beam©

POWERED BY

ISD | Institute for Strategic Dialogue

CASM technology

Beam: Defending Information

Carl Miller, CASM Technology
Melanie Smith, ISD
Dr Francesca Arcostanzo, ISD

Beam: defending information

Beam is a multi-lingual, multi-platform capability to expose track and confront information threats online, from disinformation to hate, extremism, information operations, harassment and harmful conspiracy theories. It is co-developed by CASM Technology and the Institute for Strategic Dialogue.

In 2021 Beam was the joint-winner of the US-Paris Tech Challenge for innovative approaches to counter disinformation, sponsored by the US State Department, the Department of Digital, Culture, Media and Sport, and NATO.

Acknowledgements

A large number of people have contributed to the creation and use of Beam, extending all the way from fundamental technology development through to analysis and reporting. This report explains their work, but all errors and omissions however remain the authors' own.

For CASM Technology: Jeremy Reffin, Andrew Robertson, Albertus Andito, Nestor Prieto-Chavana, Chris Inskip, Shaun Ring, David Weir, Justin Crow, Alex Krasodomski-Jones, Jack Pay and Simon Wibberley.

For the Institute for Strategic Dialogue: Sasha Havlicek, Cooper Gatewood, Jennie King, Jacob Davey, Henry Tuck, Milo Comerford and Chloe Colliver.

Information Threats

Our age is defined not just by how central information is to it, but also the many threats ranged against it. Information both surrounds us more and more, but we've probably also never been more worried by it, or suspicious of it. To both need information and be harmed by it; this is perhaps one of the most important paradoxes of our lives.

There has been a dawning awareness that information spaces can be manipulated, gamed, warped and controlled in ways that are often hidden but powerful. This has shaken our confidence not only in the health of information but also the many social and democratic practices that rely upon it. The integrity of elections are threatened by disinformation campaigns that suppress turnout or deny the result. Climate action is undermined by coordinated attempts to propagate pseudo-scientific reasons for delay or deflection. Families are split apart by anti-vaccine conspiracy theories amplified by autocratic states and those trying to oppose these information threats - from journalists and activists to Wikipedia editors - are harassed and threatened.

Frequently described as 'disinformation', information threats are much more varied than simply the propagation of falsehoods. They are technically broader, comprising a whole tradecraft of different techniques and tactics from coordinated amplification to malicious automation, manipulation of search engine results, spoofing identities and gaming social reputation systems. They are also sociologically deeper, where entire communities are now defined by their rejection of mainstream science and the experts, professionals and journalists that create and interrogate it. From autocratic state militaries to extremist political parties, there is an increasingly popular way of thinking about information that sees it as a theatre of war, where conflict is rated as surely as it is in air, sea, land or space. A for-profit industry has also emerged where companies offer online manipulation, and disinformation as a service for sale.

Information spaces need to be protected, from major events, elections and summits to local events and global conversations from climate action to sexual and reproductive rights. We also need to project the people working to protect the health of information ecosystems, from activists and investigative journalists to open-source investigators and academics.

To those who do it, the systematic manipulation of information is a route towards power and profit. We are only in the foothills of the information age, but also the foothills of the threats to information itself.

POWERED BY

ISD Institute for Strategic Dialogue

CASM technology

Contents

Responding to Information Threats

Since 2018, CASM and ISD have worked towards a common vision for how to defend information spaces, which this paper is about. It is a shared capability to detect, characterise and then confront information threats that is in parts technology, team, accumulation of experience and data and networks and relationships.

Beam blends two priorities. On the one hand, it is a capability that is technologically-driven; an attempt to harness the latest capabilities from machine learning and data science to the task of spotting and investigating information threats constantly and at great scale. On the other, it is an undertaking that is centred around humans, empathetic to the human experience and designed to sit within broad coalitions of civic society groups able to respond to information threats in a number of powerful ways.

It is shaped by 11 core principles:

1. **To be plugged into civic society** in a number of ways. Beam's direction is guided by civic society, it works in ways transparent to civic society and its outputs are designed to power civic societal responses to information threats.

2. **To be independent of any online service provider.** Beam avoids the development of capability that is only effective on a single platform or information space. Therefore…

3. **To work across as many different data sources and languages as possible**. Instead, Beam is designed to work across as many different information spaces as possible.

4. **To be sensitive to context but also to operate globally** including across the many themes, geographies, languages and bodies of information affected by information threats.

5. **To be connected to the under-served** languages, issues, communities and regions that are constantly targeted by illicit influence operations.

6. **Beam will never be in a settled state**. It must have a team of developers tasked with constantly and reactively adding new capability driven by investigators face-to-face

with the data. Important new requirements will best be discovered through the continued and practical use of the system. Hence…

7. **To be adaptable to the constantly emerging stream of new capabilities** from the world of machine learning and natural language processing.

8. **To also be adaptable to the constantly evolving tradecraft of illicit influence.** We know the threat actors we are confronting will also innovate and seize new opportunities created by emerging technology and societal change.

9. **To be as transparent as possible**, to enable peer review, share best practices and avoid black box technologies, magical algorithms or hidden techniques.

10. **To work across a range of themes and issues**, from climate summits to elections, the protection of activists and in the world of geo-politics.

11. **To power responses beyond platform removals or journalistic expose**, informing everyone from activists, to strategic counter-communications, diplomatic, political, legal and economic responses.

Beam

Beam is multi-layered and in what follows we will explain what each of these layers are and how they relate to each other. This will move from the most fundamental layer - the development of an underlying technology platform - to how Beam is used as a research and investigatory architecture, to the team that drive it, and the coalitions that use Beam's outputs in its various deployments around the world.

Layer 1: Underlying Technology development

Underlying Beam is a fundamental technology called Method52. Developed over the last ten years by CASM Technology, it was built to answer a general need for researching social media. Many subject matter experts saw social media data as relevant to their interests, but this data couldn't be handled by the research methods they were used to using. To cope with the full scale and complexity of data now available, they needed to use powerful machine learning and analytics technology, but in ways ways consistent with the social and behavioural sciences that they typically came from. Social listening dashboards for

marketing and advertising, and other kinds of 'plug and play' technology were often far too inflexible and inaccurate to fill this gap.

The building blocks of Method52 are around 100 different components. Each of these components does a different task; some allow data collection from Facebook, Reddit, Telegram, Twitter, YouTube, PeerTube, 4Chan, Discord, Mastodon, mainstream media news sources and a large number of standalone websites (where ethical and legal to do so). A large number of components make available the key analytic power of Method52: natural language processing and text analytics. This includes the training of artificial intelligence algorithms, the use of language-based network analytics, topic detection, named entity recognition and many more, explained in the Beam section below.

Method 52's users string and calibrate these components together into 'architectures'. Data flows from one component into the next, each doing something different - variously collecting, handling, storing, splitting apart, exporting data - in order to have the overall effect that the user wants. As ISD and CASM came together to respond to information threats, more and more architectures were built on Method52 that were dedicated to researching them. This collection of pipelines became easily the most complex and far-reaching that we'd ever used Method52 to build, and it was this collection of architectures that became the basis for Beam.

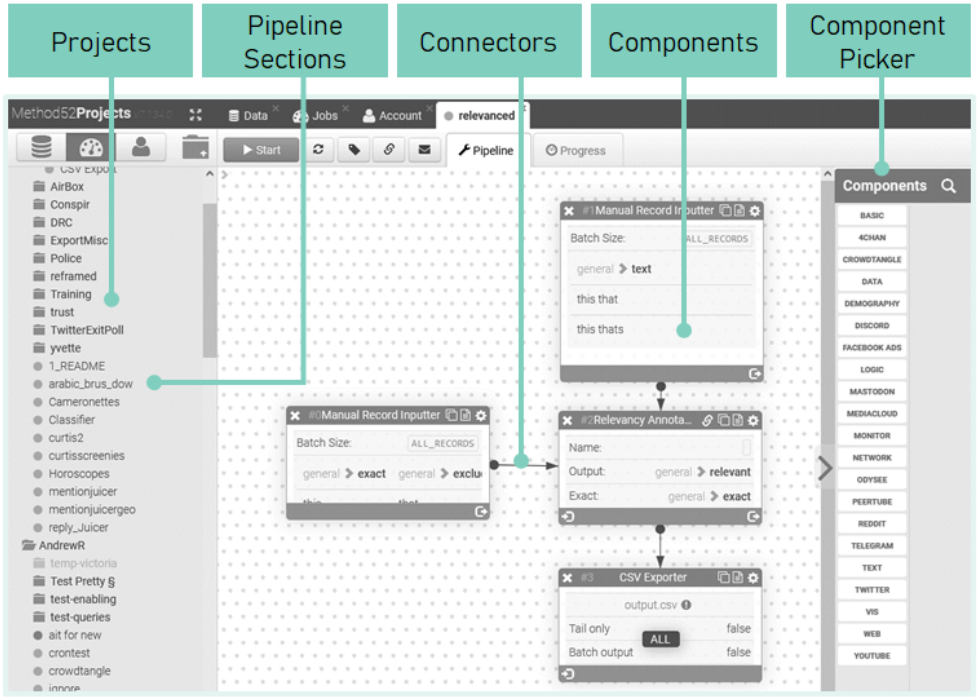

Layer 2: Beam's Capabilities

The second layer has been the continuous development of architectures on Method52 to detect, decompile, diagnose and measure the influence of information threats across languages and platforms. This has led to a series of faculties within Beam that represent distinct functionality but are also closely linked together. Practically, analysts hop back and forth between each of Beam's capabilities in an iterative way, quacking moving from discerning new narratives to collecting new data, and then looking at how certain accounts behave overall and from there identifying community-level groups which then creates new data collection opportunities. The best way of seeing these faculties are as subroutines linked by a series of common workflows that analysts move along.

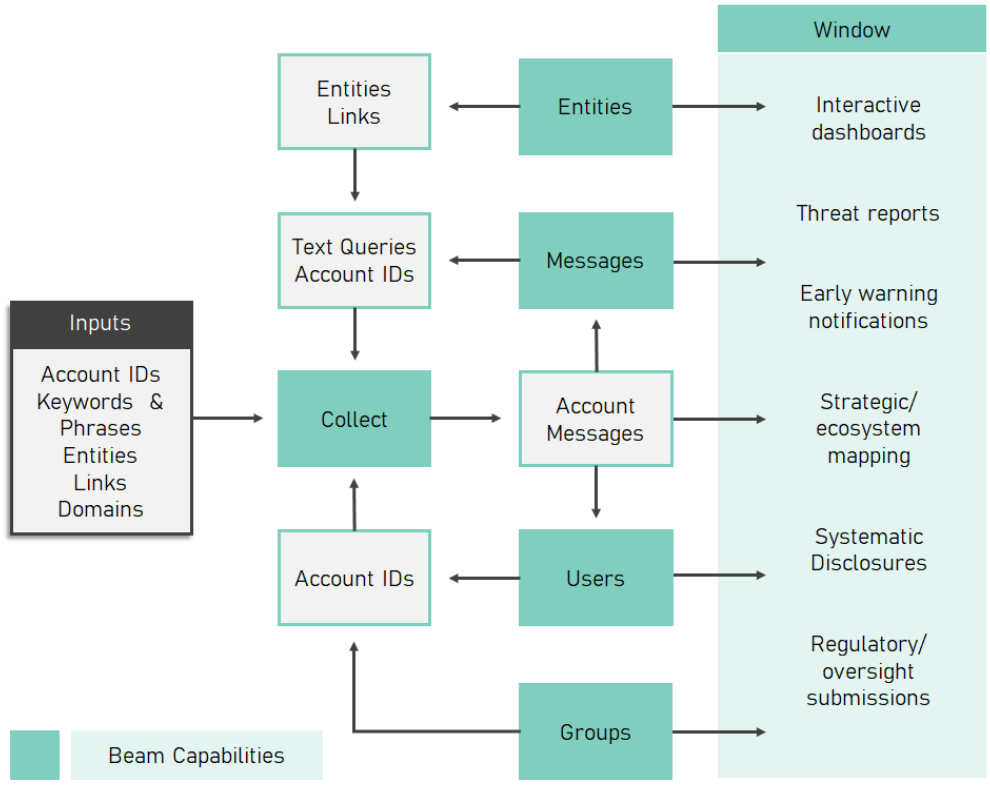

Beam Collect

- The first thing that any user of Beam needs to do is collect data; for this they turn to **Beam Collect.** It handles close-to-live data collections from social media platforms.
- It gathers data from Twitter, public Facebook Groups and Pages, Reddit, YouTube, 4Chan, Mastodon, Instagram, Telegram, PeerTube, MediaCloud, Discord and RSS feeds.
- Each platform offers different opportunities to collect data, but it is usually done either on the basis of keywords, entities, links or from specific accounts, channels or spaces.
- Data from all these different sources are brought to a single place and turned into a single canonical format to allow cross-platform analysis.
- It is integrated with Google Sheets to allow users to add new data collection inputs and narratives.
- 'Key phrase identification' is a specific capability that helps the user identify the kinds of keywords that will collect more data relevant to a given phenomenon. The is done by linguistically contrasting data already known to be relevant against a broader background of noise.

Beam Messages

As social media data is collected, researchers need ways of making sense of what is being said at the message-level. **Beam Messages** brings together technology to categorise and split apart social media messages based on the language that they contain, including specific claims, broader narratives, links and so on. This covers:

- **Link analysis.** Beam analyses the volumes and patterns of links and domains being shared, and helps to identify highly shared links and domains and produces flags whenever links are crossposted. It also visualises the evolution of links and domains sharing over time and across different communities, etc.
- **Named Entity Recognition.** This uses machine learning models to identify the places, locations and organisations being mentioned within each message. Analysts can use

this to find the key entities being mentioned overall, or to focus on messages only mentioning specific entities.

- **Geo-parsing** maps identified entities onto specific geographic locations, so analysts can then see where in the world messages are referring to or are talking about.
- **Narrative discovery.** Using a technology called topic modelling, Beam helps an analyst identify the key narratives contained within a body of collected data.
- **Narrative classification.** Once narratives are known, Beam allows analysts to train algorithms to categorise each message according to whether it belongs to a specific narrative.
- **Social velocity.** An underlying piece of machinery in Beam Messages automatically analyses the meta-data associated with each message, measuring the velocity of its engagement, sharing, viewership and so on.
- **Cross-posting.** Beam helps users find evidence of what we call 'cross-posting' - or the simultaneous propagation of identical messages or links across or within platforms.

Beam Hate Speech

Much of the classification that is done using Beam is bespoke; built specific to the context where it is needed. However, a more general capability has been created to detect hateful and offensive speech and understand whom this speech is targeting, across platforms.

To classify hate speech, we've built an 'ensemble' model. This is a process that uses a large number of other models that have already been designed to detect speech, including some from academic groups, some from commercial providers others from tech giants and some build by ourselves. This was supplemented by 27 lexicons and keyword lists.

- It has been trained to operate across Facebook, Instagram, Reddit, Telegram, Twitter, YouTube and 4chan.
- Trained to decide whether any given message is hateful or not, and the protected characteristic (such as nationality, disability, religious affiliation, sexual orientation or gender identity) that is targeted.

- Works at roughly 80% accuracy when compared to a human's judgement.

Beam Accounts

As key narratives are discovered, alongside important links, mentioned entities and so on, analysts will typically seek to move analysis to a more general level and look at how specific accounts across a number of messages that they've sent. Beam Accounts allows analysts to do this sort of account-level aggregation and sorting in the following ways:

- **Account-level summaries.** Which accounts are sharing the most narratives of interest, for instance, or the most varied collection of conspiracy theories? Do some accounts share certain kinds of links much more than others, or do some post particularly harmful messages?
- **'Spree' Posting:** Beam helps analysts see if an account is posting the same message in a number of different online venues. These 'sprees' have been one of the behaviours linked to influence operations identified in the past.
- **Account-level information.** The availability of this information depends on the platforms being researched, but can include biographical or channel information, account creation date, the total number of followers, members or subscribers and so on.
- **'Power-user' analysis.** We often see a small number of hyper-active 'power-users' on social media contribute the vast amount of the content that is actually produced. Power-user analysis helps analysts identify these accounts and separate them from others.

Beam Network and Communities

From understanding messages and accounts, analysts can also aggregate one level higher again and understand the online behaviour of groups of accounts. This is what Beam Network lets you do.

- **Threat-actor Networks**. Analysts often use Beam network to identify networks of threat actors that, for instance, all propagate certain disinformation narratives or engage in spree posting and collect all of their activity over time.

- **'Blue list' networks.** Analysts also use Beam to compile networks of accounts or spaces that are targeted by hateful or extremist actors, and to monitor the activity that is mentioning or directed at them.

- **Network expansion.** Beam network allows you to expand from a list of accounts that you know about to a larger list of accounts that are linked, in some way, to those accounts. It does this through a semi-automated process that, provided with some initial 'community definitions' (which can be either accounts, or keywords or links that they share) identifies new possible accounts that are linked.

- **Mapping Semantic Communities.** Beam uses deep learning algorithms to measure the semantic similarity of different accounts, and represent their differences as a map. This is a very powerful way to spotlight differences in accounts' behaviour based on what they actually say.

- **Follower, Mentions, Engagement Mapping** Beam creates a number of platform-specific and cross-platform networks based on the accounts', channels' or space's behaviour.

Beam Window

Beam Window produces Beam's outputs, including dashboards, data visualisations, regular reporting and alerts, that are used to power responses to information threats.

- Beam Window pushes data to **vis-analytic dashboards,** primarily build on Tableau Server. These allow top-level monitoring of large datasets, and a collection of sophisticated tables and charts to allow complex analyses to happen in a few clicks. These dashboards are usually heavily customised for each deployment of Beam.

- Continuous raw data streams, for especially data literate end-users.

- Daily and weekly snapshots.

- Flash reporting, triggered by important changes in the threat.

- Beam window also creates an early warning layer, where particular kinds of content are flagged when they have high 'velocity' in the sense of unusual levels of engagement and visibility.

- Specific thresholds and alerts to trigger focussed investigation.

- Specific threat reports - immediate threats - mentioning named institutions, individuals and so on.

Layer 3: Beam's Team

No single skillset is sufficient to confront information threats at scale. Data science alone is unlikely to uncover motivation, interests or identities. Manual analysts cannot cope with the sheer scale of social media data with which they are confronted. Beam is not just the technology, rather a whole interlocking set of different skills that come together around the investigations that we do.

- **Subject matter experts** have the deepest level of subject matter expertise which spans the actors who are seeking to conduct online manipulation, the targets of online manipulation, and the issue areas involved. They drive Beam, defining its priorities, initial landscape of threats and the most useful outputs.

- **Data journalists** are the most common users of Beam's dashboards, who spot new leads, whether anomalies, patterns, contrasts or consistencies with what is already known.

- **Open Source Intelligence practitioners** conduct targeted investigations of the most harmful, urgent and important detection that the system has made. Their role is to take leads from the system, and to make use of a separate suite of OSINT tools in order to uncover the possible identities, motivations, ownership structures, and hidden associations between online manipulation campaigns.

- **Beam Architects** are those who are most familiar with how the system and its various components works, and have a series of generic skills regarding the analysis and management of large online datasets. This is the team who is configuring and applying detections within the wider architecture.

- **Beam Developers** are the people who build and maintain the underlying system of Method52 that makes Beam possible. Their role is to react to novel analytical challenges raised by Beam's analyst and data interrogator teams, and apply backend software architectural development in response to these challenges.

Layer 4: Beam's Use

The third layer has been an emerging research practice for how Beam's technology and team come together.

Continuous Monitoring vs. Deeper Investigations

Practically, Beam is deployed in two different ways. The first is continuous monitoring, where online platforms, certain forms of language, lists of accounts, channels and spaces are monitored over time to identify information threats that are emerging, or changes in the salience, targeting or influence of ones that are already known to exist. The second is deeper investigation, where specific illicit influence campaigns are subjected to decompilation and close scrutiny to identify their targets, effects, methods and possible success.

Outputs into Inputs

A key principle of Beam analysis is that many analytical findings also represent new opportunities for collecting data. As researchers discover new disinformation narratives, for instance, they'll also discover new keywords to use for data collection. They might also discover accounts heavily engaged in the propagation of disinformation that they want to discover more about, or certain clusters of accounts that are closely linked to others that may be amplifying previously unseen narratives.

Multi-lingual

We have built Beam to function in as many different languages as possible. Beam has been used on French, German, Italian, Arabic, Dhivehi, Somali, Spanish, English, Chinese Mandarin, Vietnamese, Dutch, Farsi, Macedonian, Albanian, Russian, Bengali, and Portuguese content thus far.

Platform Agnostic

We deploy Beam across as many information environments as we can. This includes Twitter, Facebook, Instagram, YouTube and Reddit, but also emerging platforms such as Telegram, PeerTube, Discord, Diaspora and Odysee, standalone websites and Wikipedia.

Layer 5: Beam Deployments & Coalitions

Beam is deployed in many contexts around the world for different reasons: to protect an election, a summit, in the wake of a major event, in response to harassment against activists

and so on. Each of these deployments requires modifications to our approach; new technology, new skills within the team and new coalitions to confront information threats. Each deployment can have different languages, data sources, information threats, threat actors, targeted groups and coalition partners.

Coalition Responses

Beam is plugged into a whole array of end-users - journalists, lawyers, activists, fact checkers, regulators and governmental decision-makers - who can actually respond to information threats in effective ways. So far, it's been used to support over 350 civil society organizations across 10 countries to confront information threats. It has created over 90 non-public data briefings for partners, including legal, security and government partners, 28 public investigations, 150 media exposes of disinformation and fifteen reports of credible threats to the authorities. More information on the deployments of Beam to date are below.

Addressing under-served communities, languages and geographies

We work to deploy Beam outside of the contexts that are most served by research into influence operations and disinformation campaigns. This includes technical development of analytical and detection methods that are amenable to a broad set of languages and data sources.

Election Disinformation in the US

ISD and CASM have leveraged Beam over the course of three years to identify and inform the response to disinformation and weaponised hate online targeting two US election cycles: the 2020 presidential election and the 2022 midterms. This research has raised awareness of online harms among decision-makers in government, law enforcement, community groups, technology platforms, and the public. The project also aims to build an evidence base of gaps in platform policy and the enforcement of those across different types of harmful activity.

To this end, Beam has been deployed as a data-driven detection capability for disinformation, platform subversion and coordinated inauthentic behaviour (CIB) across nine social media platforms, combined with qualitative monitoring of threat actor channels. In order to deliver the insights in an actionable way and in real-time, this research is conducted in partnership with over sixty stakeholder organisations, who corroborate, verify and respond to information threats once detected. This stream of work often acts as the early warning system for information campaigns, online threats and new narratives, targets or tactics in the disinformation and extremism domains. At crisis points, Beam also becomes a vital

source of intelligence on online mobilisation for real-world action, including mass demonstrations and targeted violent attacks.

Over the course of this project, we have produced over 45 public research reports that expose the actors and platforms involved in promoting disinformation and hate online. These have covered topics such as election denial conspiracy theories, public health and vaccine misinformation, and analysis of foreign state actor activity targeting American social media users. This program has continued throughout 2022 and in the lead-up to the midterm elections using the same coalition-based model.

ISD and CASM have continuously iterated the methodologies and capabilities available for this work, which leans heavily on the flexibility of the Beam environment to detect and analyse incidents of online manipulation. This particular use case has created a collaborative practice by opening up complex data science techniques to mixed-skills teams of investigators, researchers, and subject matter experts.

Information Operations Regarding the War in Ukraine

Beam has been deployed consistently for the detection and analysis of pro-Kremlin disinformation campaigns since the invasion of Ukraine in February 2022. This project sought to examine the extent to which pro-Kremlin content and narratives were reaching online audiences, in spite of a number of governments, the EU and major tech platforms putting measures in place to curb propaganda. This research has spanned over 10 languages and countries and has leveraged next generation deep learning models, creating opportunities to optimise Beam's technological capabilities.

In one instance, the ISD and CASM teams used Beam to evaluate the authenticity of accounts promoting content from Western influencers who are known to share Kremlin disinformation about events on the ground in Ukraine across YouTube, Twitter, and Facebook. By retrieving account-level metadata and using Beam to isolate peaks in creation dates, researchers were able to determine that content from these influencers was being boosted by inauthentic accounts, as well as by state-affiliated media and government officials from both Russia and China. This secondary network was found to have propagated pro-Kremlin disinformation to over 41 million Twitter users.

In another case shortly after the invasion, CASM was able to map all of the accounts heavily sharing two popular pro-Russian hashtags: #IStandWithPutin and #IStandWithRussia. These hashtags had begun to trend on Twitter across a number of geographies around the world but the size of each, and the connections between them, remained unclear. This analysis allowed for a visualisation of the online communities engaging with it, clustered by their various languages (below).

This model placed each of the accounts heavily sharing the hashtag on a map, depending on their general use of language. This exposed clusters of accounts that had distinctively different national and linguistic identities, including pro-BJP accounts using Hindi, Farsi and Sindhi-using accounts, a Tamil-language cluster and accounts claiming to be from South Africa and Nigeria mainly using English. The investigation suggested that pro-Russian influence (in this instance) was targeting BRICS countries and more broadly the global south.

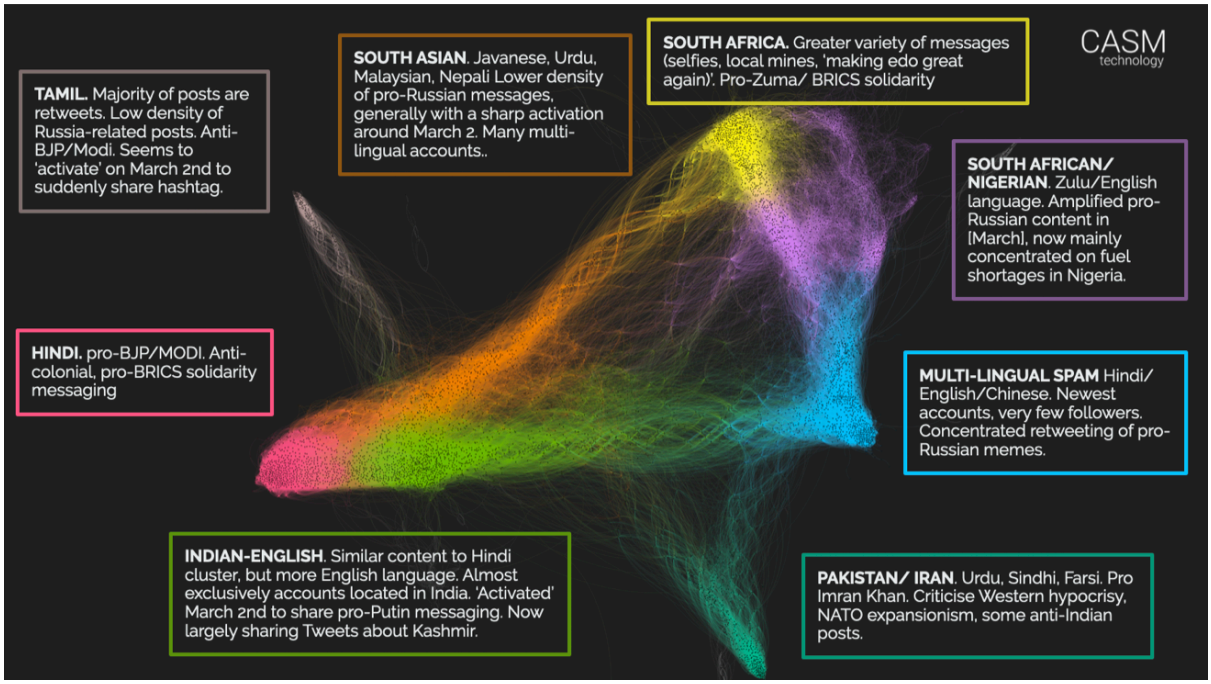

Election Disinformation in France

ISD and CASM leveraged Beam to investigate disinformation and misinformation in the months ahead of the April 2022 French presidential elections. During this time, the teams used Beam to collect data and perform various types of social media analysis and network mapping, alongside qualitative monitoring.

The French online election conversation at this time was rife with disinformation about the COVID-19 pandemic, the potential risks of the emerging Omicron variant, as well as

controversies around the vaccine pass and other sanitary measures. To better understand the online communities congregating around these topics, the teams utilised Beam to perform a network discovery exercise to hone in on a set of accounts sharing false and misleading information.

After collecting the output from these accounts over the course of a month, which amounted to over 1.7 million social media posts, the team used the link and narrative analysis modules within Beam to identify what types of content was being shared within these spaces with regard to the election. This research provides an overview of the key communities that were actively spreading disinformation and misinformation on social media platforms ahead of the elections, and the dynamics and narratives shaping these communities.

English-Language Disinformation about the Syrian Conflict

In 2022, The Syria Campaign commissioned ISD and CASM to conduct a digital investigation to help examine the extent of English-language disinformation about the Syrian conflict from 2015-2021. This analysis formed the basis of The Syria Campaign's report, Deadly Disinformation: How Online Conspiracies about Syria Cause Real-World Harm. Beam was deployed in this context to assess the narratives and targets of disinformation about Syria, and how both had evolved on social media over a seven-year timeline.

COP26/COP27 Climate Disinformation War Rooms

ISD and CASM used Beam to power a Climate Disinformation Dashboard over the course of the COP-26 Summit in Glasgow in 2021. This involved a systematised, live assessment of coordinated disinformation and malign influence activities targeting climate action. The team also developed strategies to expose, disrupt and mitigate malign coordinated influence campaigns around climate.

In 2021, the team worked with the Counter-Disinformation Unit at DCMS to integrate its monitoring with cross-Whitehall efforts. Twelve climate action NGOs – known collectively as the Climate Action Against Disinformation Alliance - were on-boarded onto the data system. This first-hand and real-time engagement allowed the teams to produce over 100 pieces of media and journalistic coverage (NPR, Guardian, BBC, New York Times) and reactive fact-checks from the COP26 Presidency.

An overview report produced in collaboration with these twelve partners gives a data-driven examination of the landscape, actors, systems and approaches that are combining to prevent action on climate. This report was presented at the United Nations Climate Change Conference and was covered by several major international media outlets and climate-

focused podcasts. An optimised version of this model will be deployed within the same coalition context for COP27, due to take place in Egypt.

Election Disinformation in Australia

CASM and ISD deployed Beam in the context of the Australian election for the purposes of training data journalists and analysts using it to detect election disinformation.

This involved compiling a list of political candidates, parties and other relevant political actors (e.g. Free movement groups), and live collecting data into structured formats that make it easily comparable across platforms.

Data were then presented to journalists into custom designed Beam-powered dashboard, produced in direct collaboration with data journalists and analysts to capture key themes and narratives of interest. These users were able to help guide the underlying analysis by constructing filters and narrative categories on the basis of keywords or natural language processing classifiers.

Throughout the election campaign, twice a week ISD and CASM ran dashboard trainings and clinics about how to interpret findings. Clinics were also used to identify promising leads emerging from the dashboard, and to design more in-depth follow-up investigations. Beyond the daily actor-based monitoring, Beam was then used by ISD and CASM's analysts to support these investigations, collecting and analysing new data multiple times a week.

As mentioned, ISD and CASM had previously used Beam to support and train the climate sector in how to understand and analyse disinformation relating to climate change, including in Australia. The use of dashboards to make Beam available to end users, like journalists and CSOs, enables monitoring of key actors and narratives threats across social media on a daily basis, including on Twitter, Facebook, Instagram, YouTube, Reddit and Telegram. is able to support key stakeholders in their own advocacy, public education, policy-making and strategic communications work.

Information Warfare on Wikipedia

ISD and CASM Technology set out to examine the ways in which Wikipedia may be vulnerable to the forms of systematic manipulation that have been exposed on Facebook, Twitter, YouTube, Reddit and a number of other information spaces. As a case study, we took the English-language Wikipedia page for the Russo-Ukrainian war, finding 86 editors who had previously made changes to the page and who had been banned from Wikipedia due to vandalism, sock-puppet accounts and so on. We used Beam to track the 794,771 revisions these 86 editors had made across Wikipedia, using network analysis to identify other pages which had seen many edits, identify the addition of state-affiliated sourcing, categorised the edits for (in some cases) pro-Kremlin bias and identified possible coordination through the introduction of common URLs.

New Zealand's Online Extremist Ecosystem

CASM and ISD used Beam for a project commissioned by the New Zealand Department of Interior Affairs that was intended to map the ecosystem of online extremism either from, or about, New Zealand. This included the online activities of extremists with a demonstrable link to the country and led to an exploration of transnational linkages and coordination between groups and individuals.

Exploring far-right, Islamist and far-left extremism as well as the growing grey area between conspiracy theories and extremism online, this research draws on data from social media sites including Facebook, YouTube and Twitter, as well as a range of 'alt tech' platforms, including Parler, Gab and Telegram, and data from stand-alone extremist websites and forums. The report formed part of the government's response to the Royal Commission of Inquiry into the 2019 Christchurch attacks.

BEAM TIMELINE

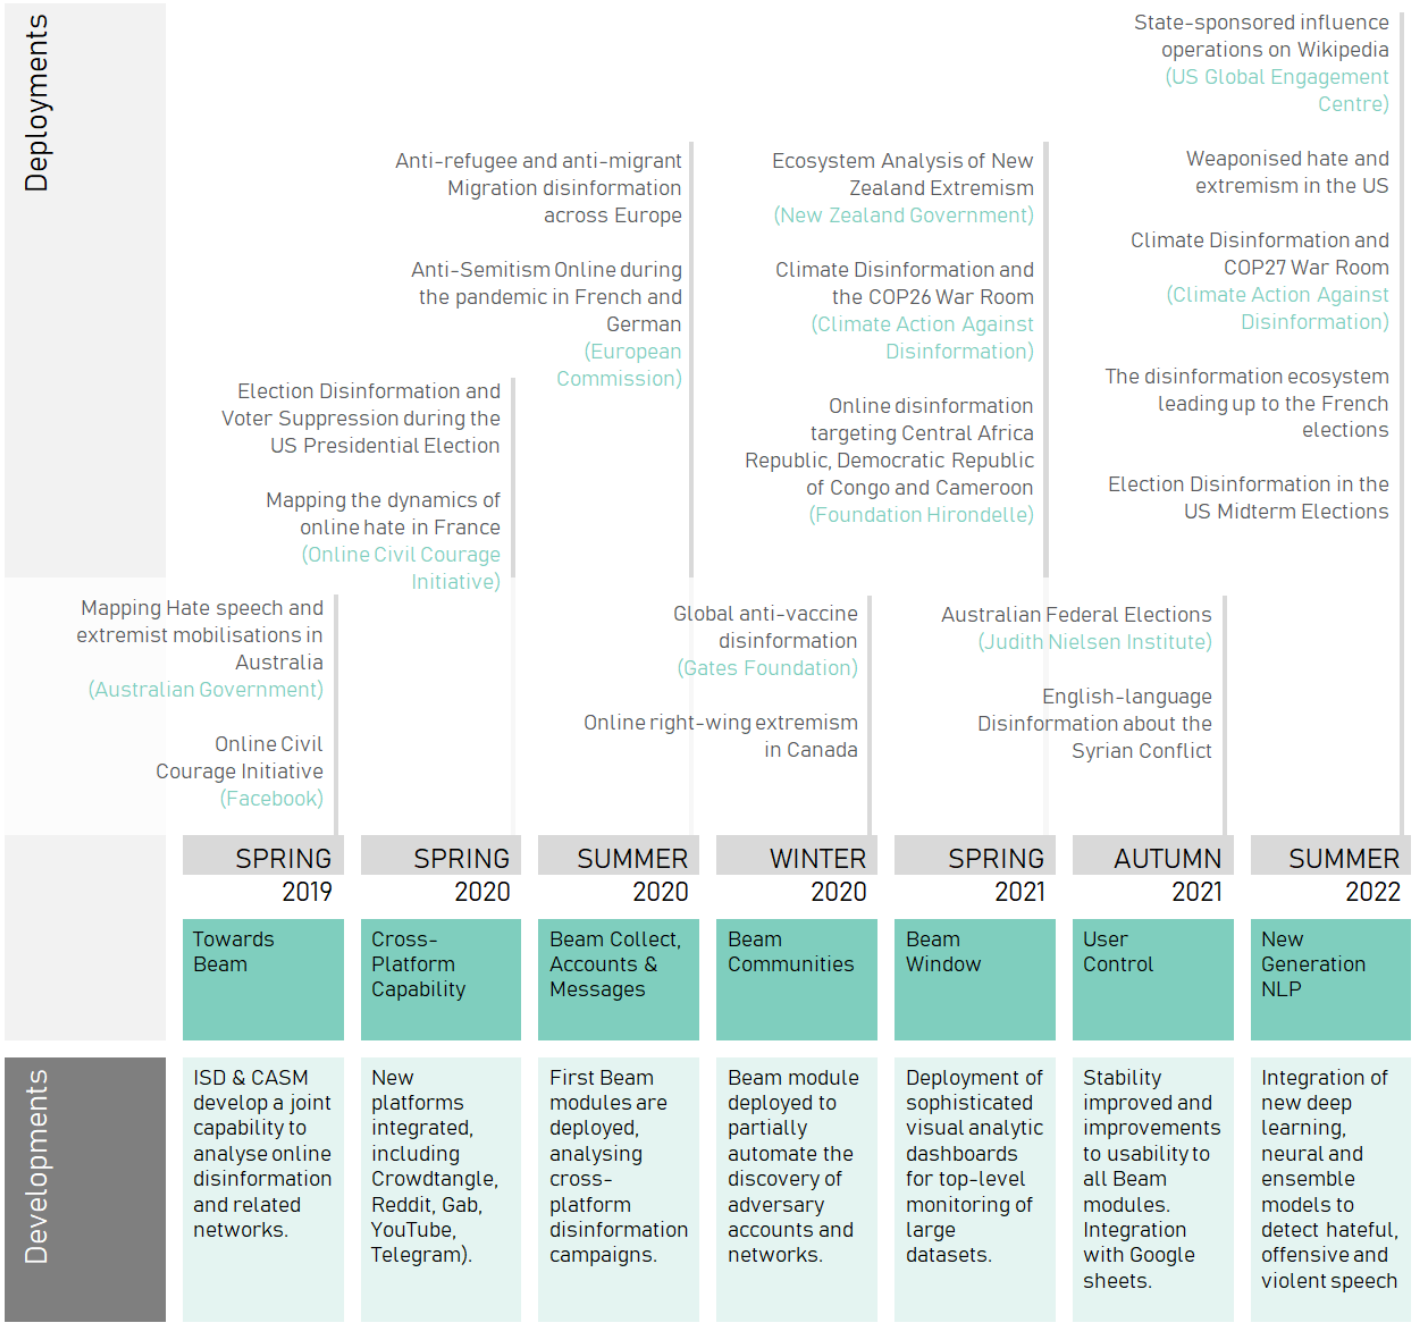

beam

POWERED BY ISD Institute for Strategic Dialogue | CASM technology