

Social Computing for Verifying Social Media Content in Breaking News

Stuart E. Middleton, Symeon Papadopoulos and Yiannis Kompatsiaris

Abstract (max 150)

Social media is the place to go for both journalists and the general public when news events break, offering a real-time source of eyewitness images and videos through platforms like YouTube, Instagram, and Periscope. Yet, the value of such content as a means of documenting and disseminating breaking news is compromised by the increasing amount of content misuse and false claims in social media. To this end, cost-effective social computing solutions for user-generated content verification are crucial for retaining the value and trust in social media for breaking news.

Keywords

Social Media; Multimedia; Multimedia Forensics; Natural Language Processing; Information Extraction; Fact Checking; Geoparsing; Verification; News

ACM Taxonomy

H.5.1 Multimedia Information Systems; I.2.7 Natural Language Processing; I.2.6.g Machine learning

Most people have a smartphone in their pocket today, so eyewitnesses experiencing an event like a terror attack will often post real-time claims, such as the numbers dead or injured in a location, to [Twitter](#) or [Facebook](#). Eyewitness images and videos will also be uploaded to sites like [YouTube](#) and [Instagram](#), or even streamed live to sites like [Periscope](#). For events such as the Paris shootings 2015 [1] the first eyewitness videos of the various shootings were posted within 5-10 minutes of the event happening. This was followed about 20-30 minutes later with verified news reports from sources such as [Le Figaro](#), [BBC](#) and [CNN](#). In other cases, verifying eyewitness or user-generated media and claims can take much longer, from hours to even days, as for instance in the case of the Malaysia Airlines Flight 17 shot down on 17 July 2014. In many cases as soon as a breaking news event starts trending on Twitter, it is accompanied by considerable amounts of false claims and content misuse [2]. This involves the use of multimedia for misinforming the public and misrepresenting people, organizations and events. Misuse practices range from publishing content that has been digitally tampered using photo-editing software to falsely associating content with an unfolding event. Figure 1 illustrates three recent real-world examples of content misuse that quickly reached wide audiences, while the survey in [2] contains an extensive discussion on the problem of rumour detection in social media.




Figure 1: Examples of social multimedia misuse: a) digitally tampered photograph of IAF F-16 deploying a single flare over Southern Lebanon; the flare was digitally duplicated to make it appear that several missiles were being fired; source: [Wikipedia](#) b) picture of a storm descending in New York, purported to be captured during Hurricane Sandy, while it was featured in [Wall Street Journal](#) one year earlier; source: [Twitter](#) c) image claiming to depict the solar eclipse in 2012, while it was in fact digital artwork originally published on [DevianArt](#). Source: [Twitter](#)

Given the grave societal and economic impact of having misused content and false claims featured in mainstream news, it becomes extremely important for news organizations to be able to verify eyewitness media in very short time. To this end, journalists are turning to social computing approaches to automatically analyse and verify [3] user generated content (UGC) in real time. The eventual hope is that cost-effective social computing can reduce the time spent on verification to timescales nearer to real time.

Social Multimedia Forensics and Supervised Verification

Methods from the field of digital forensics are often used for assessing the veracity of multimedia items (images/videos) posted online. Some methods focus on the analysis of information encoded in the metadata of multimedia content, such as Exif information which is often associated with JPG and TIFF images. Several digital manipulations, such as the use of photo-editing software, leave traces in the form of metadata unless special care is taken to remove them and analysis of these traces can detect manipulations. Unfortunately, several of the most popular social media platforms, including Facebook and Twitter, automatically remove much of the metadata from posted content, rendering metadata-based methods useless for content obtained from these platforms.

Other forensics-based methods aim at uncovering traces of manipulation in the visual content itself. In images such methods [4] can detect cases of *splicing* and *copy-move* operations, for example inpainting of a part of one image into a second, or replication of a part of an image within the same image. Methods can leverage the uniqueness of noise patterns introduced by the capturing device in order to detect whether an image contains traces from another image captured by a different device. Other methods focus on patterns associated with the colour filter arrays of modern image capturing equipment. Splice detection methods exploit traces left by the JPEG compression process, working on the basis that the splicing of two different images and the subsequent recompression will leave detectable traces in the final JPEG file.

While all of the above methods yield satisfactory results when applied on well-controlled test samples, they have been found to exhibit poor performance in real cases [5]. One of the reasons that state-of-the-art methods fail to detect manipulations from media content published on the web is the fact that such content is often the result of numerous intermediate operations, including resizing, cropping and recompression, which has an obfuscating effect on the traces of digital manipulation. An example is Twitter and Facebook, both of which automatically resize and recompress all images uploaded to them [6]. Recent work in the [FP7 REVEAL project](#), see Figure 2, addresses the poor performance of individual tampering detection methods [7] by generating tampering probability heat maps based on a number of complementary forensics analysis algorithms. The inclusion of multiple image forensics algorithms and side-by-side comparisons gives a powerful means to journalists to understand where possible digital tampering has occurred. The problem of identifying digital manipulations on video content is even more challenging compared to the case of images, and it is further exacerbated in cases where such content is sourced from video sharing and social networking platforms such as YouTube and Facebook. The [H2020 InVID project](#) is looking into resilient approaches for video forensics building upon the [TUNGSTONE](#) commercial forensics engine.



Figure 2: Digital forensics platform for image verification. Source: [Media Verification Assistant](#)

Some cases of content misuse are not detectable by using forensics analysis, for example when an image from a past event is reposted as being associated with an unfolding event. These require other methods that attempt to detect misuse by analysing contextual cues from social media sources [8]. A typical approach adopted by such methods is to extract a variety of trust-oriented features from social media posts, and the accounts generating these posts, and to use them for training machine learning models based on historically labelled cases of fake and real posts. Such methods have been shown to yield very high fake post detection accuracy [9].

Eyewitness Media and Fact Extraction

The standard workflow for automated fact checking [10] involves monitoring of data sources, fact identification, fact extraction and fact checking. The challenges for social computing mostly involve fact identification and extraction. Once a fact is extracted it can be either checked manually or automatically against databases from sites such as [PolitiFact](#), [FactCheck.org](#), [Snopes](#) and [Wikipedia](#).

Factual claims come in many forms. The most important for social computing are factual assertions, contextual statements associated with a fact, and contextual statements associated with the trustworthiness of the fact. Factual assertions themselves can be true, false, half-truths or exaggerations. Contextual statements can allow a true representation of a fact or misrepresent it by suggestions of a false location, actor or timestamp. Contextual text can also introduce ideological cues and loaded language to bias the interpretation of the fact. Lastly contextual statements can suggest

trustworthiness, such as attribution to a trusted source or claims of previous verification which themselves might be subject to falsehoods or deliberate bias.

Fact identification approaches, especially for news related sources, try to classify sentences into non-factual, unimportant factual and 'check worthy' factual statements [11] so they can be filtered prior to fact extraction. Fact extraction is a type of information extraction (IE) problem that runs alongside information extraction techniques for concepts such as event, topic, location and time. In the past approaches such as argumentative zoning [12] were applied successfully to extract factual statements on well-structured and trustworthy scientific documents. However, the text in web and social media sources is often neither well-structured nor trustworthy so new approaches are being explored.

Early work in this area focused on verb phrase patterns (e.g. 'was elected to') to extract facts via systems such as OLLIE [13]. These used parts of speech (POS) tagging, dependency parsing and distant supervision coupled with seed attributes and bootstrapping to provide unsupervised fact extraction. In particular they were able to capture the 'long tail' of factual statements which is very important for the contextual interpretation (e.g. 'Putin made a deal with the separatists'). Later advances [14], motivated by the need to answer queries in search engines, added noun phrase patterns (e.g. 'Obama's wife') very successfully. Typically, such approaches exploit large databases of attribute names and noun phrases such as [FreeBase](#) and [DBpedia](#).

Automated fact checkers either use domain specific databases (e.g. PolitiFact) or web-scale datasets (e.g. DBpedia). Recently there has been a trend for real-time crowd-sourcing of fact checking during events such as US political rallies, with the Trump/Clinton presidential debates being the latest example. Fake news sites have also been increasing in number and can easily mislead readers [15] into trusting misinformation based on a credible but false source attribution. The iCheck system [16] is a good example where domain specific heuristics extract fact types which are visualized via a crowd-sourcing interface for users to check claims and up or down vote them.

Work from the REVEAL project [17] has taken these ideas one step further to support journalists verify breaking news. Automated fact extraction using semantic grammars, seeded with linguistic phrases originating from journalists is used to extract evidence from social media content about news events such as incident reports, facts about damage and numbers dead/injured and mentioned users, locations and attributed sources. User generated content from the scene of a breaking event, not yet syndicated via news organizations, is particularly important for journalists. Supervised learning algorithms are employed within REVEAL to identify and filter posts containing eyewitness images and videos. This type of social computing is coupled with real-time visualizations, see Figure 3, allowing journalist to quickly find contextual content such as original mentions of claims for subsequent verification.

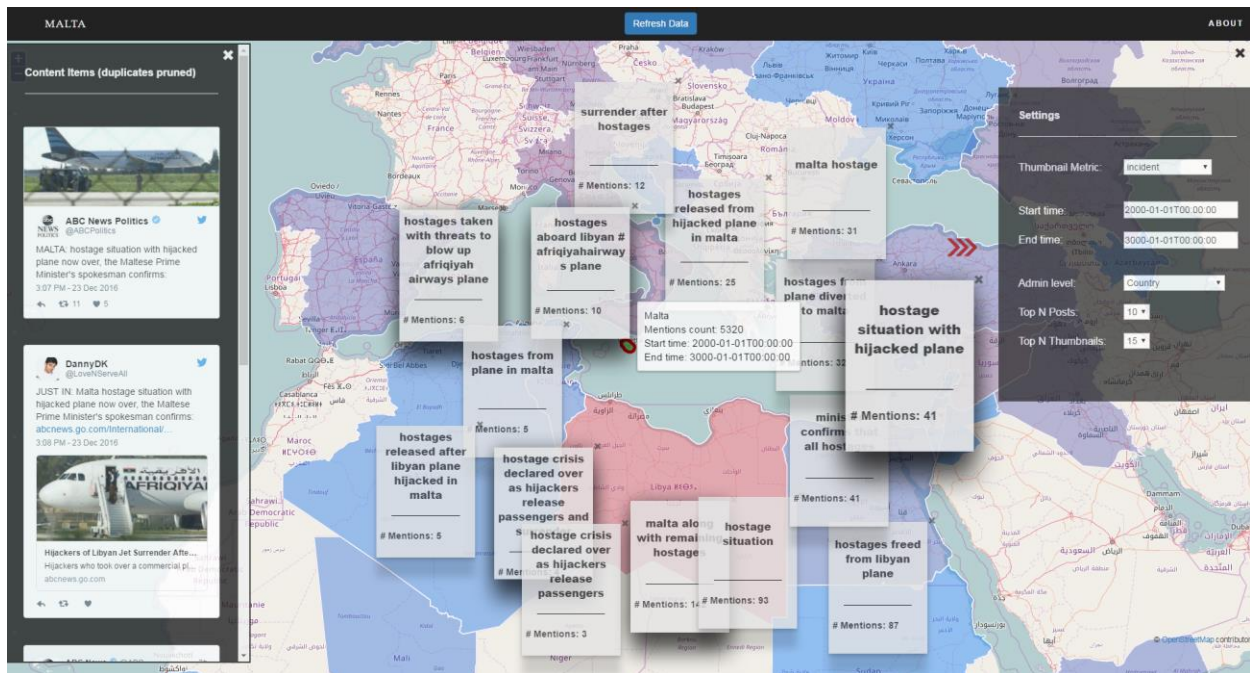


Figure 3: Interactive real-time visualization mapping extracted facts and eyewitness media in posts about the Malta plane hijacking Dec 2016. Source: [Journalist Decision Support System](#)

Considerations and Outlook

We highlighted in this article the potential of employing social computing approaches for speeding up the task of verifying user-contributed information and content sourced from social media platforms. The problem is complex and calls for a variety of approaches, each targeting different challenges stemming from the characteristics of user-generated content including high volume, inconsistent quality and a lack of provenance information. Multimedia forensics targets the actual content of multimedia. Supervised verification is best suited to cases where contextual features can be extracted, and labelled training sets of fake and real examples are available. Fact extraction and visualization approaches target text-based sources that contain references to different elements of an event, such as people, times and locations. The REVEAL project is one of the first efforts to bring together those technologies under a single platform that could provide comprehensive verification support to professional end users; details on the successful user evaluation of pilot prototypes can be found at [18]. However there is still a long way to go before such tools are widely used by newsrooms and journalists day to day. Figure 4 provides an overview of the projects and datasets useful to researchers interested in automating verification tasks for social media and news-related content.

One key challenge involved in delivering such an integrated solution is the lack of an appropriate Human-Computer Interaction (HCI) approach that would empower end users (e.g. journalists) to make

optimal use of the technologies described above. Given the extensive use of algorithms, an effective HCI approach would need to build the trust of end users by providing intuitive control and clear explanation of the results. Ultimately, end users are in charge of the whole process, and will make the final decision with respect to whether a piece of user-generated content should be considered authentic or not. Moreover, support for collaborative work among teams of journalists is another key social computing challenge that is missing from existing news provider in-house solutions, which instead employ general-purpose communication and messaging platforms such as [Slack](#) and [WhatsApp](#).

In conclusion, the problem of real-time verification of user-generated content is expected to remain unsolved in the near future, but marked improvements have already been achieved on individual parts of the verification process thanks to social computing approaches incorporating intelligent information processing. In the future, we anticipate considerable progress on this problem by incorporating the latest advances from deep learning, for instance by employing Generative Adversarial Networks [19] to build highly accurate and robust models for visually distinguishing between tampered and non-tampered regions in multimedia content, and novel HCI approaches focusing on the explainability of automatically generated results and the collaborative aspects of the verification process.

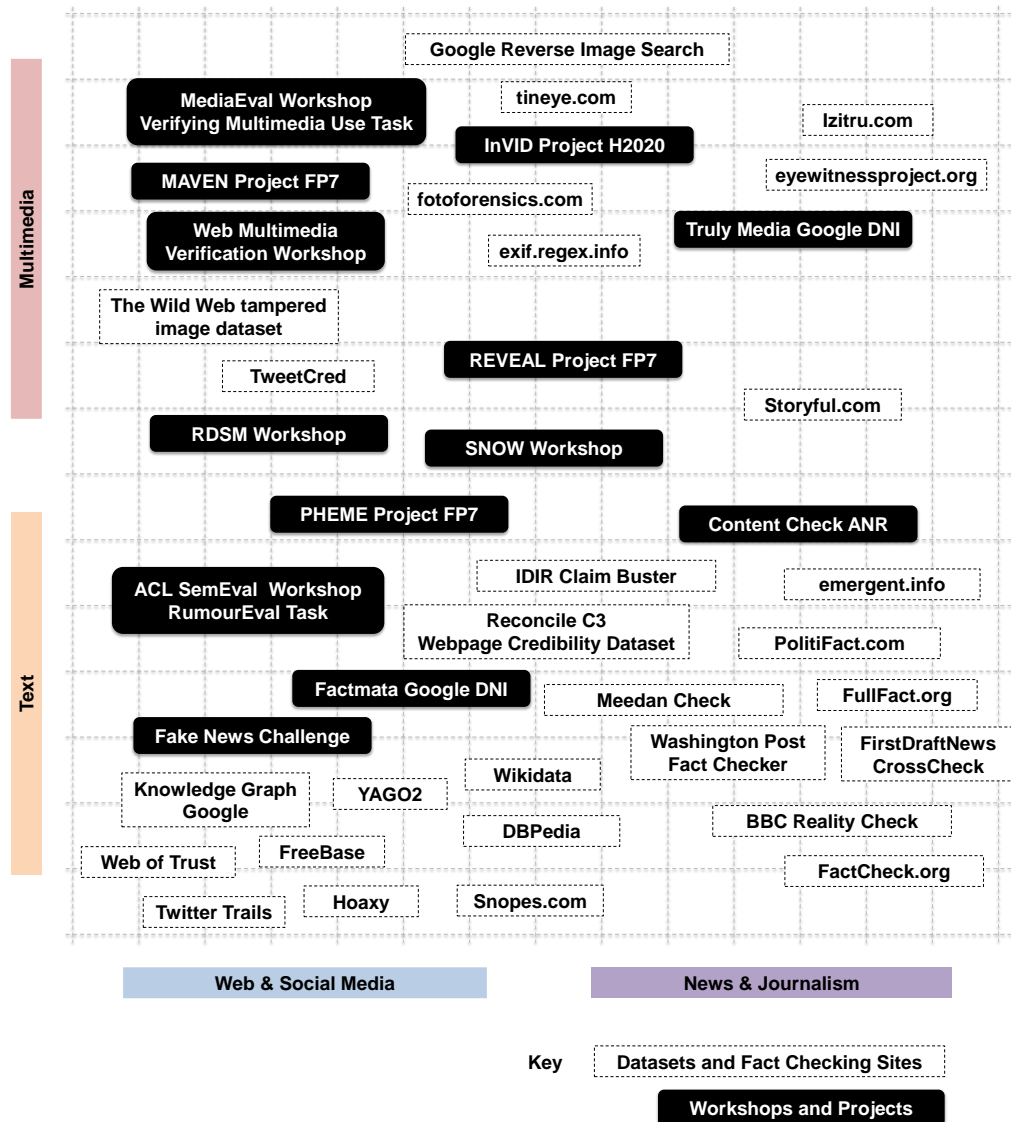


Figure 4: Research and dataset landscape for researchers interested in verification of social media and news-related content

Acknowledgement

This work was supported by the REVEAL FP7 and the InVID H2020 projects, partially funded by the EC under grant agreements 610928 and 687786 respectively.

References

- [1] S. Wiegand, S.E. Middleton, "Veracity and Velocity of Social Media Content during Breaking News: Analysis of November 2015 Paris Shootings", Third Workshop on Social News on the Web (SNOW-2016), WWW 2016, Companion of the 25th International World Wide Web Conference, 2016.
- [2] A. Zubiaga, A. Aker, K. Bontcheva, M. Liakata, & R. Procter, "Detection and Resolution of Rumours in Social Media: a Survey", arXiv preprint, arXiv:1704.00656, 2017.
- [3] C. Silverman, "Verification Handbook: a Definitive Guide to Verifying Digital Content for Emergency Coverage", European Journalism Centre, 2014.
- [4] G.K. Birajdar, V.H. Mankar, "Digital Image Forgery Detection using Passive Techniques: A Survey", Digital Investigation, 10(3), 226-245, 2013.
- [5] M. Zampoglou, S. Papadopoulos, Y. Kompatsiaris, "Detecting Image Splicing in the Wild (WEB)", in IEEE International Conference on Multimedia & Expo Workshops (ICMEW), Turin, Italy, pp. 1-6, 2015.
- [6] M. Zampoglou, S. Papadopoulos, Y. Kompatsiaris, R. Bouwmeester, J. Spangenberg, "Web and Social Media Image Forensics for News Professionals", in Tenth International AAI Conference on Web and Social Media, Social Media in the Newsroom Workshop, 2016.
- [7] M. Zampoglou, S. Papadopoulos, Y. Kompatsiaris, "Large-scale Evaluation of Splicing Localization Algorithms for Web Images", Multimedia Tools Appl. 76, 4, 4801-4834, 2017.
- [8] C. Castillo, M. Mendoza, B. Poblete, "Information Credibility on Twitter", in Proceedings of the 20th international conference on World wide web (WWW '11). ACM, New York, NY, USA, 675-684, 2011.
- [9] C. Boididou, S. Papadopoulos, L. Apostolidis, Y. Kompatsiaris, "Learning to Detect Misleading Content on Twitter", in Proceedings of the International Conference on Multimedia Retrieval. ACM, 2017.
- [10] M. Babakar, W. Moy, "The State of Automated Factchecking", Full Fact, 2016.
- [11] N. Hassan, B. Adair, J.T. Hamilton, C. Li, M. Tremayne, J. Yang, C. Yu, "The Quest to Automate Fact-Checking", Computation+ Journalism Symposium, 2015.
- [12] S. Teufel, "Argumentative Zoning: Information Extraction from Scientific Text." PhD diss., University of Edinburgh, 2000.
- [13] M.S. Mausam, S. Soderland, R. Bart, O. Etzioni, "Open Language Learning for Information Extraction", In Proceedings of Empirical Methods in Natural Language Processing, 2012.

- [14]M. Yahya, S.E. Whang, R. Gupta, A. Halevy., “ReNoun: Fact Extraction for Nominal Attributes”, Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 325–335, 2014.
- [15]W.L. Bennett, “News: The Politics of Illusion”, Tenth Edition, University Of Chicago Press, ISBN: 9780226344867, 2016.
- [16]B. Walenz, Y. Wu, S. Song, E. Sonmez, E. Wu, K. Wu, P.K. Agarwal, J. Yang, N. Hassan, A. Sultana, G. Zhang, C. Li, C. Yu, “Finding, Monitoring, and Checking Claims Computationally Based on Structured Data”, Computation + Journalism Symposium, Columbia University, 2014.
- [17]S.E. Middleton, V. Krivcovs, "Geoparsing and Geosemantics for Social Media: Spatio-Temporal Grounding of Content Propagating Rumours to support Trust and Veracity Analysis during Breaking News", ACM Transactions on Information Systems (TOIS), 34, 3, Article 16, 26 pages, 2016.
- [18]P.B. Brandtzaeg, A. Kuczerawy, P. Valcke, I. Graef, I. Lambrecht, J. Ausloos, A. Følstad, “D7.2 User Evaluation Report”, REVEAL project, 2017.
- [19]I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, “Generative Adversarial Nets”, Advances in Neural Information Processing Systems (NIPS), 27, 2014.

Bio

Stuart E. Middleton is a senior research engineer at the University of Southampton IT Innovation Centre. His main research interests are social media, machine learning, information retrieval and semantics. Stuart has a PhD in Computer Science from the University of Southampton and is a senior member of the ACM. Contact him at sem@it-innovation.soton.ac.uk.

Symeon Papadopoulos is a post-doctoral research fellow at the Information Technologies Institute (ITI), part of Centre for Research and Technology Hellas (CERTH). His main research interests are multimedia processing and indexing, information retrieval, machine learning and social network analysis. Symeon has a PhD in Computer Science from the Aristotle University of Thessaloniki. Contact him at papadop@iti.gr.

Yiannis Kompatsiaris is a senior researcher at the Information Technologies Institute (ITI), part of Centre for Research and Technology Hellas (CERTH). His main research interests are semantic multimedia analysis, indexing and retrieval, social media and big data analysis, knowledge structures, reasoning and personalization. He has a PhD in Computer Science from the Aristotle University of Thessaloniki. Contact him at ikom@iti.gr.