VERIFICATION HANDBOOK

AN ULTIMATE GUIDELINE ON DIGITAL AGE SOURCING FOR EMERGENCY COVERAGE.

European Journalism Centre

EDITED BY CRAIG SILVERMAN

EDITOR OF 'REGRET THE ERROR', THE POYNTER INSTITUTE

Chapter 3:

Verifying User-Generated Content

Claire Wardle is a research fellow at the Tow Center at Columbia University, working on a major research project on user-generated content and television news. She designed the social media training program for the BBC in 2009 and went on to train journalists around the world on social newsgathering and verification. For the past two years Claire has been working with Storyful[1]. Claire has a Ph.D. in Communication from the Annenberg School for Communication at the University of Pennsylvania. She is @cward1e[2] on Twitter and blogs at clairewardle.com[3].

In less than a decade, newsgathering has been transformed by two significant developments.

The first is mobile technology. In the summer of 2013 an important tipping point was reached. For the first time, more than half (55 percent) of all new mobile phone handsets sold were smartphones[4].

By definition a smartphone has a high-quality camera with video capability, and it allows the user to easily connect to the Web to disseminate the pictures. As a result, more and more people have the technology in their pockets to very quickly film events they see around them, and share them directly with people who might be interested, as well as more widely via social networks.

The second, connected development is the social Web. When the BBC's User Generated Content Hub started its work in early 2005, they were reliant on people sending content to one central email address. At that point Facebook had just over 5 million users, rather than the more than one billion today. YouTube and Twitter hadn't launched. Now, every minute of the day, 100 hours of content is uploaded to YouTube[5], 250,000 tweets are sent and 2.4 million pieces of content are shared on Facebook.[a] Audience behavior has shifted substantially.

Rather than film something and, when prompted, send it to a news organization, people shoot what they see and upload it to Facebook, YouTube or Twitter. Research has shown very few audience members have enough understanding of the news process to think of their footage as valuable enough to send it, unprompted, to a news organization or other entity.[b] Essentially, they're uploading the content to share the experience with their friends

---

a These stats change all of the time, but this is the most recent attempt at measuring activity on the most popular social networks http://blog.qmee.com/qmee-online-in-60-seconds/ «

b http://www.bbc.co.uk/blogs/knowledgeexchange/cardiffone.pdf «

and family.

Increasingly, at any news event around the world there are "accidental journalists": people standing in the right place at the right time with a smartphone in their hands. As Anthony De Rosa, the former social media editor for Reuters and current editor-in-chief of Circa, writes[6]: "The first thought of the shooter is usually not: 'I need to share this with a major TV news network' because they don't care about traditional television news networks or more likely they've never heard of them. They have, however, heard of the Internet and that's where they decide to share it with the world."

Similarly, during breaking news events, the audience is often more likely to turn to social networks for information, meaning first responders and emergency organizations are using social networks themselves. Unfortunately, these news events invite false information to circulate, either deliberately or by accident. Therefore, journalists and humanitarian professionals should always start from a position that the content is incorrect. During emergencies, when information can literally affect lives, verification is a critical part of the newsgathering and information dissemination process.

The importance of verification

The ability for anyone to upload content, and to label or describe it as being from a certain event, leaves many journalists, and particularly editors, terrified about being hoaxed or running with false content.

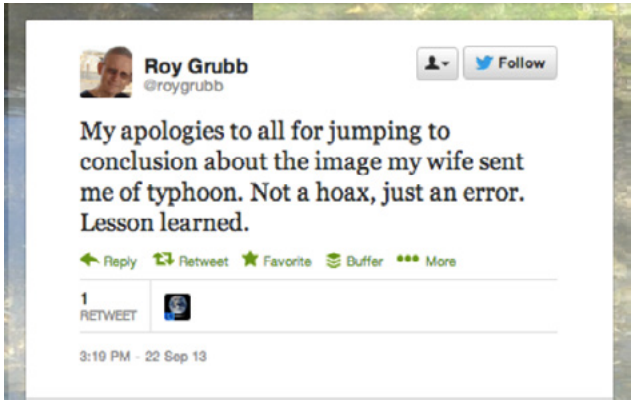

Some people go out of their way to deliberately hoax news organizations and the public by creating fake websites, inventing Twitter accounts, Photoshopping images or editing videos. More often, the mistakes that happen aren't deliberate. People, trying to be helpful, often find mislabeled content from previous news events and share it. Below is an example of a man apologizing after tweeting a photo[7] emailed to him by his wife. She had told him it

showed Typhoon Usagi as it headed toward Hong Kong; in fact it was an old image of another event.

People downloading content from YouTube and uploading it to their own accounts, claiming it as their own, cause other problems. This isn't a hoax - it's what is known as a "scrape" - but it means we have to work harder to find the original uploader of the content.

The difficulty of finding original footage was demonstrated when the U.S. Senate Intelligence Committee released a playlist of 13 videos[8] that had originally appeared on YouTube, which they had used to look for evidence related to the 2013 chemical weapons attack on East Gouta in Syria. A number of these videos were taken from a well-known Syrian aggregator YouTube channel which regularly republishes videos from other people's channels. This suggested the videos within the playlist were not the original videos and were in fact "scrapes." Using a range of different verification techniques, Félim McMahon from Storyful was able to discover the original versions of these videos. He wrote up the process here[9]. What this example shows is that these issues are no longer just a concern for the journalism community.

Verification checks

Verification is a key skill, made possible through free online tools and old-fashioned journalism techniques. No technology can automatically verify a piece of UGC with 100 percent certainty. However, the human eye or traditional investigations aren't enough either. It's the combination of the two.

When a journalist or humanitarian professional finds a piece of information or content via social media, or has it sent to her, there are four elements to check and confirm:

1. Provenance: Is this the original piece of content?
2. Source: Who uploaded the content?
3. Date: When was the content created?
4. Location: Where was the content created?

1. Provenance: Confirming the authenticity of the piece of content

If you find content on a social media profile, you have to run a number of checks on that profile to make sure it is real.

In the case of a tweet, be aware that the site lemmetweetthatforyou.com[10] makes it shockingly easy to fake a tweet, which can be then shared as a picture.

Another way people spread fake information on Twitter is by presenting the fake information as a retweet. For example: “Really? RT@JoeBiden I’m announcing my retirement from politics.” That makes it appear as if you’re simply retweeting an original tweet.

Fakers also often add a Twitter blue verification check mark to the cover photo on a faked account to make it appear legitimate. To check whether an account is actually verified, hover over the blue tick, and you will see the text “verified account” pop up. If it’s not there, it is not a verified account.

Facebook introduced a similar verification program, using the same blue tick system, for celebrities, journalists and government officials. Verified ticks can appear on Facebook pages as well as personal profiles. (As with Twitter, Facebook manages the verification program, and decides which verification requests to accept.) On Facebook pages, such as Usain Bolt’s below, the tick appears underneath the cover photo, next to the person’s name.

On personal profiles, the tick appears on the cover photo. Here’s the profile of Liz Heron, editor of emerging media at The Wall Street Journal:

It’s worth noting that, as with Twitter, people have been known to Photoshop blue ticks onto cover photos. So, as with Twitter, if you hover your mouse over the blue tick, the phrase “verified profile” will appear.

But as with Twitter, remember the verification process is far from transparent, so with less-famous people, it can be unclear whether an unverified account is a fake, or whether they’re just not famous enough to be verified!

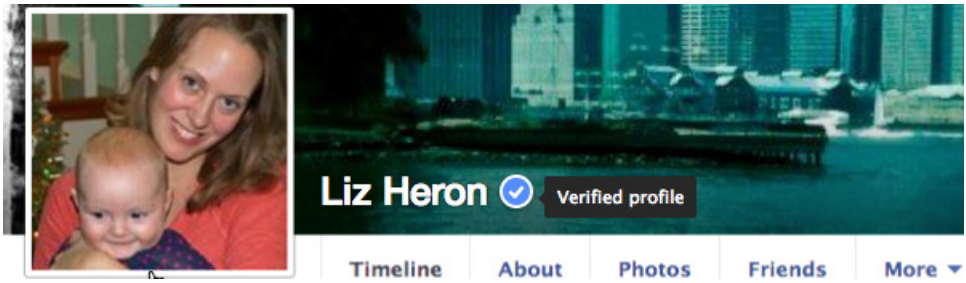

But even with these official verification programs in place, there is no quick way of checking whether an account is real, other than painstaking checks on all of the details available on the profile. Items to review include linked websites, location, previous pictures and videos, previous status updates or tweets. Who are their friends or followers? Who are they following? Do they feature on anyone else's lists?

If you're looking at a piece of rich content, such as a photo or video, one of the first questions is whether this is the original piece of footage or picture. Using reverse image search tools such as TinEye or Google Images[3] you can find out whether it has been posted online previously. (For more detail on using these tools, see Chapter 4 of this book.)

While deliberate hoaxes are rare, they do happen. In recent years there have been relatively harmless hoax videos produced by PR companies looking for publicity[11], and by students completing an end-of-term assignment[12]. There have also been deliberate attempts to create false content, particularly in Syria and Egypt, where discrediting the "enemy" can be achieved via reputable-looking content shared on social media channels.

Techniques include creating a false, but identical-looking website and claiming responsibility for a bomb attack[13], or staging a gruesome incident and blaming the other side. Manipulation is relatively easy to do today, and whether you're Nancy Pelosi trying to create a photograph of all female Congresswomen[14] even when some of them are late, or a Syrian activist group sharing video of a man appearing to be buried alive[15], any journalist or humanitarian professional has to start off by assuming a piece of UGC is false. (See Chapter 5 of this book for more detail about verifying video.)

2. Confirming the source

The ultimate goal when attempting to verify UGC is to identify the original uploader and get in touch with them.

In that conversation, the key questions involve discovering where someone was standing when they took the footage, what they could see, and the type of camera used to record the footage. (These questions provide the essential data to answer Steve Buttry's essential "How do you know that?" test outlined in the previous chapter.)

If someone is attempting to pass along false information, either deliberately or not, asking direct questions will often result in the person's admission that they did not actually film the footage themselves. Additionally, it is possible to cross-reference answers to some of these questions with available information by examining the EXIF data in a photo, or comparing video of a specific location to Google Street View, which we detail in subsequent chapters.

But first you have to find the person responsible for the content. Researching the history of an uploader can mimic the characteristics of an old-fashioned police investigation, and perhaps also make you feel more like a stalker rather than a journalist or researcher.

Some people list a great deal of information on their social profiles, and a real name (especially one that is not too common) can provide a wealth of information. As people live more of their lives on different social networks, they are often unaware how clues can be combined to build up a substantial dossier of information. A YouTube profile with little personal information listed but that includes a website URL can lead a journalist to a person's address, email and personal telephone number, via the website who.is.[c]

3. Confirming the date of the event

Verifying the date of a piece of video can be one of the most difficult elements of verification. Some activists are aware of this fact and will show a newspaper from that day, with the date clearly visible when they share their footage. This obviously isn't foolproof, but if an uploader becomes known and trusted by organizations, be they news or humanitarian, this is a helpful additional piece of information.

Be aware that YouTube date stamps its video using Pacific Standard Time. This can sometimes mean that video appears to have been uploaded before an event took place.

Another way to help ascertain date is by using weather information. Wolfram Alpha[16] is a computational knowledge engine that, among other things, allows you to check weather from a particular date. (Simply type in a phrase such as "What was the weather in Caracas on September 24, 2013" to get a result.) This can be combined with tweets and data from local weather forecasters, as well as other uploads from the same location on the same day, to cross-reference weather.

4. Confirming the location

Only a small percentage of content is automatically geolocated, but mapping platforms - Google Maps, Google Earth, Wikimapia - of the first checks that need to be performed for

---

c *A journalist should always check both of these tools. Sometimes results can emerge on one and not the other.* «

video and photos, and it is quite incredible what can be located.[d] Geolocation is always more difficult, however, when the imaging is out of date, for example in Syria, subject to damage from bombs or shelling, or on Long Island after Hurricane Sandy.

Activists who are aware of the challenges of verification often pan upward before or after filming some footage to identify a building that could be located on a map, whether that's a tall tower, a minaret or cathedral, or signpost. This is partly a result of news organizations' asking activist groups to do this,[e] as well as activists themselves sharing advice about best practice[17] when uploading UGC.

Verification as process

Unfortunately, people often see verification as a simple yes/no action: Something has been verified or not.

In practice, as described above and in subsequent chapters, verification is a process. It is relatively rare that all of these checks provide clear answers. It is therefore an editorial decision about whether to use a piece of content that originates from a witness.

Two recent academic studies performed content analysis of output on the BBC and Al Jazeera Arabic. They found that while these verification checks are undertaken by editorial staff, and considered absolutely necessary, the results of the checks are rarely shared with the audience.

As Juliette Harkin concluded in her 2012 study[18], "[n]either BBC Arabic nor Al Jazeera Arabic explicitly mentioned in any of the programs or video packages that were evaluated whether the sources were verified or were reliable. The common on air explanation of 'this footage cannot be verified,' was absent in all the content evaluated for this study."[f]

There are recent moves to increase transparency with the audience about the verification checks made by journalists when a piece of UGC is used by a news organization. The AP and BBC are both working toward making their verification processes clearer; in August 2013, the BBC said[19] that since a comprehensive study into the use of UGC during the Arab Spring[20], "the BBC has adopted new wording for all user-generated footage where independent verification has not been possible," letting its audience know what it knows.

It is likely that within the next few years, a new grammar of verification will emerge, with

---

d See this post about geolocating the position of a tank explosion in Syria: http://blog.storyful.com/2013/03/13/the-changing-nature-of-conflict-and-technology/ «

e See Harkin study «

f See Harkin study, p. 31 «

the audience expecting to be told what is known and what isn't known about a piece of UGC sourced from social media. With the audience able to see the same footage as the news organizations and others that gather material from the crowd, this level of transparency and accountability is required.

Case Study 3.1:
Monitoring and Verifying During the Ukrainian Parliamentary Election

Anahi Ayala Iacucci is the senior innovation adviser for the Internews Center for Innovation & Learning[1] and the Internews Humanitarian Media Project. Over the past four years, she has worked on the applications of technology and innovation to humanitarian crises, media development, conflict prevention and human rights around the world for organizations like the World Bank, the U.N., NDI and Freedom House, among others. She blogs at anahiayala.com[2] and tweets @anahi_ayala[3].

During the Ukrainian parliamentary elections of fall 2012, Internews Ukraine[4], a local NGO supported by the global nonprofit media organization Internews[5], ran an election monitoring project called Elect.UA. It used a mix of crowdsourcing, mobile phones, social media, professional electoral monitoring and media monitoring to oversee the electoral campaign, and possible violations of or tampering with the results.

The project was built upon a fairly complex structure: 36 journalists around the country reported stories during the electoral campaign and on election day. At the same time, three different electoral monitoring organizations had workers reporting to the same platform using SMS, online forms and emails. Elect.UA also invited Ukrainians to report about their election experience using social media (Twitter and Facebook), mobile technology (SMS and a hotline number), a smartphone app, an online form or email.

All information coming from Internews-trained journalists and electoral monitors was automatically tagged as verified, while messages from the crowd were vetted by a team of 16 administrators in Kiev.

For the messages coming from the crowd, the admin team set up a verification protocol based on the source of the information: mobile technology, social media, online form or email.

For each source, the team would try to verify the sender of the information (when possible), the content of the information and the context. For each of those components the team would also try to establish if something could be 100 percent verified, or only partly verified.

For information coming via social media, the below image shows the decision tree model used by administrators in the verification process.

For a full version of the diagram, see page 120

The first step was to perform an online search of the information and its source to identify all possible digital traces of that person, and the piece of content. (For example, we examined other social media accounts, mentions by media articles, information about university, affiliations, etc.). The search was aimed at determining if the person was a reliable source, and if there was a trace of the information they provided elsewhere online.

The second step was to use the information collected to build a profile of the person, as well as a profile of the content they provided. For each of the 5Ws - who, what, when, where and why - administrators had to carefully determine what they could prove, and what they could not.

For multimedia content, the source verification protocol was the same, but we had a different path for the content. Photos and video were verified by looking for any identifiable landmarks, and by performing an analysis of the audio (to listen for language, dialects, slang words, background noise, etc.), clothes and of light (artificial or natural), among other elements in the content.

When a piece of information could not be verified with a sufficient degree of certainty, the report was sent back to an electoral monitor or a reporter on the ground for real-time, in-person verification.

For example, on September 28, 2012, Elect.UA received an anonymous message via its website that parliamentary candidate Leonid Datsenko had been invited for a discussion by a stranger, and then was intimidated in order to force him to withdraw from the elections.

The next day, the administrators of the platform found an article[6] in a reliable media source that included a record of the exchange. We still held the report for verification, and then, on October 1, local journalists reported on a press conference about the incident[7]. Elect.UA's local journalists also conducted interviews with local law enforcement services, who acknowledged this case to be true.

Overall, the Elect.UA team managed to verify an incredible amount of information using these protocols, and also noticed that the more the administrators became familiar with the verification process, the faster they were able to work. This proves that the verification of user-generated content is a skill that can be systematized and learned, resulting in efficient, reliable results.

The decision tree model: