

# Forensic Analysis of Memetic Image Propagation: Introducing the SMOC BRISQUeT Method

**James A. Hodges**

The University of Texas at Austin  
School of Information, USA  
james.hodges@utexas.edu

**Mitch Chalet**

Harvard Kennedy School  
Shorenstein Center, USA  
mitchchalet@hks.harvard.edu

**Praful Gupta**

The University of Texas at Austin  
Cockrell School of Engineering,  
USA  
praful\_gupta@utexas.edu

## ABSTRACT

This paper introduces a mixed-methods approach for forensically reconstructing the propagation of visual media via networked digital devices. The authors present case studies drawn from political misinformation around the January 6, 2021 riots at the U.S. Capitol. Using interpretive analysis, the authors identify traces of user interfaces that remain in images being shared about the riots. Using computational analysis, the authors evaluate compression levels in digital photographs of the events in question, thus identifying which instances of the image are closer to the source (as well as which images appear to be identical). By combining these two approaches, the authors argue that SMOC BRISQUeT refines our understanding of misinformation’s memetic spread—helpful in curbing future abuses as well as in guiding the production of more effective cross-platform spread when desired.

## KEYWORDS

Social media, misinformation, digital forensics, image compression, memes.

## INTRODUCTION

This study proposes a two-step process for understanding the memetic spread of images via social media networks. The affordances of many social networking platforms facilitate more effective and widespread sharing of images when compared against video, text, or other media formats. Images are typically sourced from one post on a social platform, downloaded to a user’s device, and uploaded elsewhere online. Furthermore, in the wake of increased platform moderation around sensitive topics, controversial content is frequently spread via screenshots, archive links, and other novel methods in order to evade automatic filtering and moderation processes (Acker and Chalet, 2020). This paper demonstrates two separate methods for understanding the memetic spread of imagery, before demonstrating the combined use of both approaches. The authors begin with interpretive human identification of diegetic elements in a digital image in order to retrace a general sense of its path across multiple users and platforms. Next, the authors apply a newly trained blind, no-reference image-quality analysis algorithm in order to generate a quantified score of image characteristics that can identify the level of JPEG compression across a corpus of images.

By combining human interpretation with computational analysis of compression, the authors generate a more robust account of the images’ specific genealogical histories than either technique alone. This richer account of memetic copying and distribution offers insight into the specific vectors of spread used to propagate controversial content and evade moderation on social platforms. Differences between digitally generated images and “natural” photographs make certain images more amenable to interpretive and computational analysis, respectively. The techniques discussed in this paper can be used to forensically reconstruct the spread of damaging content online, and the combination of human and machine-driven analysis ensures that conclusions are sensitive to both cultural specificity and imperceptible visual changes, respectively.

## Background

Every time an image is uploaded to a social platform, it is compressed using some form of algorithm in order to reduce its required bandwidth and streamline user experience (Joshi and Sarode 2020). At the same time, user behaviors such as screen capture can produce new visual features related to the on-screen elements present on their device, including user interface elements like captions, commentary, or GUI elements. Therefore, each time an image is propagated from one place to another, it will undergo a subtle change in both its visual and technical composition. The paper presents a methodology for understanding both of these elements in media distribution via social platforms, using a combination of machine learning and human interpretation.

First, we develop a linear visualization of memetic spread using interpretive analysis of diegetic interface elements in a method that we name SMOC (Sourcing Memetic Online Content). We assess the spread of content using memetic propagation analysis (MPA) techniques in order to produce a linear, graphical memetic analysis (LiGMA) to visualize the spread of our examples over time. Next, we assess the visual quality of imagery and changes in

---

*84<sup>th</sup> Annual Meeting of the Association for Information Science & Technology | Oct. 29 – Nov. 3, 2021 | Salt Lake City, UT. Author(s) retain copyright, but ASIS&T receives an exclusive publication license.*

visual compression between discrete copies of memetic images using a specially trained version of No-Reference (NR) Image Quality Assessment (IQA) enabled by Mittal, et al's (2011; 2012) Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE), which we have retrained for our purposes and named BRISQUEt (an acronym for "Blind/Referenceless Image Spatial Quality Evaluator: reTrained"). Based on the study, the authors conclude by making estimates concerning the specific avenues and extent of spread afforded to different copies of key sample images.

RQ1: How can the human interpretation of diegetic image components be used to understand the spread of memetic image content online, and in what contexts is such an approach most productive?

RQ2: How can computational analysis of image compression be used to interpret the spread of memetic image content online, and in what contexts is such an approach most productive?

### **Related Work**

The ability to copy and distribute digital media is often seen as one of its key attributes. Significant existing research within media and cultural studies offers useful definitions and critical interpretations that can enrich technical research (Parikka 2007, Davison 2012, Shifman 2014, Stein et al 2014, Gaboury 2015). Computer and information science has tackled these issues by automating various forms of content identification, evaluation, and classification, with a particular focus on identifying content that is false or misleading (Farid 2006, Gupta et al 2013, Ma et al 2013, Conroy et al 2015, Rubin 2016). A more recently emerging and still-developing third thread of research also exists, which combines the aforementioned humanistic and computational approaches to understanding memetic distribution (Teyssou et al 2017, Acker and Chalet 2020, Singh et al 2020).

The present study joins recent scholarship in the small but growing field of literature that combines technical and humanistic methods. In particular, we demonstrate the value of interpretive analysis that focuses on the presence of user interface elements from a user's personal networked device in images and video (thus attesting to certain details about that user's personal information practices), as well as the computational identification of image compression (which attests roughly to the number of times an image has been re-uploaded and compressed by online social networks). By combining these threads of research, we advance a mixed-methods approach to understanding media sharing online and work toward developing a richer, more well-balanced methodological toolkit for media and information research to employ when conducting their work.

Humanistic media studies scholars have developed robust definitions and critical approaches to digital media, in particular exploring the subtle effects of certain definitional decisions. For example, while the terms "meme" and "viral" are often used interchangeably to describe the spread of digital media, scholars like Shifman (2014) point out that these terms come from different origins and possess different meanings. Viral media was originally used to refer specifically to media that is, or is thought to be, self-replicating within host machines in the sense of a biological virus (Parikka 2007). The notion of a meme, on the other hand, emerges from evolutionary studies emphasizing the role of social activity in transmitting cultural habits (Davison 2012). In this study, we prefer the notion of memetic spread, which builds on the study of memes, as well as its etymological roots in terms like mimesis and copying.

Existing information and computer science research addressing memetic spread has addressed the detection of doctored content online in considerable depth. Early in the 21st century, Farid (2006) began developing forensic tools to detect alterations and distortions in photography (2006). Gupta et al (2013) demonstrates a different approach, using metadata drawn from specific social networks (such as user details and timestamping) in order to identify fake images. Conroy et al (2015) build further on this approach by using network analysis combined with computational linguistics. Additional research using linguistic identification, such as that by Rubin et al (2016), has worked to identify misleading content that is primarily textual in nature. In the realm of video, Teyssou et al (2017) have also developed new software that can identify falsification in video, expanding upon previous research addressing images and text. Teyssou et al also highlight an important characteristic of research problems in media verification, noting "the problem of online information verification is very complex and touches upon a number of research fields, including media studies and journalism (p. 23). The present study builds on Teyssou's multi-pronged approach by drawing on both technical and humanistic approaches.

Singh et al (2020) provide another example of multi-pronged analysis of online media, by automating both visual and textual analysis. Yet where multi-pronged approaches like those of Singh et al and Teyssou et al still primarily aim to automate the detection of misleading media, the present study focuses on understanding memetic spread, rather than detecting veracity. By demonstrating a methodological approach that can help researchers forensically reconstruct the information behaviors of users that spread memetic content, this study enriches scholarly knowledge about the ethos of sharing and specific techniques employed when images are shared across multiple platforms and networks.

## CASE STUDIES

### Case Study 1: Interpretive Analysis Using the SMOC Method (Sourcing Memetic Online Content)

In order to demonstrate interpretive analysis of memetic content, this case study examines a series of social media posts (image and text) collected and preserved during the storming of the U.S. Capitol Building on January 6, 2021. After collecting the corpus, two particular screenshots posted to Twitter are analyzed for the visual characteristics that identify unique patterns of cross-platform propagation from the same original post. By identifying the specific social networks whose user’s handles and interface elements are displayed in the images, as well as evidence of the technical modalities used to reproduce the images, we are able to draw qualitative conclusions about the nature of sharing and copying among proponents of a particular memetic object. We term these visual elements “diegetic user interface elements,” using a notion of “diegesis” that is drawn from cinema studies to refer to visual markings that are contained within the frame of a selected media object. In this case, the user interface elements are *diegetic* to the social media post because they are included in the post contents, rather than being superimposed around the post as they would be when viewing an object through the social network’s user interface.

Data collected for this portion of the research was manually curated into a GitHub repository comprising screenshots of the posts in question, copies of the images included within posts, and a record of metadata for all collected tweets (Berk and Chalet 2021). For the sake of demonstrating interpretive analysis, a tweet by the Twitter user “StunningTruthSeekr” (@stunningtruthsr) was selected for further analysis (see Figure 1), alongside a separate tweet by Twitter user “RC0076” (@rwacollins). The posts in question are exemplary of the layered quality that visual elements often accumulate during the process of circulation. Like the rings of a tree, interface elements can stack up next to one another over time, with each instance attesting to a hierarchical unit of time—in this case, one act of sharing. These qualities are most clearly visible in the two images attached to the post in question. This dataset includes both screenshots and post URLs. While screenshots are generally useful for interoperability between contexts, post URLs and archived versions thereof provide richer contextual information. For more information, see Nelson (2021).

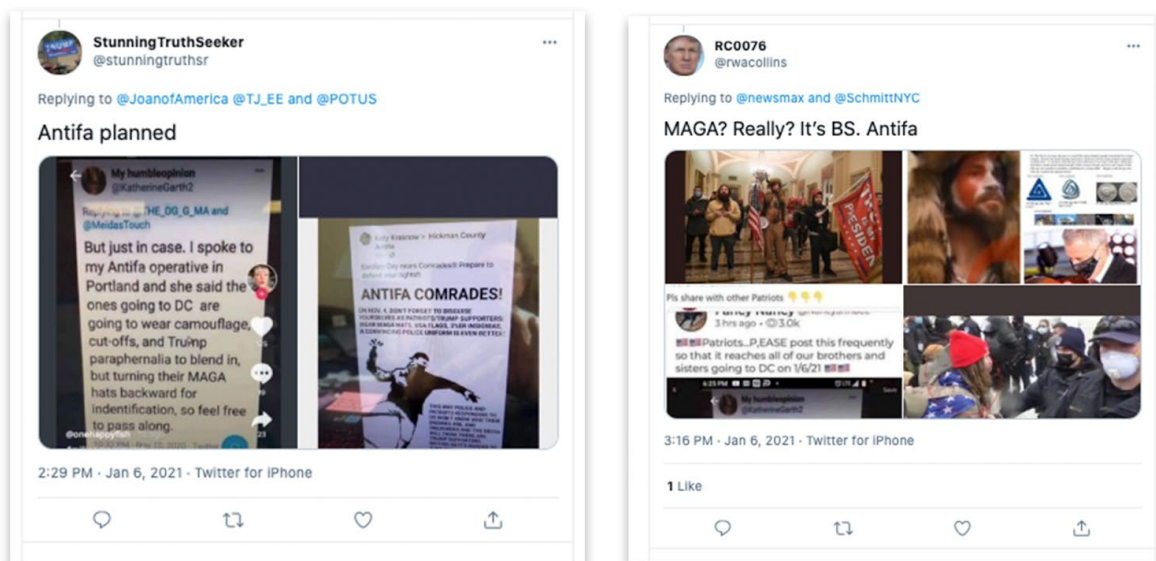
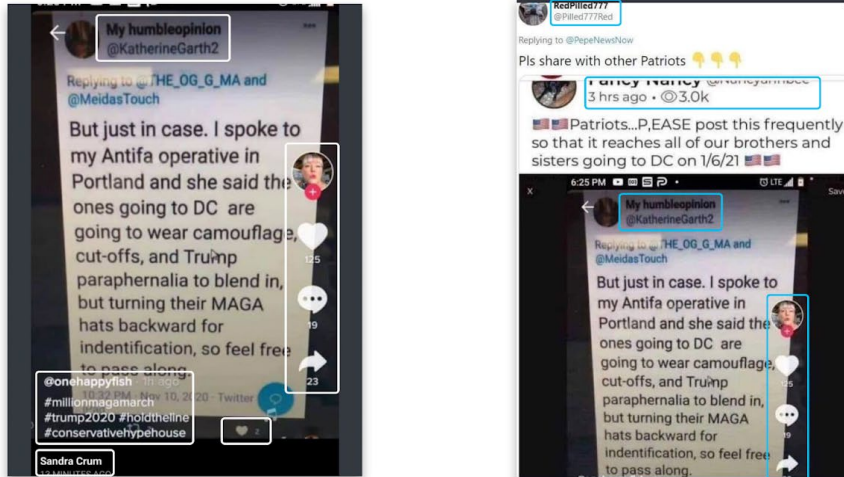


Figure 1. Tweets containing memetic content

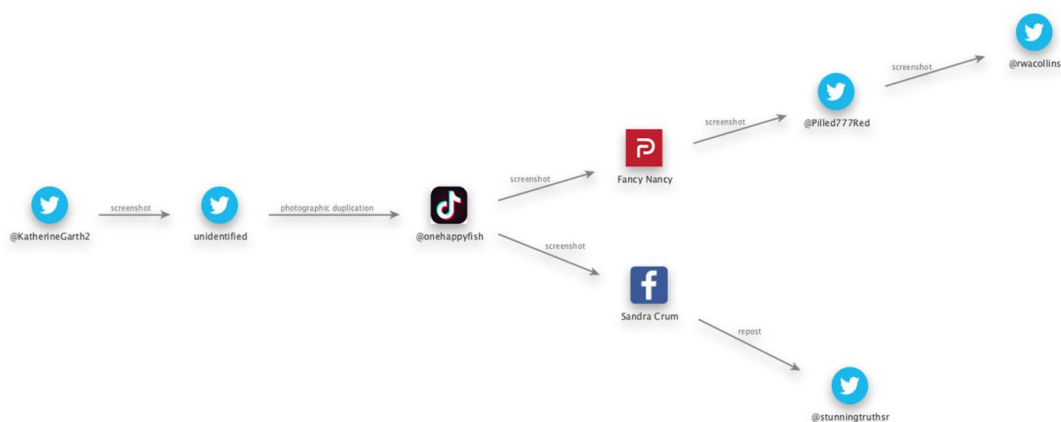
In the first Twitter post in Figure 1, StunningTruthSeekr includes two popular images that suggest members of left-wing “Antifa” (short for “anti-fascist”) groups may have planned to participate in the January 6 actions while disguised as members of right-wing “Patriot” or “MAGA” groups loyal to then-president Donald J. Trump. The second image included depicts a Twitter post by RC0076 (@rwacollins), which includes as its third attached image (bottom left) a slightly varied copy of the purported “Antifa operative” screenshot. The narrative that left-wing agitators were responsible for the unrest on January 6, 2021 was widely deployed in right-wing media as a way of deflecting associations with the insurrectionists that stormed the U.S. Capitol (Grynbaum et al 2021). By examining the specific techniques and paths used to spread this narrative through memetic means, an interpretive analysis of the post’s contents, we note the cross-platform spread, as well as the coexistence of both screen capture and photographic duplication (taking a photo of one’s screen)—two particular information behaviors that complicate attempts at automating detection of misleading content in methods established by existing research.



**Figure 2. Detail of image included with Tweets** Left image depicts original attachment to StunningTruthSeekr tweet. Right image depicts further spread of this image via screenshots, Parler, and additional Twitter reposts. Diegetic interface elements contained in memetic images are highlighted.

The first image included in StunningTruthSeekr’s Tweet (Figure 2) shows evidence of at least three acts of cross-platform sharing—first on Twitter, then on TikTok, and finally on Facebook. This finding is informed by the main content, which contains the word “Twitter” and a characteristic timestamp at bottom of the white screenshot containing the image’s main contents. It is also supported by the avatar, heart icon, comment icon, and share arrow present along the image’s right edge, all of which are characteristic of TikTok’s user interface. Additional evidence of an intermediary appearance in Twitter can be faintly identified by the presence of a twitter user interface (with black background) just beneath the main image’s white field. Finally, the black bar at bottom of the image reading “Sandra Crum” and “13 minutes ago” are consistent with the mobile interface for Facebook. Thus before StunningTruthSeekr reposted this image via Twitter, it had already been propagated across different social networks at least three times, and shared at least four times overall. Also present in StunningTruthSeekr’s first image is a skewed perspective and reflection of artificial lighting, suggesting that the image was captured by photographing the user’s screen, rather than downloading or screen capture. The second image in Figure 2 depicts another vector of transmission for the “Antifa operative” post initially spread by StunningTruthSeekr. In this image, the same post appears to have been screenshotted and re-shared in the Parler app, before being shared to Twitter by RedPilled777 (@Pilled777Red) (see Figure 2 for detail of diegetic elements). “Vernacular” information practices such as these screenshotting behaviors complicate efforts at rationalizing and predicting user behaviors and are therefore challenging to automated systems (Gaboury 2015), and therefore make the identification of diegetic evidence a useful step in expanding researchers’ kit of methods and tools for disinformation detection efforts.

The second image of StunningTruthSeekr’s post also contains diegetic interface elements and evidence of screen photography. The image appears to display a Facebook post on the “Hickman County Antifa” Facebook page by “Katie Krasnow,” while the image’s skewed perspective and surrounding desktop environment suggest that it was downloaded to a user’s Macintosh desktop before being photographed for additional sharing. The white stripes on either side of the image suggest that this photographed image was subsequently recorded via screen capture over a white background. The black stripes above and below suggest the image was screen shot once again afterwards, while the grey stripe at bottom indicates that this final act of screen capture was performed on an Apple iPhone X or later, after the manufacturer introduced the “home bar” as a replacement for its earlier “home button” designs (Patel 2017). These diegetic interface elements suggest that the original Facebook post was photographed at least once and screenshotted at least twice. They also suggest that the two images in StunningTruthSeekr’s post arrived through different histories of sharing and mimesis.



**Figure 3. Linear graphical memetic analysis (visualizing spread of imagery in question)**

By enumerating the particular acts and techniques of sharing used to spread this particular narrative, this case study shows that human users can interpret interface elements and qualitative properties of an image that may be skewed in perspective, obscured by lighting effects, related to specific interfaces and interface changes, and subject to inconsistently colored or oriented arrangement. Furthermore, by comparing findings using this method on multiple instances of the same memetic image, we demonstrate a method for tracking multiple branches of the memetic spread. This approach is most effective in identifying features and pathways for further exploration, whether automated or performed by humans (RQ1). At the same time, this approach has clear limitations. It does not track all sharing, but rather only detects sharing actions that produce some form of visible alteration to the image. For more subtle variations, computational analysis is warranted. In ideal cases, both approaches would be applied together in order to most effectively identify the paths used in distributing a particular unit of misinformation. Like a forensic crime scene investigation, reconstructing the specific timeline of events during analysis may help to identify the particular parties, environments, and tools used in carrying out the acts in question. These findings are generated using human perception and knowledge of the rapidly changing sociotechnical operating environments within which memetic content spreads. Thus this portion of study enables a reconstruction of activities that may not be readily identified using automated methods.

### **Case Study 2: Computational Analysis Using BRISQUeT (Blind/Referenceless Image Spatial Quality Evaluator: reTrained)**

While digital-born imagery such as screenshots often carry diegetic markers of their propagation histories, measuring the spread of natural imagery must be analyzed separately. We analyze the propagation of natural images by measuring the compression artifacts present in natural imagery (a form of digital decay) by applying a newly trained blind, no-reference image-quality analysis algorithm in order to generate a quantified score of image characteristics that can identify the level of JPEG compression across a corpus of images. These methods help to establish provenance, hierarchical genealogies, and fingerprinting, and expand upon simple perceptual hashing with rough estimations of spread. Such findings are not readily perceived by unassisted human users, and thus the machine-driven method we propose can help to enrich interpretive analyses with additional sources of data.

By measuring the visual quality of natural imagery, subtle changes in compression associated with reposting can be quantified, allowing for a visual fingerprinting method which functions in a similar manner to a perceptual hash, while allowing for a generalized hierarchy to be assumed. This data can be interpreted as a “genealogical” record of how close images are to the point of origin. Assuming that imagery gets more compressed as it gets shared, changes in visual compression can be used in combination with other metadata, such as timestamps, to determine patterns of sourcing and propagation.

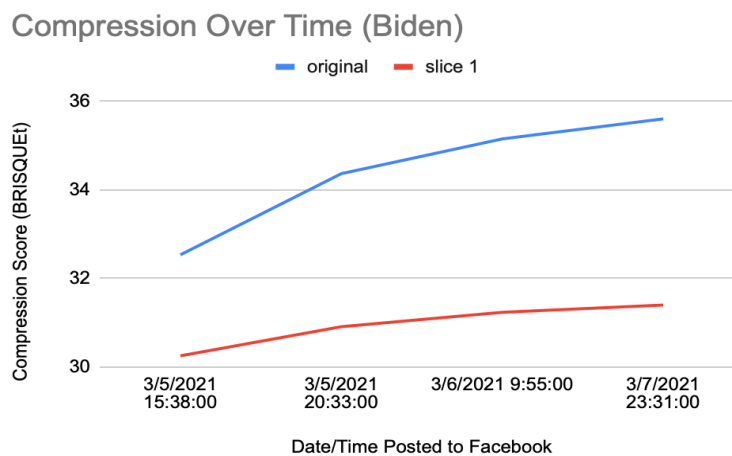
The spread of natural imagery begins with its inception: the time, date, and location, and device with which it was captured. The original image serves as a ground truth for quality. In order to assess the spread of a particular piece of content, a memetic propagation analysis can be created, which maps the various locations discrete copies of the same image have been posted. A linear, graphical memetic analysis encapsulates the spread of an image over time while associating it with other metadata in order to assume a genealogy of spread.

BRISQUeT is a “retrained” version of the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) tuned specifically for JPEG Compression. The authors selected 300 images sourced from a Facebook group which were exemplary of memetic content and features and compressed the collection of imagery at 35 discrete, increasing



levels of JPEG compression. Authors then repurposed existing MATLAB code for BRISQUE and produced BRISQUeT by retraining the algorithm on extracted features from the above dataset of compressed imagery. The initial BRISQUE algorithm was trained on the LIVE IQA Database, which featured 779 distorted images across five discrete distortions—JPEG and JPEG2000 compression, additive white Gaussian noise, Gaussian blur and a Rayleigh fast fading channel distortion. These natural scene statistics are used to compute a score which correlates with human opinion scores rating the quality of natural imagery. Next, a mapping is “learned from feature space to quality scores using a regression module, yielding a measure of image quality. The framework is generic enough to allow for the use of any regressor” (Mittal et al 2012, p. 4701). To train BRISQUeT, the authors compiled a database of 300 images sourced from an online conspiracy community. This initial dataset was then compressed at 30 separate JPEG compression levels, comprising 900 images total—30 copies of each image, each compressed at a different level. By limiting BRISQUeT to one distortion, JPEG Compression, and removing the other four, the remapped output score measures visual quality solely as a function of the amount of compression within an image. BRISQUeT is capable of measuring imperceptible differences in compression between multiple copies of the exact same image, calculating visual fingerprints for each instance as an image spreads on social media. A hierarchy can be assumed based on the increase in compression score. The output score is representative of BRISQUE’s original human opinion score mappings—the output score for BRISQUeT has been arbitrarily mapped based on the new features, so the output score functions similarly to the original BRISQUE scores while the actual output values are discrete to BRISQUeT. When all images selected for comparison are discrete copies of the same image, with no changes in resolution or aspect ratio, the original files can be compared directly using BRISQUeT. In other cases, changes in aspect ratio, user edits, diegetic interface elements, and other visual phenomenon can skew the results of BRISQUeT. When a corpus of imagery incorporates differences in aspect ratios, resolutions, etc., we align the images using automated scripts for common image manipulation software, so that a representative sample is sliced from each image where they overlap (Walter 2021). This ensures that images can be accurately compared. Next, we developed a GUI in MATLAB for the model, which allows users to easily open a single image for analysis, and returns a visual compression score. The code and dataset are now readily available and open-source (Hodges et al 2021).

A frame from a YouTube video filmed during the Biden presidency (occurring at the 01:22 mark) serves as an exemplary “ground truth” selection of determining a genealogical hierarchy from multiple discrete copies of the same image (Weekly Conversation 2021). This frame was selected as an ingredient asset within a new work, an image macro unrelated to the original video, and then spread memetically across many publicly accessible social networks (see Figure 4). We analyze four separate posts from March 5-7, 2021 on Facebook, each containing a discrete copy of the image. Since each of these images are discrete copies, we can directly compare both the original images and representative slices thereof. When the individual BRISQUeT scores are mapped to the timestamps posted, the image appears to “age” or “decay” as compression increases over time, leading to subtle decreases in visual quality. Determining a definite relationship as to the sourcing and reuploading of these examples requires further interpretive work, such as interviewing the page owners. However, using BRISQUeT a hierarchy can still be identified between images that are “upstream” (closer to source) and those that are “downstream” (further from source). From these findings further interpretive work can determine which copies came “before” or “after” the others.



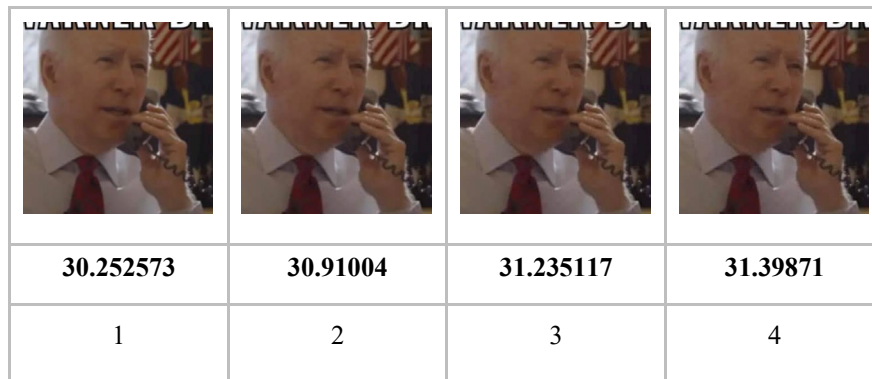


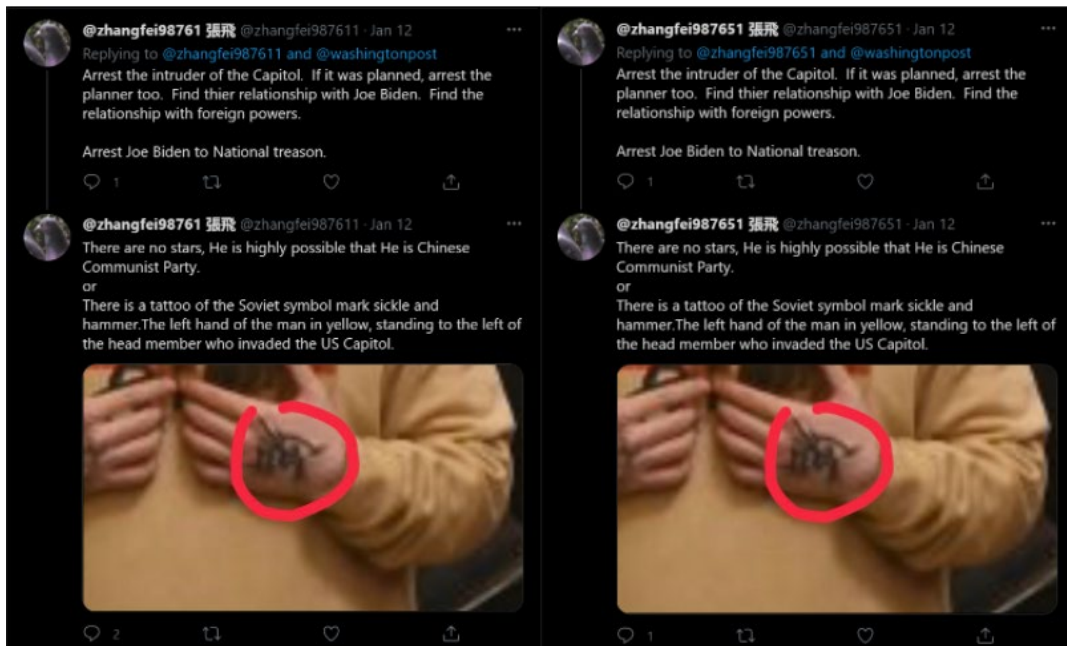
Figure 4. Example of digital compression that grows over time

Index	Source	File size	Original	Slice 1	Timestamp
1	<a href="#">MCxK2</a>	38 KB	32.53254	30.252573	3/5/2021 15:38:00
2	<a href="#">TheKinoplex</a>	30.4 KB	34.36434	30.91004	3/5/2021 20:33:00
3	<a href="#">WheresFrankOceanAsWell</a>	31.9 KB	35.14824	31.235117	3/6/2021 9:55:00
4	<a href="#">IndecentAsAllHell</a>	31.4KB	35.60043	31.39871	3/7/2021 23:31:00

Table 1. Sources of Biden image

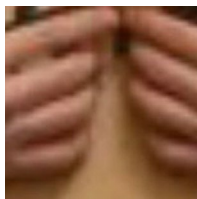

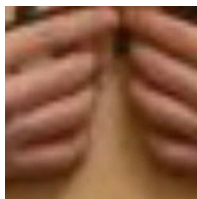

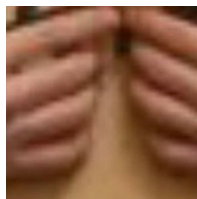
Index	Source URL	Handle	Archive URL
1	<a href="https://www.facebook.com/MCxK2/photos/a.100696408010497/464898614923606/">https://www.facebook.com/MCxK2/photos/a.100696408010497/464898614923606/</a>	<a href="#">MCxK2</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/avAM75Bg1Pi94Ojm18KuM1bQ/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/avAM75Bg1Pi94Ojm18KuM1bQ/</a>
2	<a href="https://www.facebook.com/TheKinoplex/photos/a.104157184829866/205288104716773">https://www.facebook.com/TheKinoplex/photos/a.104157184829866/205288104716773</a>	<a href="#">TheKinoplex</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/A0fiQC3YiMvCgdVlk9PqbWjj/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/A0fiQC3YiMvCgdVlk9PqbWjj/</a>
3	<a href="https://www.facebook.com/WheresFrankOceanAsWell/posts/285096592980532">https://www.facebook.com/WheresFrankOceanAsWell/posts/285096592980532</a>	<a href="#">WheresFrankOceanAsWell</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/fBIyMxaPM0jvrScc3uFUntVO/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/fBIyMxaPM0jvrScc3uFUntVO/</a>
4	<a href="https://www.facebook.com/IndecentAsAllHell/posts/280265963455011">https://www.facebook.com/IndecentAsAllHell/posts/280265963455011</a>	<a href="#">IndecentAsAllHell</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/dYtfoV8WPqBHIDkR2a0EO2Wb/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/dYtfoV8WPqBHIDkR2a0EO2Wb/</a>

Table 2. Source URLs and archived URLs for Biden image



**Figure 5. Selection of tweets and images examined with BRISQUet**

In addition to identifying hierarchies, BRISQUet can also be used to fingerprint discrete copies or slices of a corpus of imagery. This method can be used to determine content sources programmatically using visual quality alone. We detail an example of coordinated inauthentic behavior (CIB), where two different Twitter accounts, @zhangfei987611 and @zhangfei987651, post the exact same image across 5 separate posts. The similarities in these discrete handles along with the BRISQUet fingerprinting strongly suggests that these two accounts are managed from a coordinated source, and the image posted across both accounts was sourced from the same content repository. Thus we conclude that BRISQUet specifically, and computational visual analysis more generally, assists in reconstructing the memetic spread of natural imagery in ways not easily quantified by human viewers (RQ2). Whereas the SMOC method identifies born-digital characteristics like diegetic interface elements, computational analysis with BRISQUet works best in comparing natural photographic imagery.

				
<a href="#">zhangfei987611</a>	<a href="#">zhangfei987611</a>	<a href="#">zhangfei987611</a>	<a href="#">zhangfei987651</a>	<a href="#">zhangfei987651</a>
24.03049825	24.03049825	24.03049825	24.03049825	24.03049825
7	8	10	11	12

**Table 3. Comparison of image details (appearance and BRISQUet score)**



index	handle	file size	original	slice 1	slice 2	slice 3
7	<a href="#">zhangfei987611</a>	31 KB	30.89832572	28.23423216	27.2201	24.03049825
8	<a href="#">zhangfei987611</a>	31 KB	30.89832572	28.23423216	27.2201	24.03049825
10	<a href="#">zhangfei987611</a>	31 KB	30.89832572	28.23423216	27.2201	24.03049825
11	<a href="#">zhangfei987651</a>	31 KB	30.89832572	28.23423216	27.2201	24.03049825
12	<a href="#">zhangfei987651</a>	31 KB	30.89832572	28.23423216	27.2201	24.03049825

**Table 4. Complete BRSQUeT scores for hand tattoo image**

index	source	archive
7	<a href="https://twitter.com/zhangfei987611/status/1349683852357222402">https://twitter.com/zhangfei987611/status/1349683852357222402</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/EpraGkXKP9hxURJi5vwtiJKY/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/EpraGkXKP9hxURJi5vwtiJKY/</a>
8	<a href="https://twitter.com/zhangfei987611/status/1349314369990000641">https://twitter.com/zhangfei987611/status/1349314369990000641</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/IRwe8YD4Af8DiUDDXiEY9IqF/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/IRwe8YD4Af8DiUDDXiEY9IqF/</a>
10	<a href="https://twitter.com/zhangfei987611/status/1349111386194919424">https://twitter.com/zhangfei987611/status/1349111386194919424</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/rx0g4wA0vLoDmQLC9QPU1hxE/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/rx0g4wA0vLoDmQLC9QPU1hxE/</a>
11	<a href="https://twitter.com/zhangfei987651/status/1349014813901529091">https://twitter.com/zhangfei987651/status/1349014813901529091</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/f7UbOfqAOTSr7SILcWh4mpkC/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/f7UbOfqAOTSr7SILcWh4mpkC/</a>
12	<a href="https://twitter.com/zhangfei987651/status/1348934799763795970">https://twitter.com/zhangfei987651/status/1348934799763795970</a>	<a href="https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/C8ultIhT767gh3NTHgTqYiKd/">https://websatchel.com/j/pages/O83q89UjiWDjZvQOcOhMyTyk/C8ultIhT767gh3NTHgTqYiKd/</a>

**Table 5. Post URLs and archived URLs for Tweets examined with BRISQUeT**

## CONCLUSION

This study has shown a combination of methods for reconstructing the memetic spread of media content online. It has done so by highlighting the roles of specific sharing modalities related to the computing and social networking platforms that users employ in the sharing process. Practices like screen photography, cross-posting on multiple networks, capturing screenshots, and uploading in new compressed media formats all leave forensic traces of user activity in the media objects. For visual characteristics related to photography and cross-platform spread, an interpretive approach can effectively enumerate the layers of evidence created by specific sharing actions. For subtler visual characteristics related to image compression, a computational analysis can help to establish which instances of a file are further “upstream” (closer to source), and which are further “downstream” (further away from source). By combining these methods, the SMOc BRISQUeT model can assist in reconstructing the spread of imagery that is both natural and digital in origin. By combining the human and computational methods demonstrated in this paper, researchers can develop sophisticated accounts of memetic spread online, sensitive to both the subtleties of cultural practices as well as those of computational objects.

While this approach is limited by its inability to ascertain exact numbers of sharing actions and/or image re-uploads, it nevertheless provides an enriched sense for the techniques and networks used in circulating any given image. Thus while existing and current research effectively addresses the question of identifying misleading information online, our approach instead shows a method for understanding the spread of misleading information once it has been identified. Future research concerning the memetic spread of online content can be enriched by employing the methods in this paper insofar as these methods work to reveal the interrelated human and computational factors that determine information spread. Rather than favoring either cultural or computational analysis alone, this paper shows the potential for combining both.

## ACKNOWLEDGMENTS

The authors thank Dr. Alan Bovik and the UT Austin Laboratory for Image and Video Engineering for valuable input concerning the image quality evaluation algorithms used in this project. Thanks also to Emily Berk for her archival work and research support, and Jansen Derr for finalizing the software release. The authors also thank all reviewers, as well as publications and support personnel affiliated with ASIS&T during the production and review of this project.

## REFERENCES

- Acker, A., & Chalet, M. (2020). The weaponization of web archives: Data craft and COVID-19 publics. *Harvard Kennedy School Misinformation Review*. <https://doi.org/10.37016/mr-2020-41>
- Berk, E., & Chalet, M. (2021). Capitol-Riots [GitHub Repository]. <https://github.com/memeticinfluence/capitol-riots> (Original work published 2021)
- Conroy, N. K., Rubin, V. L., & Chen, Y. (2015). Automatic deception detection: Methods for finding fake news. *Proceedings of the Association for Information Science and Technology*, 52(1), 1–4. <https://doi.org/10.1002/pra2.2015.145052010082>
- Davison, P. (2012). The Language of Internet Memes. In M. Mandiberg (Ed.), *The Social Media Reader* (pp. 120–134). New York University Press.
- Farid, H. (2006). Digital doctoring: How to tell the real from the fake. *Significance*, 3(4), 162–166. <https://doi.org/10.1111/j.1740-9713.2006.00197.x>
- Gaboury, J. (2015, February 6). *On Vernacular Computing*. Art Papers. [https://web.archive.org/web/20150206051446/https://www.artpapers.org/feature\\_articles/2015\\_0102-feature3.html](https://web.archive.org/web/20150206051446/https://www.artpapers.org/feature_articles/2015_0102-feature3.html)
- Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking Sandy: Characterizing and identifying fake images on Twitter during Hurricane Sandy. *WWW '13 Companion: Proceedings of the 22nd International World Wide Web Conference*, 729–736. <https://doi.org/10.1145/2487788.2488033>
- Hodges, J., Chalet, M., & Gupta, P. (2021). SMOC-BRISQUEt [GitHub Repository]. <https://github.com/mitchalet/SMOC-BRISQUEt> (Original work published 2021)
- Jin, Z., Cao, J., Zhang, Y., Zhou, J., & Tian, Q. (2017). Novel Visual and Statistical Image Features for Microblogs News Verification. *IEEE Transactions on Multimedia*, 19(3), 598–608. <https://doi.org/10.1109/TMM.2016.2617078>
- Mittal, A., Moorthy, A. K., & Bovik, A. C. (2012). No-Reference Image Quality Assessment in the Spatial Domain. *IEEE Transactions on Image Processing*, 21(12), 4695–4708. <https://doi.org/10.1109/TIP.2012.2214050>
- Mittal, A., Moorthy, A. K., & Bovik, A. C. (2011). Blind/Referenceless Image Spatial Quality Evaluator. *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, 723–727. <https://doi.org/10.1109/ACSSC.2011.6190099>
- Nelson, M. (2021). Uncertainty in replaying archived Twitter pages. <https://www.slideshare.net/phonedude/uncertainty-in-replaying-archived-twitter-pages>
- Rubin, V., Conroy, N., Chen, Y., & Cornwell, S. (2016, June 17). Fake News or Truth? Using Satirical Cues to Detect Potentially Misleading News. *Proceedings of NAACL-HLT 2016*. Association for Computational Linguistics, San Diego, CA. <https://doi.org/10.18653/v1/W16-0802>
- Singh, V. K., Ghosh, I., & Sonagara, D. (2020). Detecting fake news stories via multimodal analysis. *Journal of the Association for Information Science and Technology*, 72(1), 3–17. <https://doi.org/10.1002/asi.24359>
- Teyssou, D., Leung, J.M., Apostolidis, E., Apostolidis, K., Papadopoulos, S., Zampoglou, M., Papadopoulou, O., & Mezaris, V. (2017). *The InVID Plug-in: Web Video Verification on the Browser*. 23–30. <https://doi.org/10.1145/3132384.3132387>
- The White House. (2021, February 6). A Weekly Conversation: On the Line With Michele. <https://www.youtube.com/watch?v=52gVNZY-r1U>