Centre for Emerging Technology and Security
EXPERT ANALYSIS

Artificial Intelligence, OSINT and Russia's Information Landscape

Charlie Winter, John Gallacher
and Alexander Harris

The views expressed in this article are those of the authors, and do not necessarily represent the views of The Alan Turing Institute or any other organisation.

Introduction

Over the last year, the Russian invasion of Ukraine has brought the value of open-source intelligence (OSINT) – the structured collection and analysis of publicly available information to achieve a targeted investigative outcome[i] – into sharp relief.

This article explores the uses of AI and data science for analysing open-source intelligence related to Russia's war in Ukraine – and considers how future capabilities can be leveraged most effectively.

From investigators tracking the military build-up[ii] before the invasion was officially announced to ongoing efforts to map the conflict,[iii] uncover war crimes[iv] and identify hostile information operations online,[v] OSINT has repeatedly demonstrated its utility.

At the tactical level, the Ukrainian military is reported to have used Instagram and TikTok content to locate Chechen forces.[vi] Online groups – both civilian activists and formalised organisations – have been tracking Russian Navy deployments using publicly available satellite imagery and estimating total Russian military hardware losses from images posted to social media.[vii]

Due to the vast amount of data available via public sources, OSINT collection and analysis has typically been a time- and resource-heavy process. However, increased use of data science and machine learning is substantially improving the efficiency and scale of such investigations. Crowd-sourced reporting of enemy troop movements,[viii] especially when combined with increasingly precise computer vision models[ix] and enhanced availability of commercial satellite imagery,[x] has improved situational awareness and allowed for near real-time tracking of the conflict. Geographic information system (GIS) software has helped identify areas of Ukraine littered with unexploded ordinance to prioritise for de-mining.[xi] Other notable advances include the extraction of information from swathes of unstructured data such as text,[xii] the deployment of computer vision models on commercial satellite imagery to provide situational awareness,[xiii] and the use of statistical modelling and simulations to forecast events under multiple scenarios.

Machine-augmented techniques have also radically altered how intelligence work is done within the information battlefield. In this context, advances in scaled data access, network science and analytical techniques have enabled OSINT researchers to trace the contours of malign influence campaigns more accurately and dynamically – including, crucially, when the principal focus of these campaigns is not Western audiences. This allows for more targeted, tailored and impact-driven interventions.

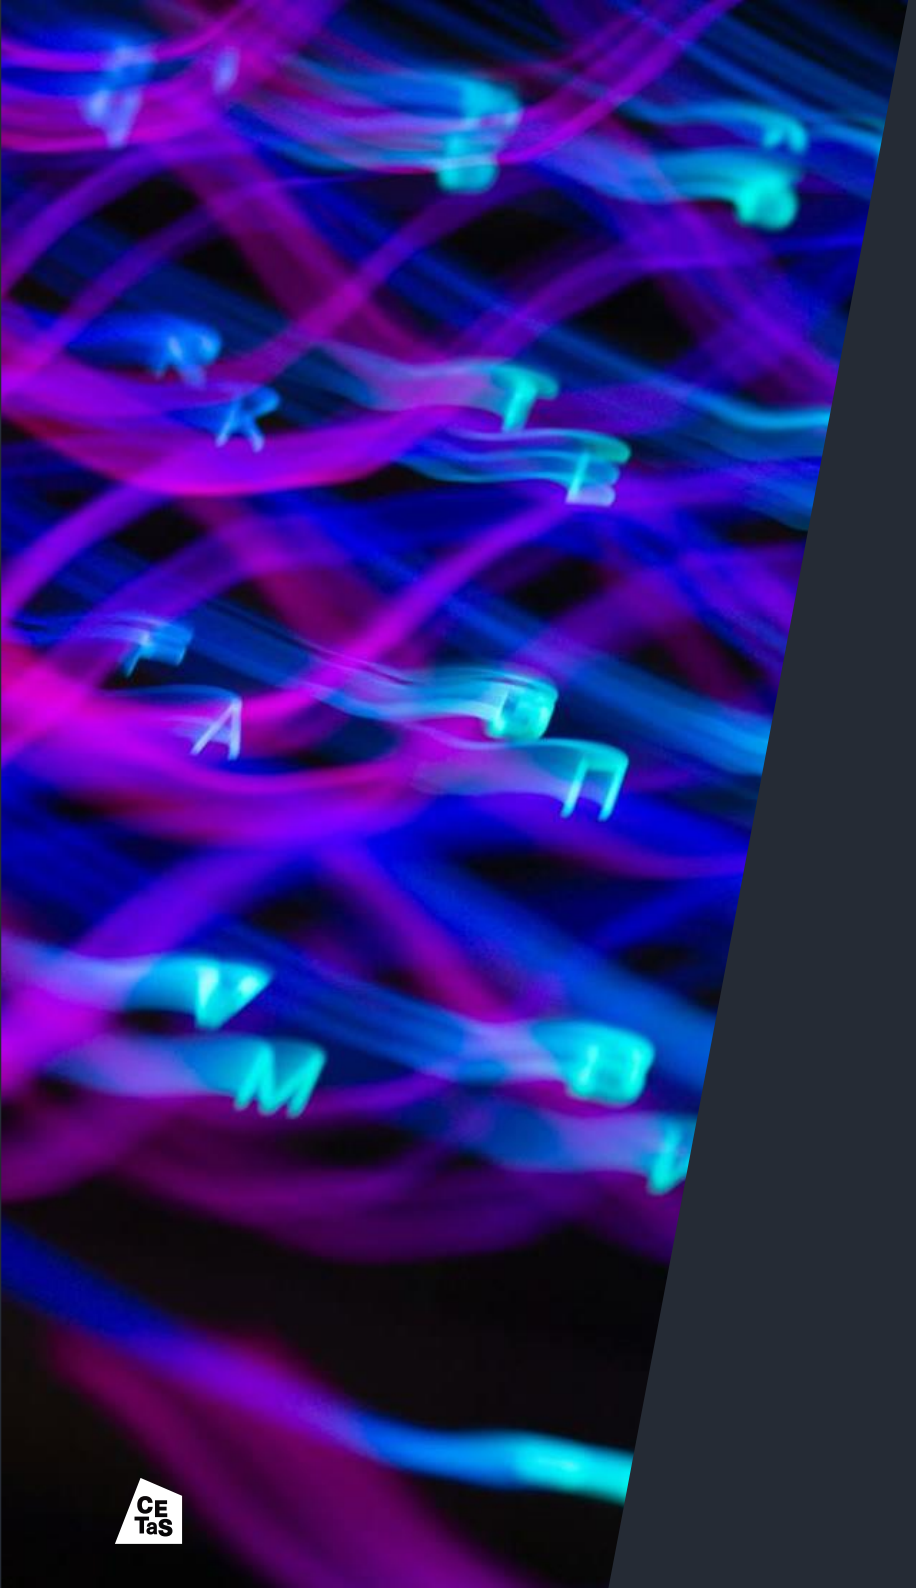

A shifting information landscape

While the strategic deployment of information operations has long been a core component of Russian hybrid warfare,[xiv] exploitation of information on the part of both state and state proxy actors has been supercharged within the context of the Russian invasion of Ukraine.[xv]

This was clear at the outset of the invasion, which saw a strategic information campaign unfold directly in line with Russia's war aims. Importantly, it was targeted in such a manner that it did not correspond with what have conventionally been assumed to be the Kremlin's adversarial priorities[xvi] – audiences based within NATO countries. Instead, this campaign was inwards-looking and focused on pro-war narratives in Russia and eastern Ukraine.[xvii]

In early 2022, the already popular instant messaging platform Telegram[xviii] benefited from a post-war information dynamic that saw access to mainstream social media platforms severely restricted for Russians. In the wake of the Kremlin's efforts to block Western-owned platforms in March 2022, Telegram's daily audience in Russia grew from 25 million people in January to 41.5 million people in July,[xix] with an additional 13 million users joining between February and March alone. This means that Telegram now ranks alongside long-dominant platforms like VKontakte[xx] in both usage and popularity.

From a functionality perspective, Telegram has been developed in such a way that it is ideal for information consumption and distribution.[xxi] Users and media organisations can share large amounts of bandwidth-heavy content rapidly and mono-directionally via channels, groups and super-groups – the latter two of which effectively act as 'discussion boards' where users can directly interact.

A shifting information landscape

A recent DFRLab study found that from July to September 2022, Russian users were sending more traffic to Telegram than any other national population. In August 2022, two of the five most popular Russia-based search topics on Telegram were 'news' and its Russian equivalent, 'новости'. The same study also demonstrated that nine of the ten most popular political Telegram channels among Russian speakers are explicitly Kremlin-aligned, and all of them had been found to amplify pro-war propaganda and dis/misinformation. Importantly, just three of these nine channels (which together preside over millions of subscribers) are official communications feeds for Russian political figures – specifically, the Chechen leader Ramzan Kadyrov, Kremlin propagandist Vladimir Soloviev, and former Prime Minister Dmitry Medvedev. The other six are putatively supporter-run news aggregators and content distribution hubs operated by pro-Kremlin Russian nationalists.

This rapid growth in alleged supporter-run and supporter-consumed Telegram channels represents a trove of open-source data which can provide valuable insights into changes in Russian sentiment and attitudes towards the war. One such system enabling this is ExTrac,[xxii] which operates by tracking real-time and historical conflict actor content and chatter online. To maximise relevance and minimise noise, it relies on human-in-the-loop curation which only processes and analyses data from analyst-selected sources. Sources are only included when they are assessed to be explicitly aligned with the conflict actor in question – in this example, aligned groups include the Russian State, Russian Armed Forces, Ukrainian Separatist organisations and proxy movements.

While individual components of Russia's information campaigns could previously be identified through manual analysis[xxiii], it was only with the application of machine-augmented tools that available open sources could be identified and analysed at scale to reveal a fuller picture of the shape and strategic intent of the Kremlin's activities and objectives in the information space. However, an authoritative view of strategic intent is only really possible when OSINT is *combined* with penetrative secret intelligence, rather than used in isolation.

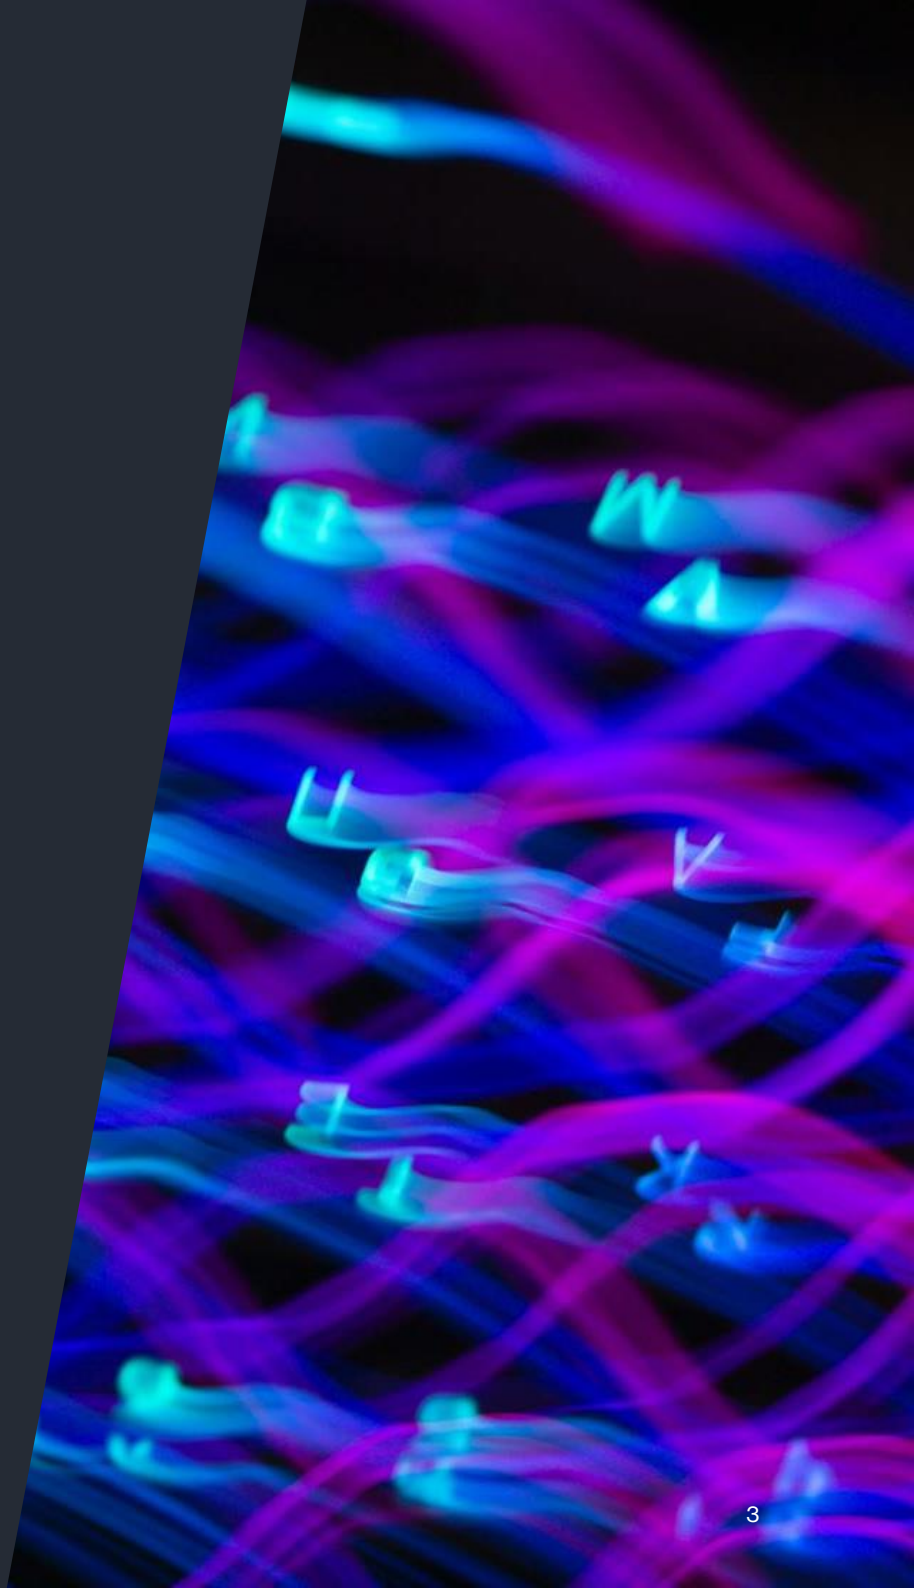

Analysis

ExTrac's data ingests were initially trained on a constellation of 427 Telegram groups and channels, including separatist group discussion boards like 'Donbassr', OSINT-focused feeds like 'Rybar',[xxiv] and paramilitary media hubs such as the Wagner-linked 'Reverse Side of the Medal'.[xxv]

Figure 1 shows all posting activity on ExTrac's pro-Russia Donetsk and Luhansk-based feeds since January 2019. It indicates that between January 2019 and January 2022, the 33 channels and groups in question generally posted between 100 and 500 messages each day. Across the first six weeks of 2022, their posting activity gradually – and seemingly organically – increased. This is perhaps logical given the mounting prospect of war during that period. But on 17 February 2022 there was an unprecedented and demonstrably inauthentic surge in posting behaviour, that saw a 15-fold increase in collective activity over the days that followed. Notably, this surge in activity commenced four days *before* Putin's Security Council address on 21 February, when he announced that Russian troops were to engage in 'peacekeeping operations' in Donetsk and Luhansk (and a week before the full-scale invasion was declared).

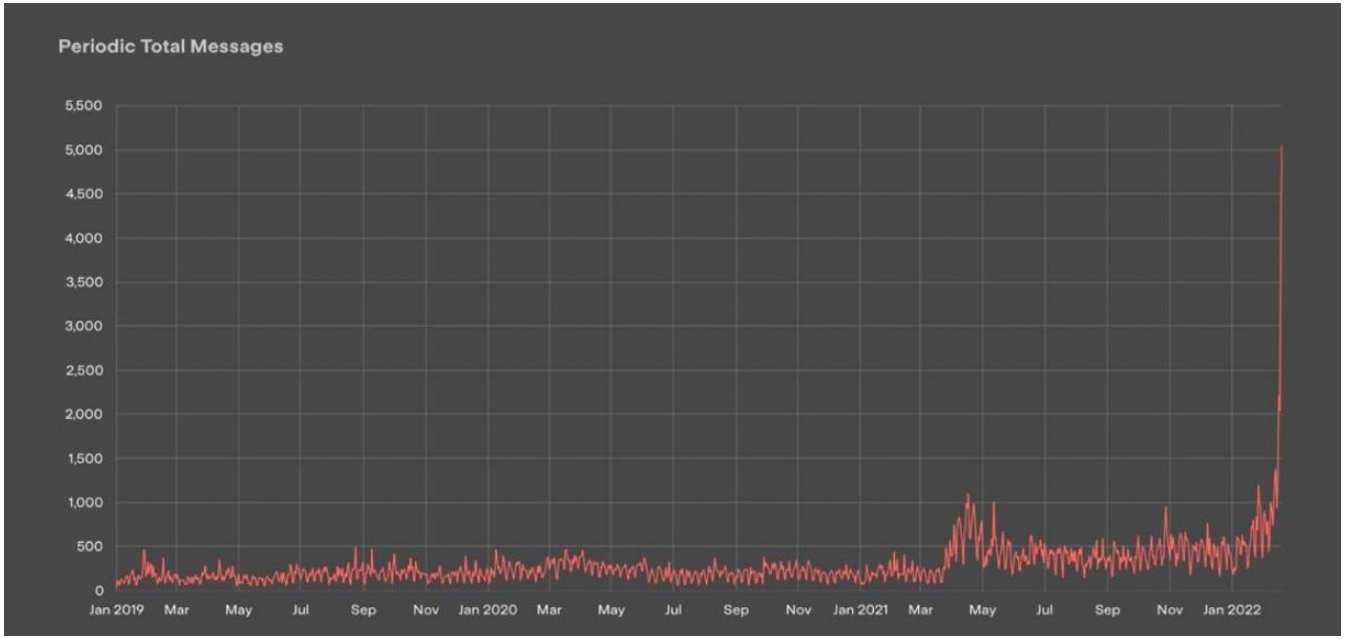

Figure 1. Posting activity of 33 pro-Kremlin LPR- and DPR-based groups on Telegram, January 2019 to February 2022.

Analysis

Figure 2 overlays Ukrainian Ministry of Defence data on Russian ceasefire violations in the eastern territories atop this pro-war chatter. It indicates that this surge started on the very same day that there was a 10-fold increase in ceasefire violations by Russian and pro-Russia forces.

When the data parameters are expanded to include other core components of the Kremlin-aligned media ecosystem, in addition to the Donetsk and Luhansk-focused/based feeds, this surge dynamic is even more pronounced. As seen in Figure 3, the Kremlin-aligned channels and groups sampled – which typically shared 6,000 posts per day from January 2019 to January 2022 – became 8x more productive within the space of a few days. This was communications hyperactivity of a scale not before seen in the context of this network and – as it turns out – a dynamic that has not yet been repeated.

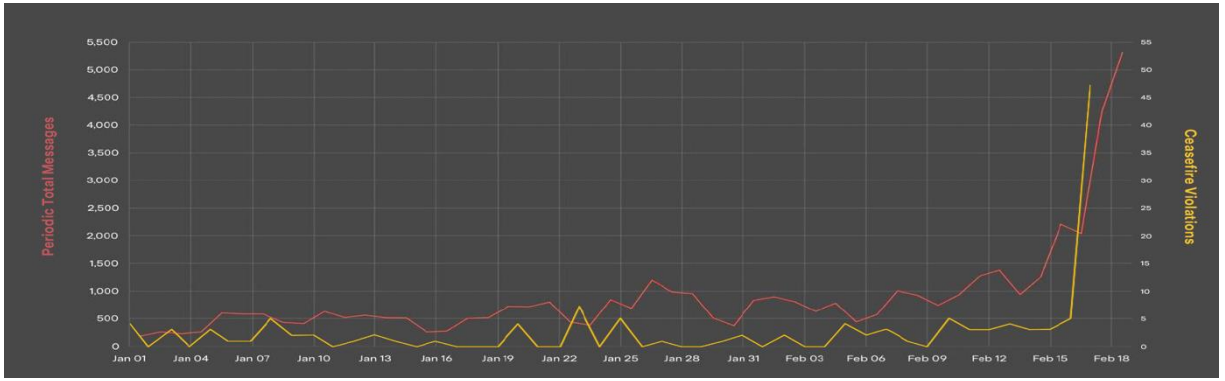

Figure 2. DPR/LPR comms activity (red) and Russian/pro-Russian ceasefire violations (yellow) 1 January 2022 to 18 February 2022.

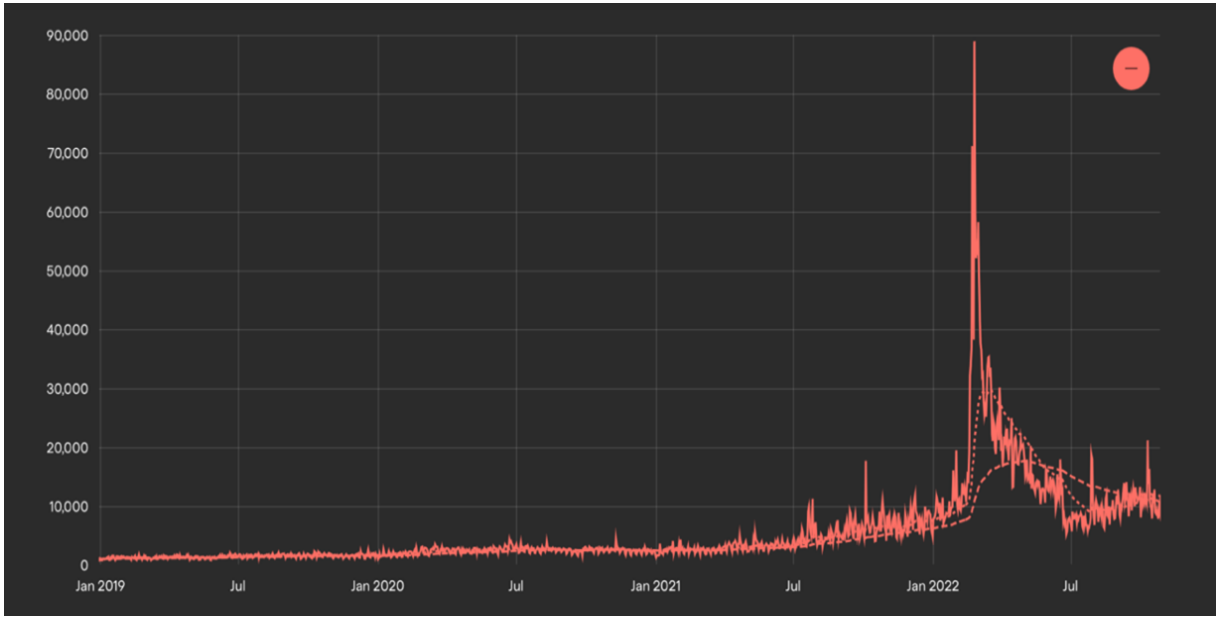

Figure 3. All posting activity on Kremlin-aligned ecosystem sample, January 2019 to April 2022.

Analysis

This escalation was driven by a combination of behaviours. The data indicates it resulted from a sudden increase in both unofficial and official pro-Russia chatter and was driven by a wave of newly produced pro-war multimedia content.

This activity, which was led more by pro-Kremlin influencers than it was by the Kremlin itself, was in lockstep with Putin's soon-to-be-stated war aims, and likely to be in part directly choreographed with a view to supporting Moscow's broader military objectives in Ukraine.

This semi-decentralised approach to influence activity has characterised the Russian war communications effort. It has seen pro-war influencers like Alexander 'Sasha' Kots being effectively deputised[xxvi] by the Kremlin, emerging as its first line of both defence and offence on the information battlefield. When the invasion was in its earliest stages, they set the triumphalist tone, and, when progress began to slow, they furnished audiences with explanations geared towards shoring up support for the flagging 'special operation'.

Perhaps most crucially, at points of inflection in the conflict, whether strategic setbacks like the routing of Russian forces at Hostomel or atrocities like Bucha or Kramatorsk, it fell to this community to navigate the rest of the ecosystem through times of tribulation.

The digital response to the Bucha massacre[xxvii] is an example of this approach. When revelations about the killings first appeared, a limited carefully defined line[xxviii] emerged and evolved within the pro-war Kremlin-aligned online ecosystem. This was a line that had been cultivated by a handful of influencers but propagated to an audience of millions within minutes – and, critically, hours before the Russian state or armed forces had weighed in. It started with denialism but – as more information emerged – shifted to allegations of staging and ultimately a claim that the killings were the result of a NATO-engineered false flag operation.

The use of machine-augmented OSINT made it possible to track this evolution in real-time as it was introduced into the ecosystem. The data was collected, processed and analysed as the narrative itself was being set among Kremlin apologists. Key nodes in the network could be identified and the nature of their ability to interface with official and overt Kremlin information infrastructure assessed.

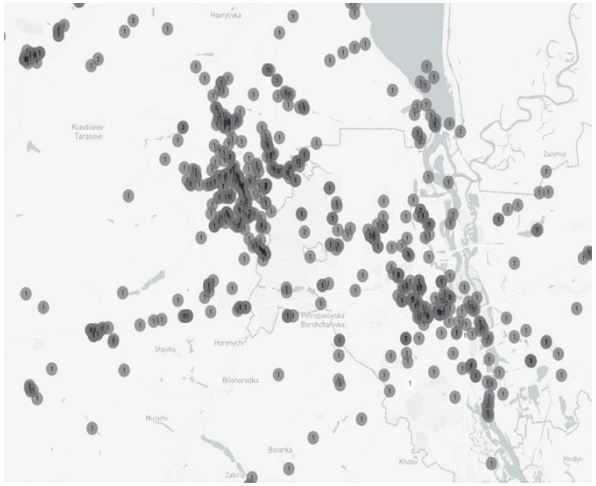

Figure 4. Conflict events detected around Bucha and Kyiv through ExTrac's event detection model. January 2022 to January 2023.

Future capabilities

The analysis above demonstrates the value of data science techniques within OSINT analysis workflows. As these technologies continue to advance – and become increasingly accessible for non-technical users[xxix] – they will enable practitioners to perform a wider range of tasks more quickly and accurately, as well as opening up new capabilities, such as the identification of patterns and trends that would be impossible for humans to detect manually. **Three key areas where machine learning will augment OSINT are natural language processing (NLP), computer vision, and forecasting / prediction.**

Natural Language Processing (NLP)

Large language models can be used to help analysts understand and process large amounts of text-based information, such as that found on social media or within unstructured datasets. For instance, OpenAI's suite of GPT language models can format unstructured data[xxx] into an easily digestible tabular format. More recent models have shown good performance at generating short summaries from book-length texts, which could potentially save hours of analyst time. Question-answering models[xxxi] allow analysts to ask natural language queries over a large body of unstructured text data and receive relevant and accurate answers, making it easier to find information without having to sift through large volumes of text manually. However, many of these large language models are trained on relatively well-written sources, such as books, journals and news articles. For these models to perform well within the social media domain, specific training (and agile retraining) on emojis, memes, misspelling and slang will likely be required.

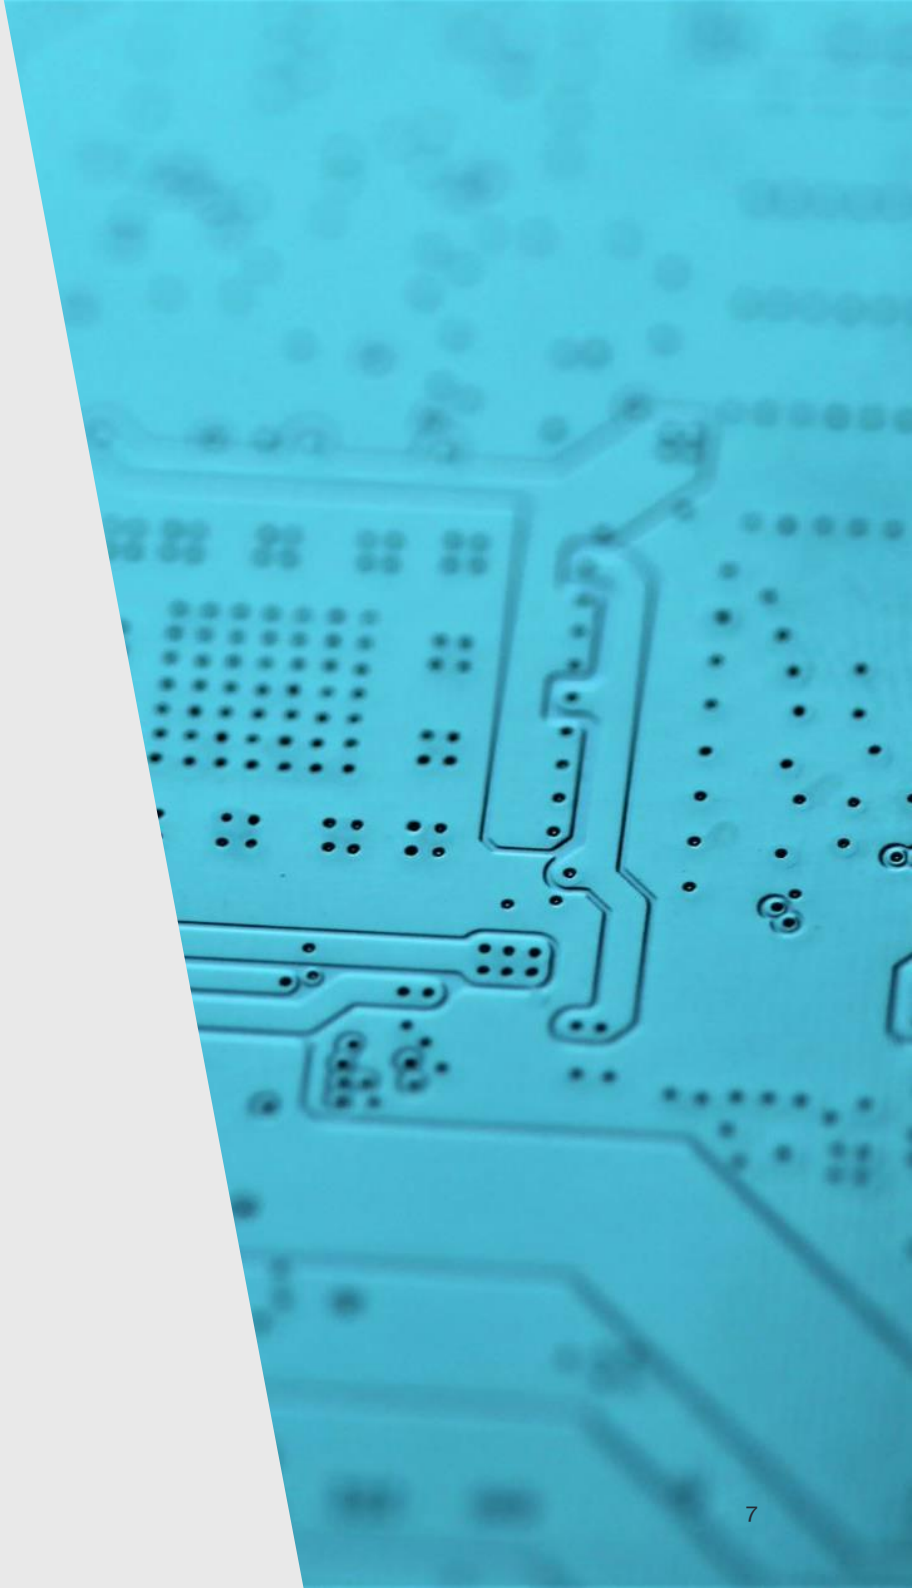

Future capabilities

Computer Vision

Computer vision models serve as valuable tools for OSINT analysts by automatically analysing and interpreting visual data. For example, image classification and object detection models can be used to identify images of weapons within a large collection of images,[xxxii] the type of weapon presented, or the frame where a weapon is visible within a video. Applying computer vision models to commercial satellite imagery has helped to track conflict damage to civilian infrastructure;[xxxiii] similar approaches can be used to assess damage caused by natural disasters[xxxiv] and inform emergency response.

The latest generation of large pre-trained computer vision models[xxxv] have also shown impressive zero-shot capabilities for image geolocation[xxxvi] (identifying where a photo was taken without any additional training data) and chronolocation (identifying when the photo was taken).[xxxvii] These capabilities are not yet at the level of human performance, but they offer scalability beyond human means. For instance, they could be used to generate time and location estimates for a large imagery dataset which would be too time-consuming for a human analyst to process – this triaging process could then be used as a first step to narrow down an analyst's manual search.

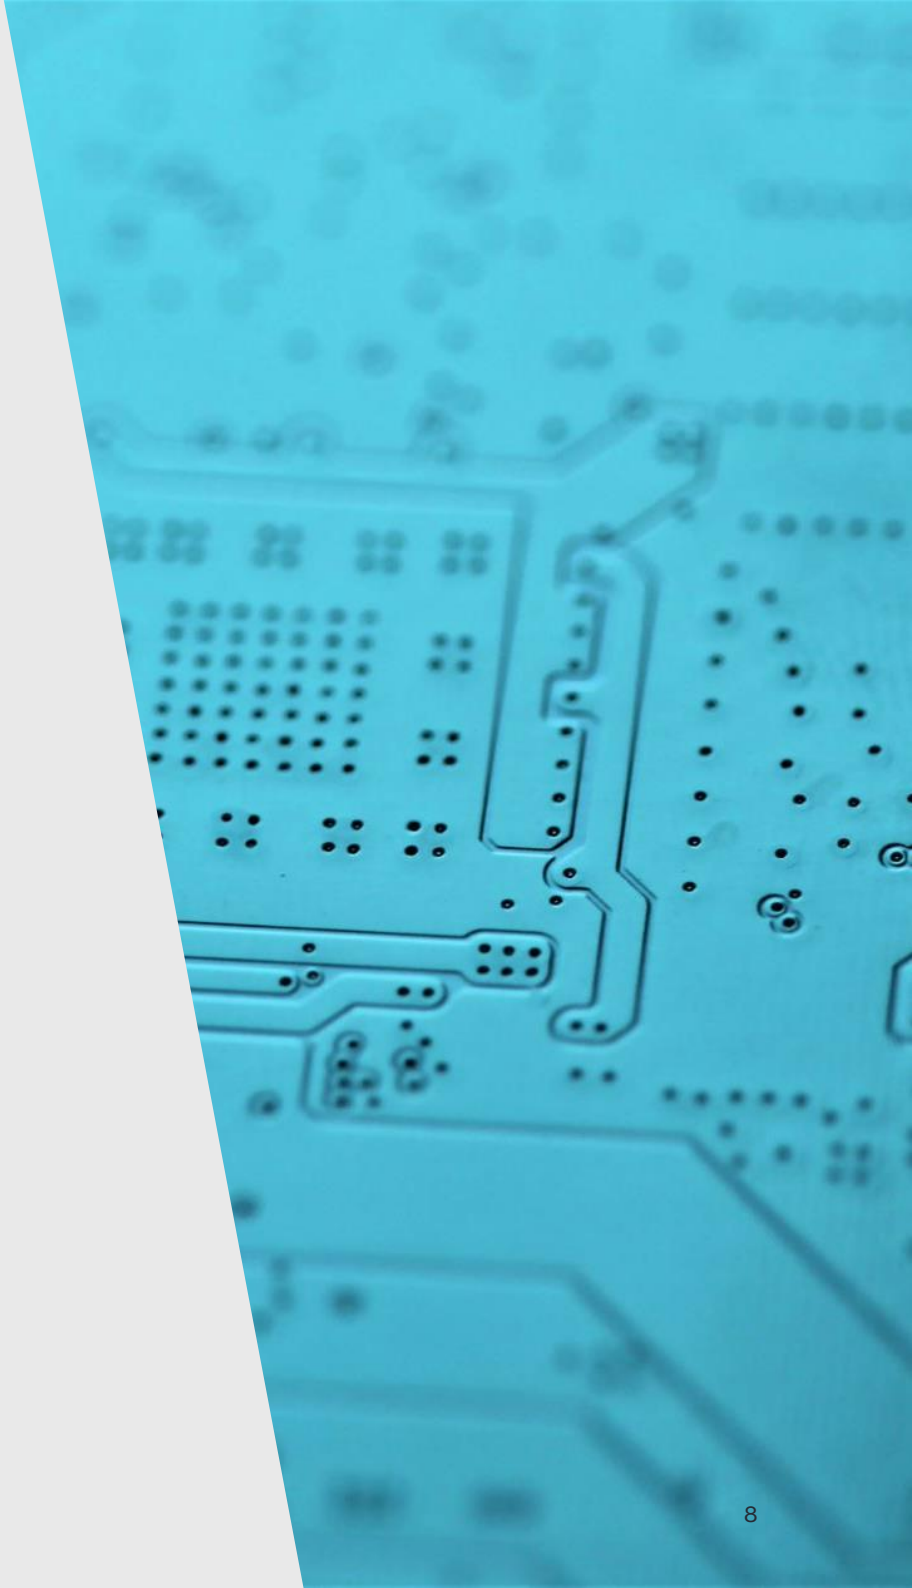

Future capabilities

Forecasting and prediction

Advances in machine learning and data science have also enabled data-driven predictions and the generation of forecasts using public data.

Accurate forecasting relies on both effective data collection and human analysis. In recent years, researchers have been experimenting with data collected by the Armed Conflict Location & Event Data Project (ACLED),[xxxviii] which compiles publicly available data on conflict violence. This dataset has been used as a basis for experimental early-warning tools for regional and sub-regional conflict, including volatility and risk predictability indices. By applying data science techniques, researchers have attempted to forecast the number of fatalities in impending state-based conflicts up to three years in advance.[xxxix] Notably, while these efforts are methodologically interesting and have yielded some interesting results, they are impossible to validate until the time-period or conflict in question has come to pass – so uncertainty must be embraced and communicated accordingly in relation to their results.

Another approach that may hold potential relates to human-driven, machine-enabled forecasting of tactical or strategic shifts within an ongoing conflict. Among these have been short- and medium-term predictive analyses building on data relating to the targeting of civilians by Wagner mercenaries in the Central African Republic and Mali,[xl] and – though not directly conflict-related – recent work exploring the likelihood of regime change in Iran based on historical and current protest trends.[xli] Research has also explored the use of human expertise, data science and machine learning techniques to predict offline extremist violence at the level of individual events – for example, ISKP's attack on Kabul airport in August 2021[xlii] – by analysing publicly available data from social media.[xliii] The use of data analytics to reliably forecast this kind of human behaviour is extremely challenging and should be approached with some scepticism.

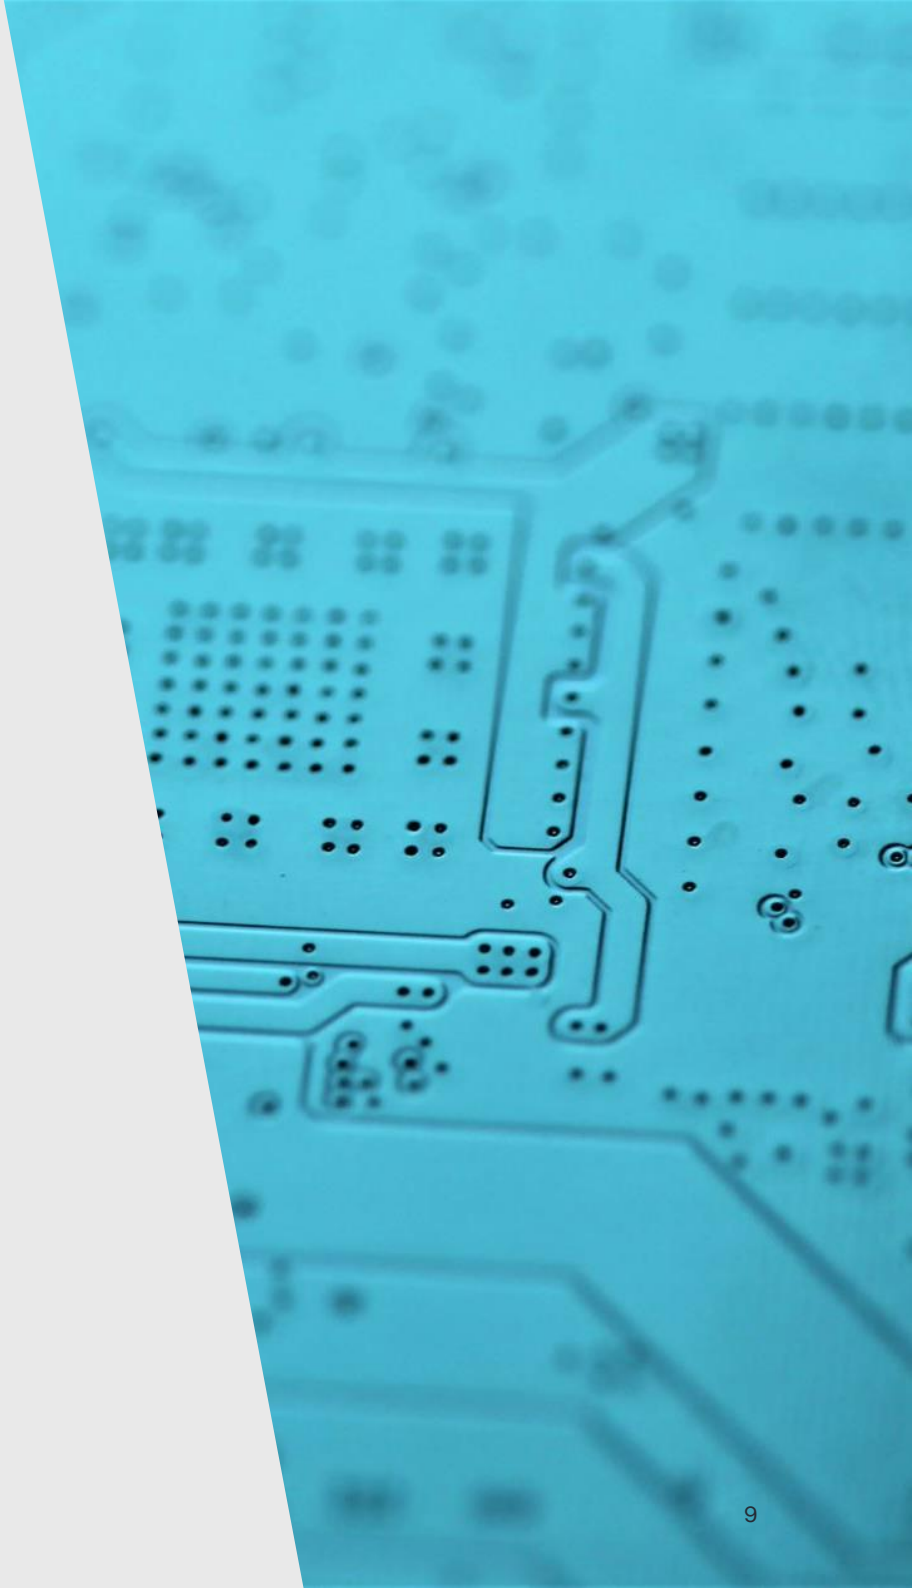

Conclusion

The last 12 months of conflict in Ukraine have continually highlighted the importance and impact of OSINT within wider intelligence and security analysis.

Developments in the use of machine learning and data science are certain to continue to accelerate in months and years to come, leaving analysts increasingly able to expedite and augment time-consuming and incrementally more complex tasks.

Recent estimates from the UK's Defence Intelligence (DI) suggest that open-source analysis contributes to approximately 20% of their current capability with 80% coming from closed sources – General Sir James Hockenhull, UK Commander of Strategic Command, expects this ratio to reverse in the coming years.[xliv] Crucially, the West's adversaries and competitors are also following this same trajectory and collaboration across sectors will be vital to maintain an advantage. A reversal as aggressive as this will require a step change in risk calculation when it comes to OSINT. The opportunity costs of not moving fast enough can quickly outweigh the benefits of adhering to fixed legislative policies designed in a previous era.

Conclusion

In many cases, the barriers to progress are not technological. To best exploit OSINT, government must foster a more flexible approach to intelligence – engaging with a wider range of partners in new and agile ways. In practical terms this will involve closer collaboration between government and the private and non-profit sectors, allowing for the two-way sharing of intelligence requirements, methodological developments and data. If correctly leveraged, this type of dynamic engagement would benefit both government and the open-source community. A unified and enabling cross-sector approach will be fundamental to help mitigate existing barriers and maximise the utility of any technical innovations as they emerge.

About the Authors

Dr Charlie Winter *is Director of Research at the conflict intelligence platform ExTrac. He is also an Associate Fellow at the International Centre for Counter-Terrorism and a member of the RESOLVE Network Research Advisory Council.*

Dr John Gallacher *is the founder and director of Alethio, a tech startup building machine learning tools for open-source intelligence. Prior to this, he worked as a research scientist for the UK Government, served with the British Army, and worked for NATO. He is also a visiting researcher at The Alan Turing Institute's Defence and Security Programme and holds a PhD from the University of Oxford.*

Alexander Harris *is a Research Associate in The Alan Turing Institute's Defence and Security Programme. Alex works across a research portfolio exploring the application of data science and artificial intelligence to a range of national security and defence challenges, working closely with national and international partners.*

References

i Harris, A., Janjeva, A. and Byrne, J. 2022. 'The Future of Open Source Intelligence for UK National Security.' Available at: https://cetas.turing.ac.uk/publications/future-open-source-intelligence-uk-national-security

ii The Economist. 2022. 'Russia's military build-up enters a more dangerous phase.' Available at: https://www.economist.com/interactive/2022/02/11/russias-military-build-up-enters-a-more-dangerous-phase. See also Middlebury Institute of International Studies. 2022. 'On Google Maps, Tracking the Invasion of Ukraine.' Available at: https://www.middlebury.edu/institute/news/google-maps-tracking-invasion-ukraine

iii Bellingcat. 2022. 'The Eyes on Russia Map.' Available at: https://eyesonrussia.org

iv Bellingcat. 2022. 'Tracking the Faceless Killers who Mutilated and Executed a Ukrainian POW.' Available at: https://www.bellingcat.com/news/2022/08/05/tracking-the-faceless-killers-who-mutilated-and-executed-a-ukrainian-pow/

v ExTrac. 2022. 'Evidence of Kremlin campaign to damage perception of Ukrainian refugees.' Available at: https://public-assets.extrac.io/reports/ExTrac_Evidence_of_Kremlin_campaign_to_damage_perceptions_of_Ukrainian_refugees.pdf

vi Joshi, S. 2022. Twitter post [17 December 2022]. Available at: https://twitter.com/shashj/status/1604103063391830017

vii Mitzer, S. and Janovsky, J. 2022. 'Attack on Europe: Documenting Russian Equipment Losses During The 2022 Russian Invasion Of Ukraine.' Available at: https://www.oryxspioenkop.com/2022/02/attack-on-europe-documenting-equipment.html

viii Ministry of Defence, Strategic Command, and General Sir Jim Hockenhull KBE ADC Gen. 2022. 'How open-source intelligence has shaped the Russia-Ukraine war.' Available at: https://www.gov.uk/government/speeches/how-open-source-intelligence-has-shaped-the-russia-ukraine-war

ix Kermode, L., Freyberg, J., Akturk, A., Trafford, R., Kochetkov, D., Pardinas, R., Weizman, E. and Cornebise, J. 2020. 'Objects of violence: synthetic data for practical ML in human rights investigations.' Available at: https://arxiv.org/pdf/2004.01030.pdf

x Agence France-Presse, 2022. 'Before-and-after satellite imagery will track Ukrainian cultural damage, UN says.' The Guardian. Available at: https://www.theguardian.com/world/2022/oct/27/before-and-after-satellite-imagery-will-track-ukraine-cultural-damage-un-says

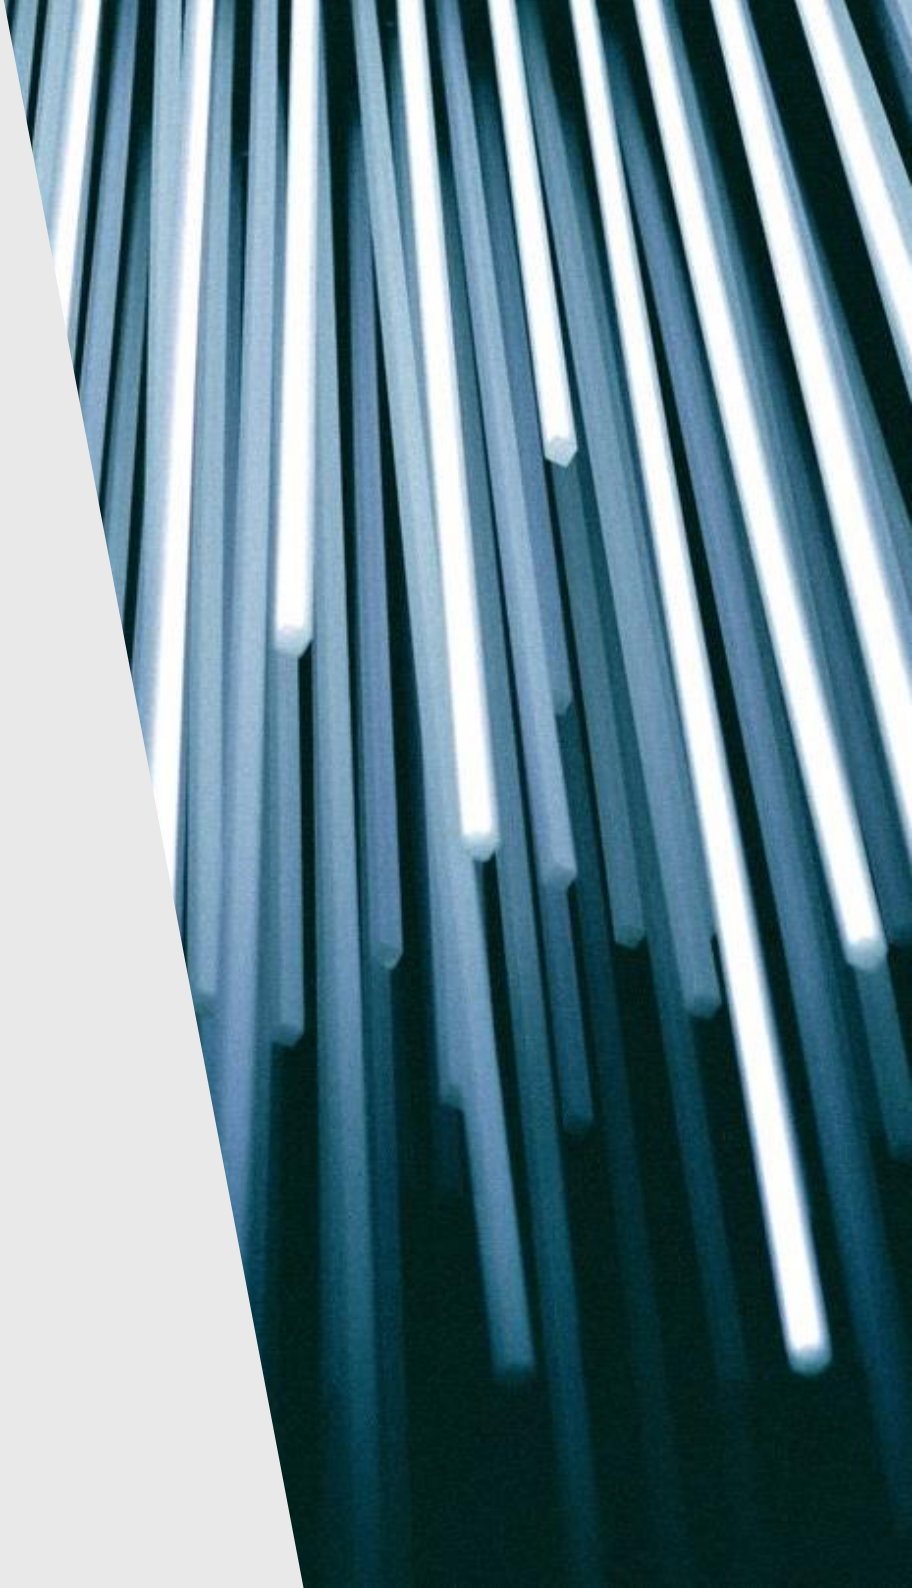

References

xi The Halo Trust. 2022. 'The Halo Trust partners with ESRI.' Available at: https://www.halotrust.org/latest/halo-updates/news/the-halo-trust-partners-with-esri-in-ukraine/

xii Rocha, R. 2022. 'Getting tabular data from unstructured text with GPT-3: an ongoing experiment.' Available at: https://robertorocha.info/getting-tabular-data-from-unstructured-text-with-gpt-3-an-ongoing-experiment/?utm_source=pocket_mylist

xiii Conflict Observatory. N.d. Available at: https://hub.conflictobservatory.org/portal/apps/sites/#/home/. See also Conflict Damage Assessment. N.d. Available at: https://sea.security.copernicus.eu/categories/conflict-damage-assessment/

xiv Marovic, J. 2019. 'Wars of Ideas: Hybrid Warfare, Political Interference, and Disinformation.' Available at: https://carnegieeurope.eu/2019/11/28/wars-of-ideas-hybrid-warfare-political-interference-and-disinformation-pub-80419

xv Batbayar, M. & Deelan, E. 2022. 'What we've learned in Ukraine about combating disinformation.' Available at: https://www.data4sdgs.org/blog/what-weve-learned-ukraine-about-combating-disinformation

xvi Blinken, A. J. 2022. 'Targeting Russian Elites, Disinformation Outlets, and Defense Enterprises.' Available at: https://www.state.gov/targeting-russian-elites-disinformation-outlets-and-defense-enterprises/

xvii Scott, M. 2022. 'As war in Ukraine evolves, so do disinformation tactics.' Available at: https://www.politico.eu/article/ukraine-russia-disinformation-propaganda/

xviii Telegram Homepage. N.d. Available at: https://telegram.org

xix Buziashvili, E. 2022. 'What Russia reads on Telegram.' Available at: https://medium.com/dfrlab/what-russia-reads-on-telegram-a45d90e3bdc1

xx VK Homepage. N.d. Available at: https://vk.com

xxi Herasimenka, A., Bright, J., Knuutila, A. and Howard P. N. 2022. 'Misinformation and professional news on largely unmoderated platforms: the case of telegram.' Available at: https://www.tandfonline.com/doi/full/10.1080/19331681.2022.2076272

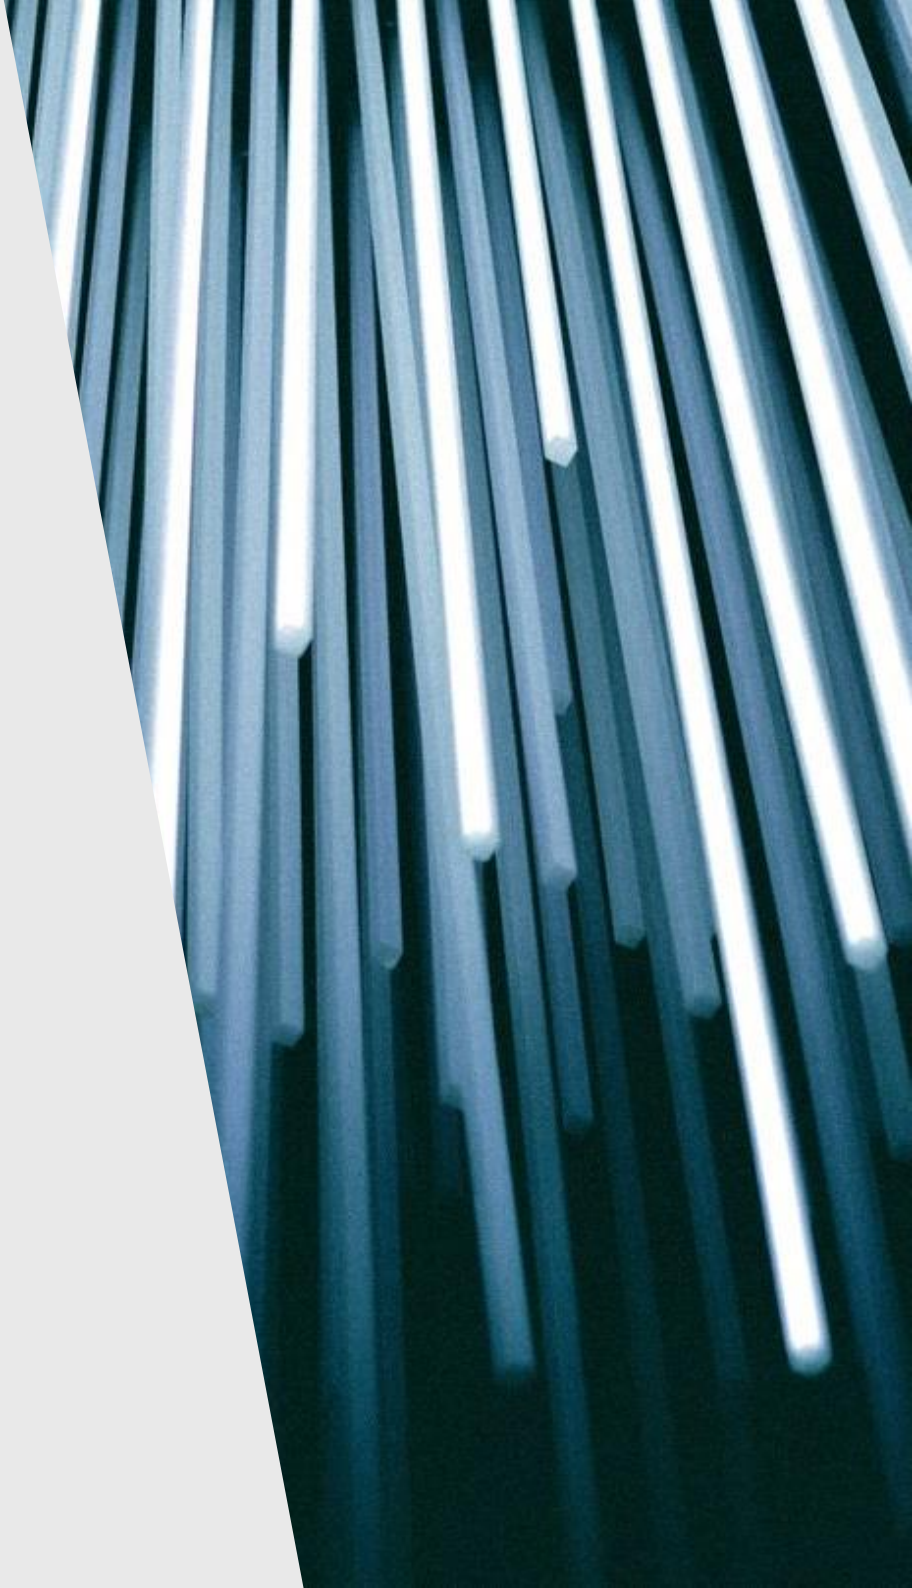

References

xxii Extrac Homepage. N.d. Available at: https://www.extrac.io

xxiii Scott, M. 2022. 'Digital Bridge: Russia's online propaganda war – Meta lobbying targets – Google vs Apple on privacy.' Available at: https://www.politico.eu/newsletter/digital-bridge/meta-lobbying-google-vs-apple-russian-disinformation/

xxiv The Bell. 2022. 'Unmasking Russia's influential pro-war 'Rybar' Telegram channel.' Available at: https://thebell.io/en/unmasking-russia-s-influential-pro-war-rybar-telegram-channel/

xxv Ghazaryan, K. 2020. 'Wagner-Affiliated Telegram Channel Trolls Nagorno-Karabakh Conflict Analysts.' Available at: https://www.bellingcat.com/news/uk-and-europe/2020/10/07/wagner-affiliated-channel-trolls-nagorno-karabakh-conflict-analysts/

xxvi Tchoubar, P. 2022. 'How pro-Russian bloggers are covering the war in Ukraine.' Available at: https://observers.france24.com/en/tv-shows/the-observers/20220503-how-pro-russian-bloggers-are-covering-the-war-in-ukraine

xxvii Al-Hlou, Y., Froliak, M., Khavin, D., Koettl, C., Willis, H. Cardia, A., Reneau, N. and Bronw, M. 2022. 'Caught on Camera, Traced by Phone: The Russian Military Unit That Killed Dozens in Bucha.' Available at: https://www.nytimes.com/2022/12/22/video/russia-ukraine-bucha-massacre-takeaways.html

xxviii ExTrac. 2022. Twitter post [5 April 2022]. Available at https://twitter.com/Ex_Trac/status/1510974341453910017

xxix GitHub. N.d. 'GitHub Copilot - Your AI pair programmer.' Available at: https://github.com/features/copilot

xxx Chiu, M., Roberts, R., and Yee, L. 2022. 'Generative AI is here: How tools like ChatGPT could change your business.' Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/generative-ai-is-here-how-tools-like-chatgpt-could-change-your-business

xxxi Hugging Face. N.d. 'Question answering.' Available at: https://huggingface.co/tasks/question-answering

xxxii DiRenzo, N. 2022. 'Identify Military Vehicles in Satellite Imagery with TensorFlow.' Available at: https://python.plainenglish.io/identifying-military-vehicles-in-satellite-imagery-with-tensorflow-96015634129d

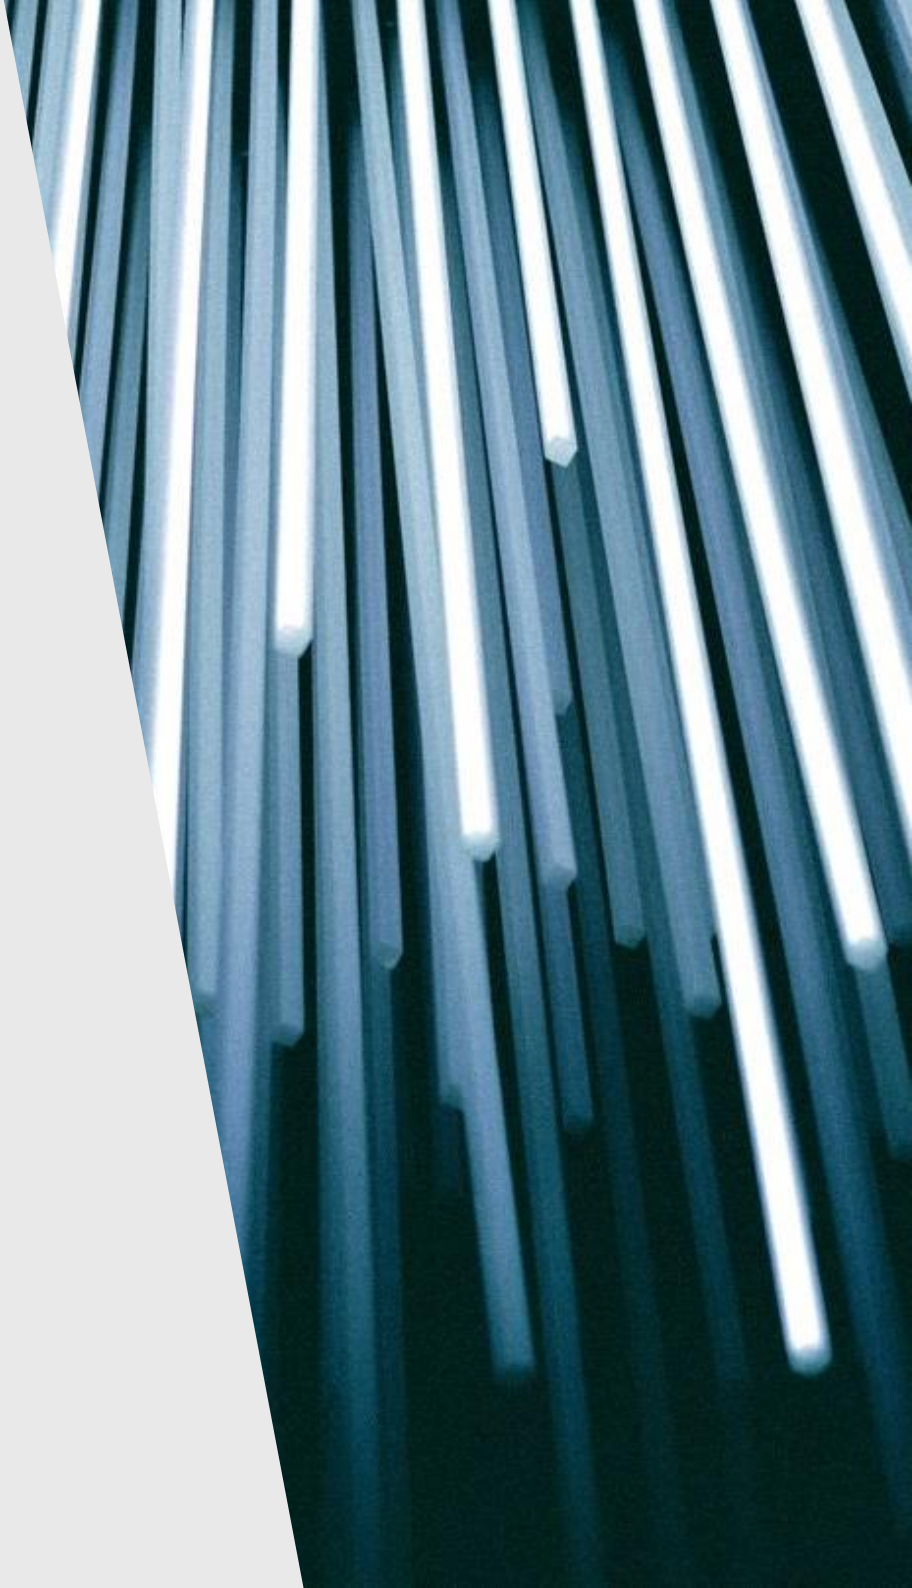

References

xxxiii Gupta, R., Reed, C., Rohrbach, A. and Darrell, T. 2022. 'Accelerating Ukraine Intelligence Analysis with Computer Vision on Synthetic Aperture Radar Imagery.' Available at: https://bair.berkeley.edu/blog/2022/03/21/ukraine-sar-maers/

xxxiv Kim, D., Won, J., Lee, E., Park, K.R., Kim, J., Park, S., Yang, H. & Cha, M. 2022. 'Disaster assessment using computer vision and satellite imagery: Applications in detecting water-related building damages.' Available at: https://www.frontiersin.org/articles/10.3389/fenvs.2022.969758/full

xxxv OpenAI. 2021. 'CLIP: Connecting Text and Images.' Available at: https://openai.com/blog/clip/

xxxvi Noufal, S. 2021. 'Photo Geolocation with Neural Networks.' Available at: https://github.com/kvsnoufal/ImageGeoLocation

xxxvii Van der Weide, Y. 2020. 'Using the Sun and the Shadows for Geolocation.' Available at: https://www.bellingcat.com/resources/2020/12/03/using-the-sun-and-the-shadows-for-geolocation/

xxxviii ACLED Homepage. N.d. Available at: https://acleddata.com

xxxix Department of Peace and Conflict Research. N.d. Available at: https://www.pcr.uu.se/research/views/

xl Serwat, L., Nsaibia, H., Carbone, V., and Law, T. 2022. 'Wagner Group Operations in Africa: Civilian Targeting Trends in the Central African Republic and Mali.' Available at: https://acleddata.com/2022/08/30/wagner-group-operations-in-africa-civilian-targeting-trends-in-the-central-african-republic-and-mali/

xli The Economist. 2022. 'Protest movements at deadly as Iran's often end in revolution or civil war.' Available at: https://www.economist.com/graphic-detail/2022/12/08/protest-movements-as-deadly-as-irans-often-end-in-revolution-or-civil-war

xlii ExTrac, 2022. Twitter post [25 August 2021]. Available at: https://twitter.com/ex_trac/status/1430526914482032649

xliii Gallacher, J.D., Heerdink, M.W. and Hewtsone, M. 2021. 'Online Engagement Between Opposing Political Protest Groups vis Social Media is Linked to Physical Violence of Offline Encounters.' Available at: https://journals.sagepub.com/doi/full/10.1177/2056305120984445. See also Williams, M.L., Burnap, P., Javed, A., Liu, H. and Ozalp, S. 2020. 'Hate in the Machine: Anti-Black and Anti-Muslim Social Media Posts are Predictors of Offline Racially and Religiously Aggravated Crime.' Available at: https://academic.oup.com/bjc/article/60/1/93/5537169?login=false

xliv See viii