, the resulting segments may have poor overlap with the manipulated region. See Figure 4 for an example of the automatic segmentation of a training image with NCuts. However, across the entire data set, with 20 cuts, only $10.55\%$ of the images contained segments that lied fully within a manipulated region.

To analyze the effectiveness of the extracted features, we train a Fisher's linear discriminant to project the high dimensional feature descriptor into a single dimension that best separates the segment types. The separation of the two classes is visualized in Figure 3.

We perform further evaluation of the methods using classification. The images in the data set are segmented using NCuts or as a lattice, and we randomly select a maximum of $10,000$ segments, $70\%$ of which are used for training. The results in Table I summarize the results of the automatic segmentation, and the number of segments by type that are used in the evaluation.

TABLE I: Number of segments trained on for each type of segmentation by type

| | Segment Type | Set 1 | Set 2 | Set 3 | Set 4 | Set 5 | Set 6 | All Sets |
|---|---|---|---|---|---|---|---|---|
| NCuts n=20 | Authentic | 222 | 217 | 241 | 234 | 275 | 253 | 1442 |
| | Manipulated | 51 | 55 | 55 | 55 | 18 | 27 | 231 |
| | Unaligned | 327 | 328 | 304 | 341 | 307 | 320 | 1927 |
| Blocks k=4 | Authentic | 6443 | 6270 | 6787 | 7762 | 7749 | 7445 | 7264 |
| | Manipulated | 3241 | 3390 | 2879 | 2041 | 2052 | 2330 | 2493 |
| | Unaligned | 316 | 340 | 334 | 197 | 199 | 225 | 243 |
| Blocks k=5 | Authentic | 3595 | 3395 | 3795 | 7425 | 7184 | 6557 | 7058 |
| | Manipulated | 1917 | 2051 | 1727 | 2133 | 1643 | 2228 | 2425 |
| | Unaligned | 368 | 434 | 358 | 442 | 413 | 455 | 517 |

From the results in Table II, we firstly learn that the lattice segmentation technique performs more robustly in disambiguation between two data sources as given by the results for each of the independent sets. However, NCuts outperforms lattice cuts on the full dataset. This may be explained by how a model based on the small blocks tend to capture the statistical properties introduced by the capture device, while NCuts with its segmentation along object boundaries can learn to model the irregularities that occur at manipulation boundaries.

We also learn that increasing the image pyramid depth decreased the mean Type I error by $31.7\%$ and the mean Type II error by $31.8\%$ across the individual sets, which may be attributable that each layer added increases the segments by a factor of four giving a larger sample that allows us to better approximate the true distribution parameters.

There is the difficulty of comparing the performance of these methods because of the difference in the number of training examples generated. NCuts generates a total of 20 segments per image while the blocking methods generate thousands per image so the size of the set of examples used in training, though from the same number of images.

TABLE II: Classification error

| | Error Type | Set 1 | Set 2 | Set 3 | Set 4 | Set 5 | Set 6 | All Sets |
|---|---|---|---|---|---|---|---|---|
| NCuts n=20 | Type I | 31.3 | 22.7 | 34.3 | 25.6 | 34.9 | 29.0 | 11.8 |
| | Type II | 37.5 | 23.6 | 41.1 | 62.5 | 16.7 | 22.2 | 28.6 |
| Blocks k=4 | Type I | 8.03 | 6.20 | 13.6 | 8.31 | 10.3 | 9.33 | 14.5 |
| | Type II | 16.9 | 20.1 | 32.1 | 27.1 | 17.9 | 23.3 | 30.5 |
| Blocks k=5 | Type I | 1.73 | 1.06 | 7.50 | 10.6 | 5.90 | 12.2 | 21.0 |
| | Type II | 23.4 | 8.56 | 20.1 | 9.99 | 6.08 | 22.0 | 32.3 |

## V. Discussion

### A. Image Level Classification

Smaller segments resulting from techniques such as image blocks leads to better alignment within manipulated regions, which allows the statistical properties of the sources to be captured. This particular technique seems to perform well

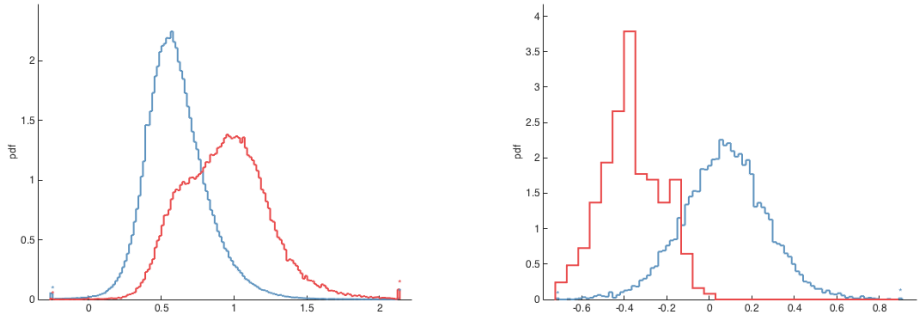


((a)) Histogram with lattice segmentation ((b)) Histogram with NCuts

Fig. 3: Separability of segment types using higher order statistic features

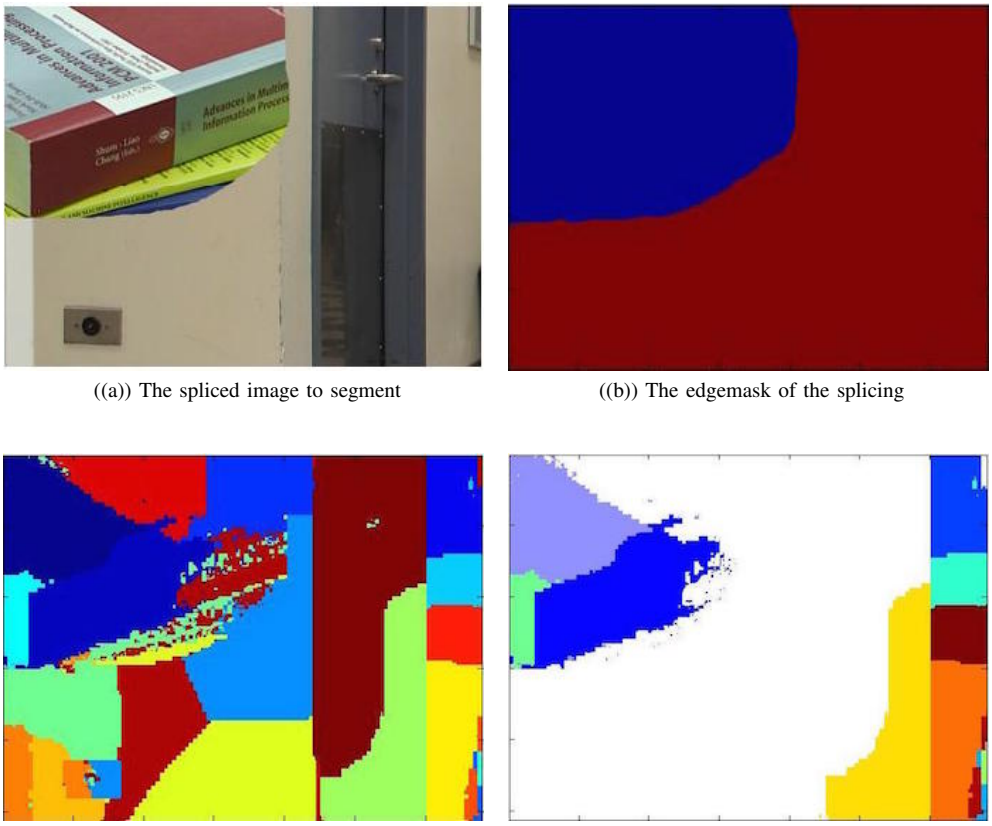

((a)) The spliced image to segment ((b)) The edgemask of the splicing

((c)) Mask output from NCuts with $N = 20$ ((d)) Unaligned segments removed from mask

Fig. 4: Results of automatic image segmentation.

for the task of source identification of segments, which is important for forensic analysis.

The next task in forensic analysis is synthesizing segment level evidence, and concluding whether a manipulation occurred. The task of classification allows us to aggregate our evidence from our various methods to provide a confidence score. In the case of image forensics, the localization of the manipulation is equally as important as to provide a means of confirmation and auditing.

One method of synthesis would follow from the logical OR of all the segment level classifications. If $p$ is the probability of a segment being incorrectly classified and there are $n$ segments, the probability that an image is incorrectly classified is $1 - (1 - p)^n$ assuming a model in which the segment level classifications are independent. There is an exponential decrease in performance with the number of segments that are generated Figure 5.

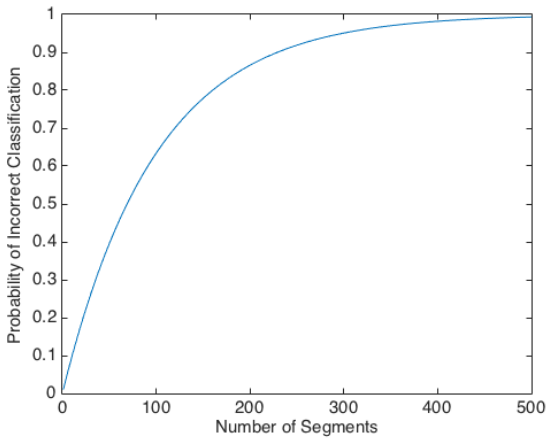


Fig. 5: Even with 99% block level accuracy, the image level classification error rate is unusable for images with a large number of segments. For a 1024x768 pixel image, there are 3,072 16x16 segments.

## VI. Future Work

### A. Markov Fields for Image Level Analysis

The assumption that each segment is independently classified is rather strong and fails to account for the spatial dependencies that exist in an image. Markov random fields can model these dependencies, and discriminative random fields have proven to be effective in image analysis [10] [11]. Using a model such as discriminative random fields would allow to collate and smooth segment level features, as well as integrate evidence from multiple observation domains.

With such a framework in place, we can integrate the various analysis techniques implemented in this capstone project into a single model that leverages observations in multiple domains for segments, and uses spatial consistency from these different techniques to boost image level classification performance. This would be a clear next step.

### B. Image Segmentation

A lattice type image segmentation, even in combination with better modeling assumptions is still suboptimal. Most image splicing manipulations would involve introducing a foreign object into the image. Manipulations that alter the background information tend to be less common, and furthermore less interesting. An interesting follow up would be to use object detection algorithms such as Multiscale Combinatorial Grouping (MCG) [12] to form segment proposals in an image that could be analyzed. Precise localization of the manipulated region would only serve to amplify statistical discrepancies, and improve detection. It would also dramatically reduce the number of image segments to be analyzed which would lead to more accurate image level predictions.

### C. Color Domain

Given that most statistical analysis is performed on the luminance channel, it would be interesting to repeat this analysis in the YCrCb domain. This is the domain in which the sensor records data so we can more accurately measure statistical irregularities introduced by manipulation rather than measuring the statistical correlations between the colors given their computation from the YCrCb channels.

## References

[1] E. Kee, J. F. O'Brien, and H. Farid, "Exposing photo manipulation from shading and shadows," vol. V, pp. :1–21, to be presented at SIGGRAPH 2014. [Online]. Available: http://graphics.berkeley.edu/papers/Kee-EPM-2014-XX/

[2] C. Chen, Y. Q. Shi, and W. Su, "A machine learning based scheme for double JPEG compression detection," in *2008 19th International Conference on Pattern Recognition*. IEEE, Dec. 2008, pp. 1–4. [Online]. Available: http://ieeexplore.ieee.org/articleDetails.jsp?arnumber=4761645

[3] H. Farid and S. Lyu, "Higher-order wavelet statistics and their application to digital forensics." [Online]. Available: http://doi.ieeecomputersociety.org/10.1109/CVPRW.2003.10093

[4] M. K. Johnson and H. Farid, "Exposing digital forgeries through chromatic aberration," in *Proceedings of the 8th workshop on Multimedia and security*. ACM, pp. 48–55. [Online]. Available: http://dl.acm.org/citation.cfm?id=1161376

[5] A. C. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," vol. 53, no. 10, pp. 3948–3959. [Online]. Available: www.cs.dartmouth.edu/farid/publications/sp05a.html

[6] E. P. Simoncelli and E. H. Adelson, *Subband transforms*. Springer. [Online]. Available: http://link.springer.com/chapter/10.1007/978-1-4757-2119-5_4

[7] R. W. Buccigrossi and E. P. Simoncelli, "Image compression via joint statistical characterization in the wavelet domain," *IEEE Trans Image Proc*, vol. 8, no. 12, pp. 1688–1701, Dec 1999.

[8] J. Shi and J. Malik, "Normalized cuts and image segmentation," vol. 22, no. 8, pp. 888–905. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=868688

[9] Y.-F. Hsu and S.-F. Chang, "Detecting image splicing using geometry invariants and camera characteristics consistency," in *International Conference on Multimedia and Expo*.

[10] Y.-F. Hsu and S.-F. Chan, "STATISTICAL FUSION OF MULTIPLE CUES FOR IMAGE TAMPERING DETECTION." [Online]. Available: http://www.ee.columbia.edu/ln/dvmm/publications/08/asilomar2008_hsu.pdf

[11] S. Kumar and M. Hebert, "Discriminative Fields for Modeling Spatial Dependencies in Natural Images." [Online]. Available: http://www.cs.cmu.edu/~skumar/modDRF.pdf

[12] J. Pont-Tuset, P. Arbelaez, J. T. Barron, F. Marques, and J. Malik, "Multiscale Combinatorial Grouping for Image Segmentation and Object Proposal Generation," Mar. 2015. [Online]. Available: http://arxiv.org/abs/1503.00848

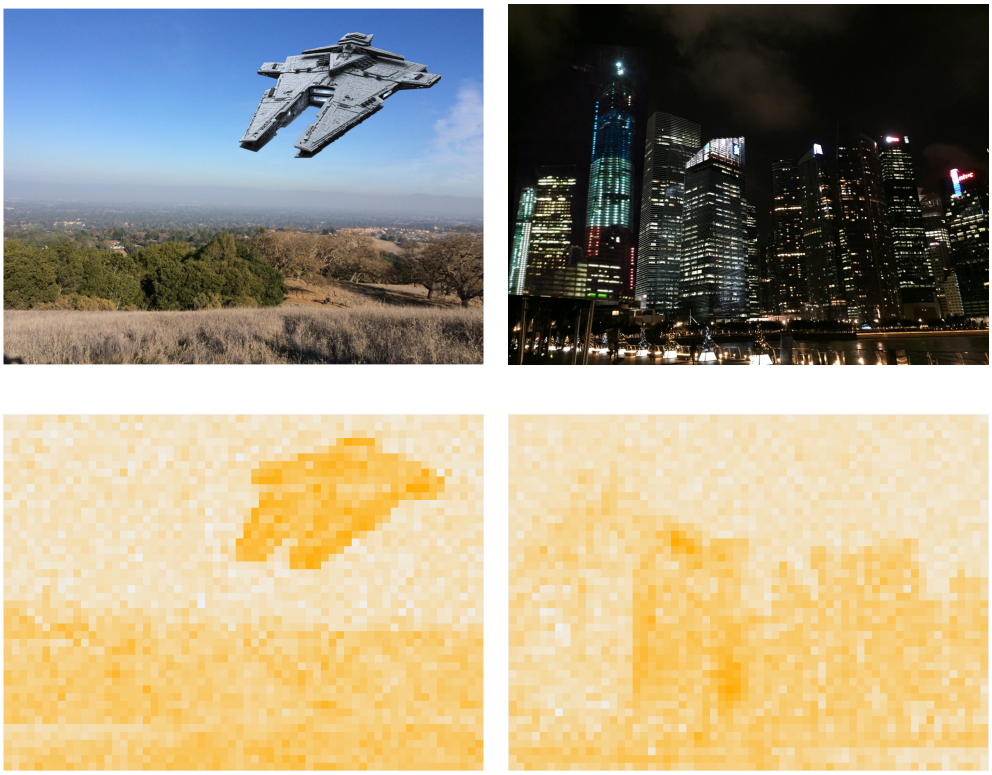

((c)) Clear discrimination of the spliced region ((d)) Unclear discrimination of spliced region

Fig. 6: Results of image level analysis.

# Concluding Reflections

In conclusion, we succeeded in delivering a system that could automatically classify manipulated images. However, the detection capacity of our system did not meet the expectations that we had envisioned.

I am quite pleased with the outcome of the higher order statistics and how well it worked given its use as in image splicing detection has not been explored. However, we were unable to bring a working implementation of a fusion classification system to life as we originally intended due to time constraints. Building this detector required implementing multiple classification techniques, but instead of building out multiple methods, we focused on iterating and the three classification systems we initially sought. Perhaps the scope of the project was too large.

Image forensic analysis is a niche speciality that has a steep learning curve. From a project management perspective, I learned that to efficiently engineer and implement methods requires a comprehensive understanding of the subject material, and a very well defined scope. In retrospect, it would have been more advisable for our first project deliverable to be a literature review to narrow our scope, rather than a preliminary implementation of the methodologies on which we would iterate and improve. It would delay the output of technical results, however, it would increase the quality of our end result.

If someone were to pick up from where our team left off, they would firstly need to familiarize themselves with the relevant literature. I believe the references in our technical contribution, as well as our literature reviews would be instructive in doing so. Our code is well documented and easy to contribute to.