Digital Verification and its Discontents: Investigating Tear Gas Abuse in a Digital Age

1. Digital Verification and its Discontents: Investigating Tear Gas Abuse in a Digital Age

Rebekah Lyndon and Michael Gyan Nyarko (Amnesty International Digital Verification Corps, Cambridge University)

In this contribution, Lyndon and Nyarko highlight the processes involved in verifying relevant material in open source investigations. Focusing on the worldwide abuse of tear gas – used by authorities as a so-called "less-lethal" weapon – the authors also outline some of the pitfalls and ethical concerns faced by researchers. Beyond the promising approaches offered by digital verification techniques, Lyndon and Nyarko advise that we recognise and carefully negotiate the tensions involved in order to attain a higher standard of digital ethics.

Introduction

Amidst selfies and memes, social media hosts content that serves different purposes: recent years have seen a proliferation of recordings by eyewitnesses of human rights violations. This user-generated content (UGC) includes evidence of human rights violations, a source that human rights researchers increasingly look to for evidence of human rights violations (Aronson 2018). Offering critical documentation to a digital public, these witnesses expose how law enforcement bodies' use of weapons can amount to punitive rather than safe, legal, necessary, and proportional deployment. Their uploads are reshaping the nature of human rights investigations. One of our recent projects addressed the worldwide abuse of tear gas as recorded by eyewitnesses – a substance whose status as a so-called "less-lethal" weapon allows it to be used by police forces against crowds of peaceful civilian protestors. In this paper, we highlight the processes involved in verifying relevant material and some of the pitfalls and ethical concerns faced by researchers in open source investigation.

An Evolving Context: Human Rights Violations, Online Documentation, and OSINT

Human rights practitioners have increasingly turned to social media and other open-source content as part of their pursuits of accountability, whether raising public awareness or as evidence in legal contexts (Minogue and Makumbe 2019). This provides a complement to interviews of witnesses that have historically often been framed according to the priorities of a distanced interviewer rather than the interviewee's experiences, and that can risk re-traumatisation. Trauma may affect memory, and images and videos can draw attention to the survivor or witness' perspective, while often adding context on spatial and group dynamics.

Amnesty International's Digital Verification Corps (DVC) – trained volunteer groups currently based at human rights centres at seven universities globally: Pretoria, Iberoamericana, Hong Kong, Essex, Berkeley, Toronto, and Cambridge – work alongside its full-time Crisis Response Team, primarily by assessing the validity of visual evidence of possible human rights abuses from the crises that it tackles. In 2018, we began identifying cases of tear gas abuse, culminating in an interactive online incident map released in 2020 (Amnesty International 2020c).

To reach this stage, verification was a fundamental step; accuracy is critical when the normalisation of misinformation and disinformation – and, just as importantly, our wary expectations of these – has made it easy for people in power to detract or distract from legitimate issues. Rather than presuming that online eyewitness accounts are not trustworthy, the verification process seeks to ensure that this kind of advocacy is impervious to attempts by governments or other powerful figures to dismiss people's valid claims by crying "fake news." This has allowed previous DVC work to challenge Hong Kong police's denial of abuses against protestors and to obtain acknowledgment on the part of the US-led coalition of some of the indiscriminate damage it inflicted through airstrikes on civilian areas in Raqqa, Syria (Amnesty International 2019a, 2019b). Nonetheless, experience in verification work only serves to highlight that the process is far more complex and patchwork than the straightforward true-false binary that the term implies.

In short, open source investigations involve collecting information from publicly available sources and analysing it to draw meaningful conclusions. Our focus lay in applying a mixture of digital tools and human analysis specifically to content publicly uploaded online. These methods reflect not only the growing range of digital investigative options available to human rights researchers but also a divergence from some of the traditional ends of open-source investigations. The term OSINT – open-source *intelligence*, the findings gleaned from investigations with open source *information* – originally emerged in US defence sectors, indicating information that was not classified, or that did not involve covert or clandestine collection (Stottlemyre 2015; Williams and Blum 2018). Its usage in military contexts continues into modern times, but today we also see growing interest in OSINT from the private sector. This often relies on broad automated tracking, including problematic sentiment analysis, to gauge industry trends rather than the close human assessment of specific individual instances needed to contextualise possible human rights abuses. Human rights practitioners find themselves uncomfortable outliers in fields dominated by the private and public sectors, including murky areas of military surveillance, while simultaneously relying on – and seeking to reshape – tools defined and developed by those sectors.

Human rights activists face a range of challenges online, from internet shutdowns to the increasing threat of "deepfakes" which will potentially make verification even more challenging in the future (Witness Media Lab 2019). Yet the ever-growing swathes of online data alone present multiple obstacles to human rights researchers. At a practical level it is hard to keep track of, sift through, and make sense of all the relevant information uploaded online. Perhaps more dangerous is the resulting false sense of security that this prevalence can foster regarding the documentation of abuses – at worst, the assumption that if not provable, or not online, something must be untrue. As human rights investigators increasingly turn to online content for evidence, we must neither allow this to undermine the value placed on verbally-shared testimonies nor relent in sensitive campaigning against abuses that are perhaps harder to document on social media: in less public spaces (McAvoy 2021); where stigma may be a greater factor; where people face political or legal intimidation for publicly "slandering" those in power; where recording would be life-threatening.

The Digital Verification Process

Identifying any possible relevant cases requires practical steps to navigate the reams of online material, and often involves a reliance on applications' own advanced search functions, Boolean search terms, and third-party filtering applications like TweetDeck. We began research with text searches for terms related to tear gas in various languages. Twitter and YouTube generally offer the best starting points for locating evidence of human rights abuses; well-known and widely used, they not only host significant levels of content, but eyewitnesses are well-aware of possible broad audiences when seeking acknowledgment of the abuses they have documented. Importantly, these are sites commonly used for the public sharing of content and understood as such, whereas other sites may seem hazier in terms of privacy settings or patterns of usage, raising ethical questions around scrutinising and sharing content. The research also mutually supported some of our concurrent projects, such as investigations into police violence against those protesting for economic justice in Chile or for Black Lives Matter in the USA, where tear gas abuse was one among many forms of violence (Amnesty International 2020a, 2020b).

As we collated social media posts relevant to tear gas misuse, we created lists of buzzwords and hashtags, with careful attention to language and avoidance of automated translation. Delegation to a team member with relevant linguistic knowledge or consulting the expertise of a native speaker can ensure that details are not overlooked – for instance, including "cra." as an alternative to "carrera" in noting references to Bolivian and Colombian roads, or recognising different vocabulary for key infrastructure even across countries that share languages, such as police stations in Sudan and Syria. This critical mentality has been equally crucial when handling English-language material; regardless of language many crises involve evolving political references and slang. Collating evidence on the misuse of tear gas against Nigerian protestors in October 2020 required careful attention to alternative spellings of key hashtags – #EndSARS, #EndSarsNow, #EndPoliceBrutalityinNigeria, etc. – and terms most commonly used alongside them. In most projects, initial vocabulary lists garnered from online content in turn fed further searches, as did any key terms mentioned verbally within footage and any significant locations or timeframes we were able to identify through the verification process. Per post, we would carefully analyse any accompanying text and comments and the uploader's profile to better understand questions of provenance, social networks and related posts; to seek any indication of location or date; and to rule out obvious malicious bot involvement. Additional footage that can be identified as from the same event as another might be included, regardless of whether it depicted the specific abuse in question, as consideration from multiple angles can be extremely helpful – for instance, one might include close-up detail of an injury, while another shows no abuse but more of the environs, aiding geolocation. This understanding of the environment is a crucial stage, not only allowing confirmation of where an event took place but also an assessment of any enclosed or limited spaces where citizens would be unable to escape tear gas – despite its explicit purpose for crowd dispersal – and where its usage may arguably constitute torture. Across the globe, eyewitness footage indicates that security forces knowingly deployed tear gas in illegitimate contexts including indoors public spaces, on bridges and bottlenecks, in dead-ends, and where civilians' paths were otherwise obstructed by fixed objects.

Potential evidence should be preserved as soon as possible. The novelty of digital landscapes makes it seem incongruous to approach online material as historical evidence, but like any form of documentation, it risks damage and degradation or editing. Managing storage is crucial, as both hardware and software can rapidly become redundant (Ng 2020). Within open source investigations for human rights, a plethora of reasons complicate the longevity of online content and necessitate thoughtful archiving (Piracés 2018). A depiction of violence might be deemed as violating a platform's terms of service by an unclear algorithm or human content moderator, themselves facing infringements of their labour rights (Crider 2020). In other cases, direct political censorship is involved. Some websites experience link rot. Countless tools are available for different media, from Archive.org's WaybackMachine website preservation tool to the independent, open-source Youtube-dl software for downloading audio-visual content, accessible on GitHub. Physical drive-based storage often provides better security, while cloud-based options are easier for collaboration. Ultimately investigators must find a structured approach to research while still disallowing complete routinisation. This is due to a need for space to query rather than accept ethical issues around the non-neutrality of data, digital tools, and the online infrastructure itself, as well as to adapt to changing technologies and patterns of internet usage.

Often, content needs scrutinising multiple times, focusing on different elements separately before verification. Carefully ascertaining what is happening with the gas is essential to determine whether it is being abused; identifying key individual officers can allow investigation of the broader chain of command beyond the event; noting any close-up detail of canisters can facilitate further analysis of shipments, trade, and manufacturing. Local identifiers can help with geolocation – these can include anything from the style of road signage and vehicle registration plate formats to dialects spoken. Typically relying on web searches, researching details like these is the simplest and most intuitive part of the process, yet still requires discerning and conscious use of search engines. Options include setting up a fresh browser profile cleared of history and location, to mitigate skewing results to a personalised footprint, and using Boolean search terms to render more appropriate results. Countless tools and databases exist to help with identification of details, preservation, image analysis, search organisation, analysis of public data, and geolocation, but websites like OSINT Essentials and the Digital Methods Initiative supply excellent collations to start with, while Amnesty's Citizen Evidence Lab offers blog-style guides for open source research.

When developing methodologies, there is often an attraction towards complexity, perhaps especially when operating in collaboration with academic spaces where constant pressure is placed on proving research relevance. Yet human rights investigators must be wary of tools that may appear more compelling merely due to novelty, impressiveness, or exclusive training required; rather, simplicity and replicability are crucial to ensuring that this work remains accessible not only in its basis in open source information but in terms of the tools and methodologies used. Many open source researchers archive content with programmes that require a basic knowledge of the command-line or python, but tools that require no coding knowledge are often as effective and these questions should not distract from the fundamentals of secure storage and good organisation – preservation is undermined if it is difficult to locate a piece of evidence. It can be challenging to estimate how time-consuming footage may be to geolocate but also to determine where additional experts may need consulting in terms of weapons used, for instance, but organisation and communication over these issues need not be overly complicated. While we sometimes use tools like the valuable platform Truly Media to systematise our verification, a well-organised shared spreadsheet works just as well for our purposes. At a minimum, categories should focus on confirming time and location of a recorded incident – the crux of verification work – as well as media type, links to the original upload and preservation details, description of the action, and progress with investigation. Though seemingly obvious, the importance of consistent, agreed-upon labelling cannot be overstated, including communication on delegation and steps taken thus far with a given video – especially when working with broad teams across different locations. More broadly, when navigating the evolving meaning of "open source," researchers might take into consideration the role of privatisation and paywalls, the surrender of personal data, and user accessibility globally and in terms of neurodivergence; these questions should certainly inform research methods.

After consolidating our initial evidence base, we would start conducting reverse image searches. With videos, this involved taking screenshots of the clearest frames with the most potential for recognisable content – perhaps a unique building or identifiable landmark. Image search engines like TinEye, Google, Yandex, and Bing are particularly effective at revealing where an identical image has previously surfaced online, in anything from journalism or stock image databases, allowing us to rule out occasional cases of misinformation or instances where persons have paired their image-less descriptions of recent events with older content as a form of illustration. InVid is a particularly useful tool offering the functionality to auto-screenshot images from videos and complete reverse image searches. Yet using multiple sites, flipping an image, and changing resolution is best as each software emphasises different aspects of an image. Where content is original, pixelating out persons to emphasise background can sometimes hone results in attempting to find similar images for geolocation, but ultimately the key search engines are skewed towards content from the global North and overwhelmingly provide unhelpful suggestions. Search engines are obscure about their automated processes and never failsafe.

Human assessment of the physical built environment in any media depicting tear gas abuse was triply significant: for geolocation, chronolocation, and assessments of abuse. As such, we were meticulous in our observation of imagery, paying as close attention as possible to the surroundings of events – often having to train ourselves to counterintuitively ignore the incident in question in favour of its background – and meticulously take note of buildings, structures, and street layouts. This was primarily central to our geolocation process. Yet once we could pinpoint a video on the map, this also facilitated our chronolocation by allowing us to identify any discrepancies or parallels with recent changes visible in any other sources, such as ongoing construction or a shop's new paint job. Thirdly, we could take this more intimate analysis of the physical spaces, enhanced with satellite imagery and other online content, to confirm any enclosed spaces preventing crowd dispersal, as outlined above.

The need for attention to layered aspects of footage – geography, language, climate, movement, weaponry, and more, often in individual freeze-frames – raises the likelihood of repeated exposure to violent imagery. Research indicates that intimacy with this kind of material has cumulative, subconscious psychological effects on viewers that can culminate in vicarious trauma (Dubberley and Grant 2017). Alongside personal and cultural appreciation of mental health, practical measures include fostering a healthy, structured work environment. Content labels are not only helpful for maintaining digital hygiene and communication while pursuing multiple tangents simultaneously, but also for mental preparation. Deliberate scheduling and a clear division in closing all digital workspaces before personal use is crucial, particularly in remote and work-from-home contexts where investigators may not be as readily surrounded by support. Though it is tempting for investigators to downplay these measures and the effects of footage, particularly in comparison to the undeniably more manifest results of the depicted actions, desensitisation is not resilience, and is not sustainable. Indeed, in order to continue appropriately recognising the weight of the content at hand and to avoid a cynicism that downplays effects of footage – initially on oneself but ultimately of the content itself on eyewitnesses and survivors – investigators must recognise that psychological resilience and self-care are central components of effective verification.

Social media posts with no description or context still provide some starting points for confirming the date of an incident; a time-stamp is default, though unreliable. When visual content is uploaded to one of the main social media sites, it is stripped of its metadata (embedded information including time taken and device and settings used). This leaves any observer reliant on the platform's time stamp, though this indicates upload time only – in the viewer or the uploader's time-zone, depending on the site – rather than reflecting the content creation. Any number of reasons may distance the timing of an online upload from the incident it purports to depict, from the physical threat when faced with armed personnel to simply waiting to reconnect to Wi-Fi.

Yet material is often uploaded soon after an incident, and it makes sense to work backwards from the time of posting, querying whether anything disallows the possibility of something having been filmed within a given timeframe. For instance, if a clip vastly contradicts the historical weather record for a given location, either the chrono- or geo-location is likely flawed. Meanwhile, daylight and shadow length can be assessed with the help of tools like SunCalc, which allows an exploration of the position of the sun in the sky relative to any point on the earth's surface on any date since 1900. Originally designed to assist photographers in making decisions about lighting, it reflects how much of this work involves an evolving toolkit of repurposed software.

Geolocation typically constitutes the most time-consuming and arguably most fundamental stage of verification. As aforementioned, local identifiers, like language or officers' uniforms, narrowed down the respective countries, and physical features within the recordings were invaluable: station signage led us to cycle through Hong Kong subway stops; deciphering the name of a supermarket meant we could identify its chains within Guayaquil. Usually clues were more subtle and entailed hours of systematically scouring satellite imagery to identify where the physical landscape lined up with the imagery. With tear gas typically used against protests taking place on city streets, Google Street View was a rewarding tool, as well as the map function and more collaborative projects like OpenStreetMap. However, the efficacy of Street View is often hampered by its limited coverage of some neighbourhoods or countries. Well-developed countries with high internet penetration experience easily accessible and up-to-date Street View functionality, while underserviced areas may attract limited to no coverage. Geolocation is often facilitated if a team member has previous knowledge of the environment depicted in the material. For instance, Nigerian team members, familiar with specific buildings and streets depicted in some videos, helped the Pretoria team to geolocate material to Abuja far more rapidly than they would otherwise. Local knowledge remains the most useful resource for verification, and prioritising it can help reduce any problematic sense of distance between analysts and witnesses.

Google Earth Pro, the free desktop version of the software, offers advantages over the website, particularly in terms of historical imagery and annotation options. By including sliders that toggle between current and previous satellite imagery of a given area, changes in the physical environment can confirm that a recording took place subsequent to a given date. Where there is no significant change, even slightly different angles in the aerial view can provide additional insight helpful to confirming the location of an event, especially where Street View is unavailable. Meanwhile, options to create, save, and share annotations and measurements of distances – whether the width of a crowd's exit route or the approximation of a building's shadow to help gauge time of day – facilitate collaborative work and clear documentation.

Yet this reliance on satellite imagery brings us back to questions of how human rights researchers relate to other investigative motivations. It is important to recognise continuities even within fast-evolving methods, and to reflect upon how they may affect both the tools we use and the way we frame and conceptualise issues. While we cannot take responsibility for the purposes to which others use the same tools, it is worth reflecting on what they were initially designed for. For instance, Google Earth owes its main functions to Keyhole EarthViewer, a mapping software backed by the CIA's venture capital fund In-Q-Tel before Google acquired Keyhole, Inc., in 2004 (Garfield 2015; Levine 2018). More broadly, much of Silicon Valley's prowess is rooted in the federally-funded development of technologies for electronics and communications for defence purposes during the Cold War and Second World War (O'Mara 2019). StreetView faces significant criticism for privacy violations (Zuboff 2019). It is incumbent upon researchers to pursue awareness about technological power and global positionality, raising questions about the reasons for the quality and frequency of imagery in certain areas and not others. Researchers should be wary that the novelty of digital techniques does not displace reliance on invaluable local knowledge. The techniques that human rights practitioners rely on in the digital age are not always as new as we might first assume; nonetheless, perhaps what can be innovative are the ways that we question and diversify their dominant usages.

Final Reflections and Ongoing Dilemmas

Amassing this much evidence exposed disturbing patterns about the use of tear gas worldwide. Above and beyond a single country or police force, tear gas is regularly used in ways that it never should be – with few repercussions. Though billed as a safer alternative to other weapons used by security forces, these eyewitness recordings together indicate that instances of its violent or disproportionate use against crowds are by no means unique to a single force, but rather that a lack of training is widespread and that tear gas is systemically used in illegitimate ways.

On a broader level, this also is a testament to the immense courage of civilians in seeking accountability through their documentation of those abusing their power. The video clips spotlighted on Amnesty's website reflect how social media can be employed as a valuable repository of eyewitness content. Though social media platforms are plagued with fundamental problems in their management of issues ranging from content moderation to advertising and data privacy, our discussion of these issues should not exclude a recognition of the different purposes to which users may post and of the significance of public expression for users who may be disempowered by other means and in different contexts; those who experience human rights violations are not passive victims but have stories to tell.

After taking hundreds of individual videos through these processes, a selection of those with confirmed locations, dates, and content were visualised within a map, an online public resource to raise the visibility of human rights violations in a digital age. Amnesty's Webby award-winning wider multimedia site helps to make expert knowledge transparent and accessible, including interactive diagrams, videos, and text about the health consequences of tear gas exposure, the chemical components involved, current legislation, and the role of unregulated manufacturing. Altogether the site seeks to help users instantly comprehend the geographical spread and make sense of information they otherwise may not come across, or may assume involved an isolated occurrence. Yet, perhaps inevitably, preparing this involves something akin to an editorial vision; not only is selective filtering of data required to conduct an investigation, but any public presentation of conclusions involves parsing people's stories through a particular lens. At the very least, investigators can set a precedent of transparency by explicitly offering their motivations – in this case, to campaign for the regulation of the manufacture, trade, and use of tear gas as well as to increase visibility of some of the ways that it is misused and the experiences of those at the receiving end. Original sources should be acknowledged and linked, to allow recognition of the authors as well as to avoid decontextualisation. Yet seeking consent of actors involved is more challenging; in particular, in these high-pressure and rapidly-escalating environments, the individual recording an abuse does not have time to consult those depicted, and often does not know them. In what ways can investigators ensure informed consent, and preserve the anonymity of civilian participants while maintaining the integrity of civilian footage? Meanwhile, where content is public and open source, is it unethical to keep any meaningful collation of it private?

Difficult issues around value-judgements must not be avoided, in our own and similar projects. Sometimes we encountered clips with no identifiable geographic context, or context that multiple investigators working for hours were unable to locate. Sometimes clips depicted harrowing distress, violence, or pain inflicted that did not constitute illegal police activity and as such did not meet Amnesty's criteria for inclusion. Researchers must be wary of broader assumptions; though the capacity to record abuses of authority is often portrayed as empowering and democratising, vast imbalances in terms of bandwidth, latency, and digital literacy complicate any characterisation of possession of a smartphone as inherently equalising (Mawere and Van Stam 2020). Undertaking verification work not only requires making difficult calls but itself involves implicit understandings about authority and veracity, and often follows political and media narratives that assume an understanding of truth based on proof, and particular to parts of the global North. Does holding the powerful to account necessitate the uncomfortable tension of shaping advocacy according to the investigators' perspectives rather than those of people experiencing disempowerment? What alternatives can investigators offer? It is hard to avoid these questions in digital verification work – and this is its key strength. Beyond the promising approaches that these techniques offer for investigative practice, recognition of the tensions involved and careful, inclusive negotiation of them may just help to push others – and, crucially, ourselves – to a higher standard of digital ethics.

Bibliography

Amnesty International. 2019a. "Syria: Unprecedented Investigation Reveals US-led Coalition Killed more than 1600 Civilians in Raqqa 'Death Trap.'" Accessed March 15, 2021. https://www.amnesty.org/en/latest/news/2019/04/syria-unprecedented-investigation-reveals-us-led-coalition-killed-more-than-1600-civilians-in-raqqa-death-trap/.

Amnesty International. 2019b. "War in Raqqa: Rhetoric Versus Reality." Accessed March 8, 2021. https://raqqa.amnesty.org/.

Amnesty International. 2020a. "Black Lives Matter Protests: Mapping Police Violence Across the USA." Accessed March 8, 2021. https://www.amnesty.org/en/latest/news/2020/06/usa-unlawful-use-of-force-by-police-at-black-lives-matter-protests/.

Amnesty International. 2020b. "Eyes on Chile: Police Violence and Command Responsibility During the Period of Social Unrest." Accessed March 8, 2021. https://www.amnesty.org/en/latest/research/2020/10/eyes-on-chile-police-violence-at-protests/.

Amnesty International. 2020c. "Tear Gas: An Investigation." Accessed March 8, 2021. https://teargas.amnesty.org/.

Aronson, J. D. 2018. "The Utility of User-Generated Content in Human Rights Investigations." In New Technologies for Human Rights Law and Practice, edited by Land, M. K. and Aronson, J. D., 129–148. Cambridge, UK: Cambridge University Press.

Crider, C. 2020. "Why I'm so Keen to Talk to Facebook Content Moderators." Foxglove, July 22). https://www.foxglove.org.uk/2020/07/22/why-im-so-keen-to-talk-to-facebook-content-moderators/.

Dubberley, S. and Grant, M. 2017. "Journalism and Vicarious Trauma: A Guide for Journalists, Editors and News Organisations." First Draft News, April. Accessed March 8, 2021. https://firstdraftnews.org/wp-content/uploads/2017/04/vicarioustrauma.pdf.

Garfield, L. 2015. "The CIA's EarthViewer was Basically the Original Google Earth." Business Insider, December 30. Accessed March 8, 2021. https://www.businessinsider.com/the-cias-earthviewer-was-the-original-google-earth-2015-11?r=US&IR=T.

Levine, Y. 2018. "Google's Earth: How the Tech Giant is Helping the State Spy on Us." The Guardian, December 20. Accessed March 8, 2021. https://www.theguardian.com/news/2018/dec/20/googles-earth-how-the-tech-giant-is-helping-the-state-spy-on-us.

Mawere, M. and van Stam, G. 2020. "Data Sovereignty: A Perspective From Zimbabwe." WebSci '20: 12th ACM Conference on Web Science Companion: 13–19. doi:10.1145/3394332.3402823.

McAvoy, L. 2021. "Centering the 'Source' in Open Source Investigation." Open Global Rights, January 21. Accessed March 8, 2021. https://www.openglobalrights.org/centering-the-source-in-open-source-investigation/.

Minogue, D. and Makumbe, R. P. 2019. "Digital Accountability Symposium: Harnessing User-Generated Content in Accountability Efforts for International Law Violations in Yemen." OpinioJuris, December 18. Accessed March 16, 2021. http://opiniojuris.org/2019/12/18/digital-accountability-symposium-harnessing-user-generated-content-in-accountability-efforts-for-international-law-violations-in-yemen/.

Ng, Y. 2020. "How to Preserve Open Source Information Effectively." In Digital Witness: Using Open Source Information for Human Rights Investigation, Documentation, and Accountability, edited by Dubberley, S., Koenig, A., and Murray, D., 143-164. Oxford, UK: Oxford University Press.

O'Mara, M. 2019. The Code: Silicon Valley and the Remaking of America. New York, NY, USA: Penguin Press.

Piracés, E. 2018. "Collecting, Preserving and Verifying Online Evidence of Human Rights Violations." Open Global Rights, January 30. Accessed March 15, 2021. https://www.openglobalrights.org/collecting-preserving-and-verifying-online-evidence-of-human-rights-violations/.

Stottlemyre, S. A. 2015. "HUMINT, OSINT, or Something New? Defining Crowdsourced Intelligence." International Journal of Intelligence and Counterintelligence 28: 578–589. doi:10.1080/08850607.2015.992760.

Sweeney, E. 2020. "OSINT Essentials." Accessed March 8, 2021. https://www.osintessentials.com.

University of Amsterdam. n.d. "DMI Tools." Accessed March 8, 2021. https://wiki.digitalmethods.net/Dmi/ToolDatabase.

Williams, H. J. and Blum, I. 2018. "Defining Second Generation Open Source Intelligence (OSINT) for the Defence Enterprise." RAND Corporation. Accessed March 15, 2021. https://www.rand.org/content/dam/rand/pubs/research_reports/RR1900/RR1964/RAND_RR1964.pdf.

Witness Media Lab. 2019. "Prepare, Don't Panic: Synthetic Media and Deepfakes." Accessed March 18, 2021. https://lab.witness.org/projects/synthetic-media-and-deep-fakes/.

Zuboff, S. 2019. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. London, UK: Profile Books.