CANADIAN DIGITAL MEDIA RESEARCH NETWORK

Incident Report

INCIDENT REPORT

Stolen Faces: An Inauthentic TikTok and Instagram Network Inflaming the U.S.–Canada Trade War

Aengus Bridgman, Esli Chan, David Hobson, Sequoia Kim, Mathieu Lavigne, Manon Paquet, Chris Ross, Ben Steel

November 2026

CENTRE FOR MEDIA, TECHNOLOGY AND DEMOCRACY

MEDIATECHDEMOCRACY.COM

Key takeaways

- **A network of at least 293 TikTok and Instagram accounts is using deepfakes of real influencers, some Canadian, to push politically divisive content about the U.S.–Canada trade war to Canadians.** Since August 7, 2026, it has posted 3,201 videos that have been seen 22.1 million times. It started during tense trade talks and a run of votes: Quebec's election, Alberta's referendum, British Columbia's election, and the U.S. midterms. It was still posting about 100 videos a day when we stopped collecting on November 4 and all accounts were still active as of November 11, 2026.
- **The network steals real people's faces, then gives them AI voices and AI-redrawn lips to read scripts that are often AI-written.** TikTok's AI label almost never catches this manipulation and flagged only 4% of the videos built on real creators' footage as AI-generated. None of the network's videos disclosed AI use.
- **The network targets Canadians.** Most of its videos (84%) are about Canada, and the trade war alone (tariffs, trade talks, boycotts and the economic fallout) accounts for 71% of its views. Most of its TikTok followers are registered in Canada (16,055) or the United States (14,467). The most-watched videos we checked contain clear falsehoods. Its most-played story, a "48-hour ultimatum" from Trump, is fabricated and was viewed 5.1 million times.
- **The network is almost certainly run from outside North America.** Its posting and registration times do not fit North American time zones. Activity drops off during the night in East and Southeast Asia, and some accounts carry Chinese-language settings or a registration in China, the only lead we have on a country. We do not know what the network is for, but its operators are willing to steal faces and spread falsehoods that set Canadians against the United States, against their own government and against each other.

Points saillants

- **Un réseau d'au moins 293 comptes TikTok et Instagram utilise des hypertrucages de véritables influenceurs, dont certains sont Canadiens, pour diffuser auprès des Canadiens du contenu politiquement clivant sur la guerre commerciale entre les États-Unis et le Canada.** Depuis le 7 août 2026, ce réseau a publié 3 201 vidéos, qui ont été visionnées 22,1 millions de fois. Il est apparu dans un contexte de négociations commerciales tendues et d'une succession de scrutins : l'élection au Québec, le référendum en Alberta, l'élection en Colombie-Britannique et les élections de mi-mandat aux États-Unis. Le réseau publiait encore environ 100 vidéos par jour lorsque nous avons cessé la collecte de données, le 4 novembre, et tous les comptes étaient toujours actifs en date du 11 novembre 2026.
- **Le réseau s'approprie les visages de personnes réelles, puis leur attribue des voix générées par l'IA et des mouvements de lèvres recréés par l'IA pour leur faire lire des textes qui sont souvent eux-mêmes rédigés par l'IA.** L'étiquette de contenu généré par l'IA de TikTok ne détecte presque jamais ces manipulations : seulement 4 % des vidéos utilisant des images de véritables créateurs ont été signalées comme générées par l'IA. Aucune des vidéos du réseau ne divulguait le recours à l'IA.
- **Le réseau cible les Canadiens.** La majorité de ses vidéos (84 %) portent sur le Canada; la guerre commerciale à elle seule (p. ex., droits de douane, négociations commerciales, boycotts et répercussions économiques) représente 71 % des visionnements. La plupart des abonnés TikTok du réseau sont enregistrés au Canada (16 055) ou aux États-Unis (14 467). Les vidéos les plus visionnées que nous avons examinées contiennent des affirmations manifestement fausses. Le récit le plus populaire du réseau, portant sur un prétendu « ultimatum de 48 heures » lancé par Trump, est entièrement fabriqué et a été visionné 5,1 millions de fois.
- **Le réseau est presque certainement exploité depuis l'extérieur de l'Amérique du Nord.** Les heures de publication et d'enregistrement ne correspondent pas aux fuseaux horaires nord-américains. Les activités diminuent pendant la nuit en Asie de l'Est et du Sud-Est, et certains comptes utilisent des paramètres en langue chinoise ou indiquent un enregistrement en Chine, le seul indice dont nous disposons quant au pays d'origine. Nous ignorons les objectifs de ce réseau, mais ses opérateurs sont prêts à voler des visages et à propager des faussetés qui dressent les Canadiens contre les États-Unis, contre leur propre gouvernement et les uns contre les autres.

Context

On June 17, 2026, Ashley Hetherington, an American content creator, posted a four-minute TikTok about her thoughts on dating as a Christian woman. Months later, on October 25, 2026, a near identical video of Hetherington appeared again from an Instagram account called Raven, which Instagram labels as based in China. In the new video, Hetherington is seen talking about Canada's job losses, U.S. tariffs, and Alberta's referendum. Raven had taken the first 21 seconds of the original video and used AI tools to digitally alter her lips to match the new AI voiceover.

But Hetherington was not alone. By tracing the same format, titles, captions, and typos across Instagram and TikTok, we found a network of at least 293 accounts (36 on Instagram and 257 on TikTok) posting the same kind of videos. Between August 7, 2026 and November 4, 2026, the 293 accounts posted 3,201 videos with 22.1 million plays (82% on TikTok, and 18% on Instagram). During this time, posting activity of the network grew quickly, from three active accounts posting two videos a day to 157 accounts posting 108 videos a day. This period coincided with the Alberta referendum, the Quebec provincial election, the U.S. midterms, and an ongoing trade war between Canada and the United States.

This network joins a growing number of documented foreign inauthentic networks targeting Canadians: in April 2026, we found 20 YouTube channels selling secession narratives to Albertans, and in June 2026, we found a network of 340 Facebook pages monetizing AI-generated fake news. This new network stands out for its lack of any visible way to make money, its wide reach and considerable growth, its theft of creators' likenesses, and the political focus of its activities. In this incident report, we investigate the network's inauthentic activity, describe the content it is producing, assess its likely origins, and conclude with what this means for the Canadian information environment.

All counts in this report likely underestimate the network's full reach, because it was still posting and expanding when we closed data collection.

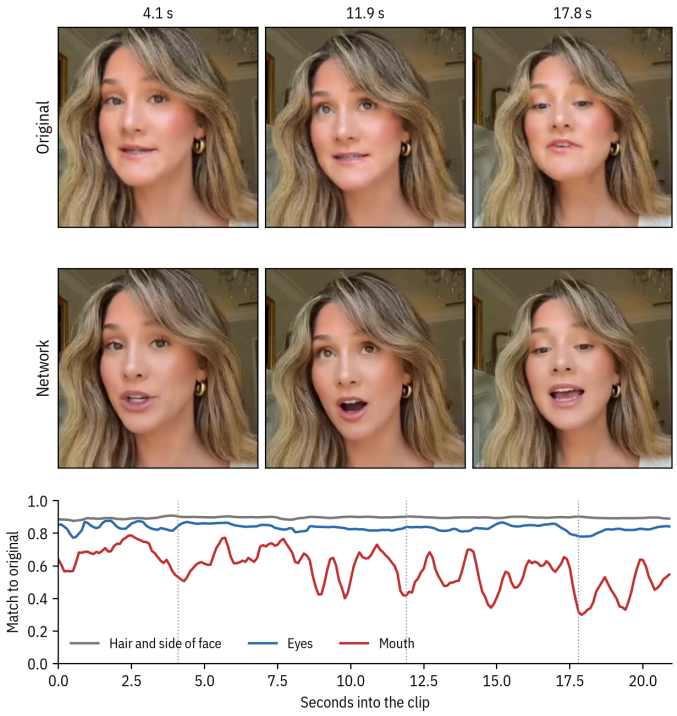


Figure 1: Stills from Ashley Hetherington's TikTok posted on June 17, 2026 (top) and from Raven's Instagram reel posted on October 25, 2026 (bottom) at the same moments. We aligned the two videos frame by frame and measured how closely three parts of the face match, pixel by pixel (1 = identical). The median match was 0.90 for the hair and side of the face, 0.84 for the eyes and 0.61 for the mouth. The head and eyes are copied directly, but the mouth and lower face have been redrawn.

How do we know the accounts are fake?

Raven's reel is typical of how the network makes many of its videos. Figure 1 shows stills from Hetherington's original video along the top row and stills from Raven's inauthentic reel on the bottom row. Her head, hair, and eyes match the original frame for frame, but her mouth and lower face do not: they have been redrawn.

The same technique runs through the network. Between October 23 and 30, 2026, the inauthentic Instagram account @anasxpongofurrov43 posted six reels, all stolen from the same clip of creator Andrei Jikh, with six distinct scripts. Across all six, the head, eyes and shoulders move the same way, but the mouth moves differently (Figure 2). Across these videos, the message is inconsistent: The same Jikh is made to say that Canada's economy is collapsing in one reel, and that Canada is building a future beyond the United States in another.

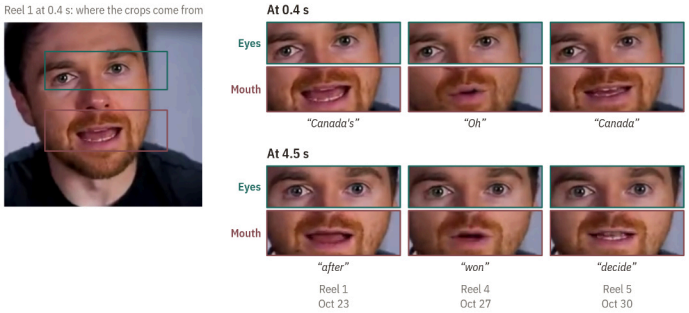
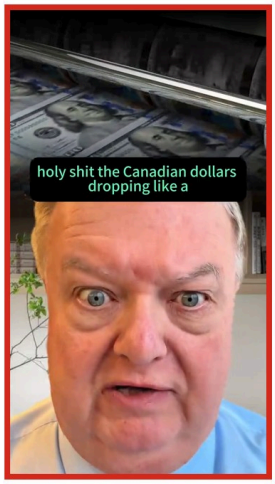


Figure 2: Six different scripts on the same video clip. The eyes and mouth from three of the six reels posted by @anasxpongofurrov43 (handle as of October 31, 2026) between October 23 and 30, 2026, each built from the same Andrei Jikh clip, shown at the same two moments. The eyes match but the mouth and the word being spoken (from automatic transcription) differ.

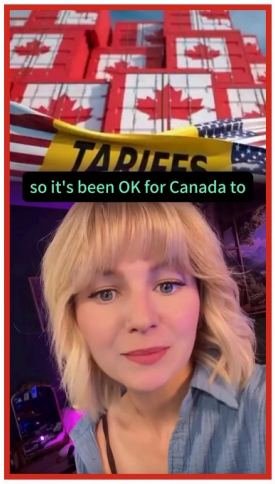
Almost every account (283 of 293) is built on footage of a single real person, stolen, where we could trace it, from that person's own videos. We counted 226 distinct people and have been able to identify 104 of them. We also found three TikTok accounts that use fully AI-generated avatars.

The faces in the network overwhelmingly appear to be white women. Using FairFace, an open-source face-attribute classifier that estimates perceived gender and race (Kärkkäinen and Joo 2021), we found that of the 226 faces, 187 appear white. These 187 white faces account for 90% of the network's plays. Another 23 appear to be East Asian, and the rest (16) appear Middle Eastern or Southeast Asian. None appear to be Black, South Asian, or Latino. Women outnumber men by five to one (188 to 38).

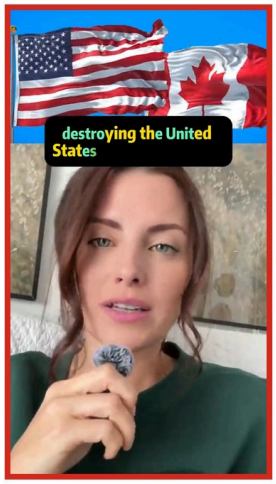
Based in Canada



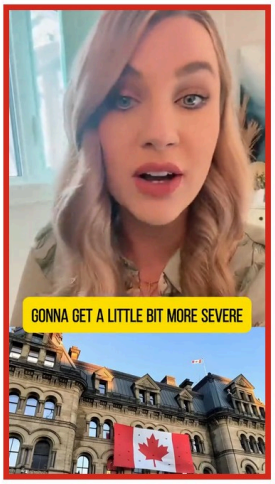
Footage of Ron Butler
posted by @fyjmbtwcbbb
(TikTok)



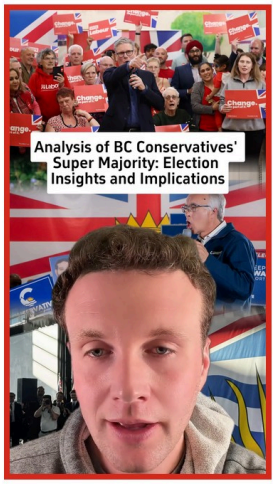
Footage of Kris Collins (Kallmekris)
posted by @lena765848
(Instagram)



Footage of Kels (UnfilteredWith Kels)
posted by @zsdtfduwpos
(TikTok)



Footage of Jasmin Laine
posted by @am_la248
(TikTok)

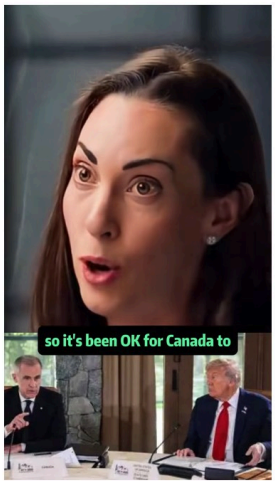


Footage of Cole Bennett
posted by @utehobw3376
(TikTok)

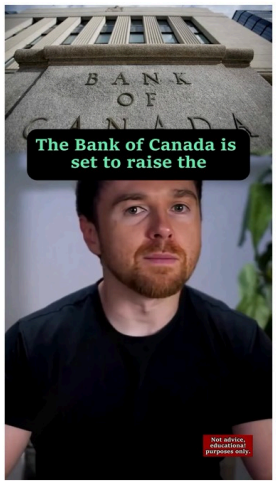
Based elsewhere



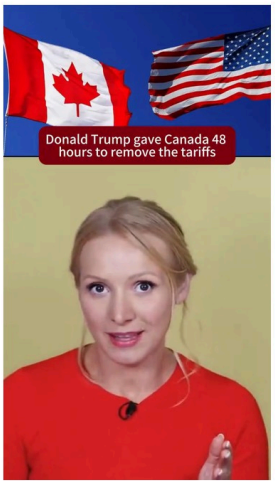
Footage of Robert Greene
posted by @aa_dam37
(TikTok)



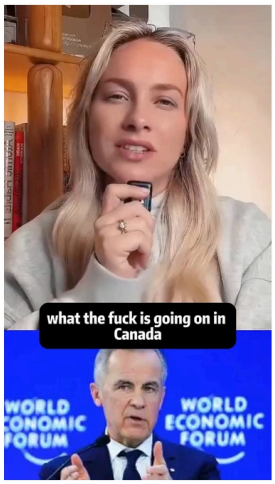
Footage of Vanessa Van Edwards
posted by @shvh3157xno
(TikTok)



Footage of Andrei Jikh
posted by @anasxpongofurrov43
(Instagram)



Footage of Lucy Earl (English with Lucy)
posted by @hsjxjymd4xj
(TikTok)



Footage of Erin McGoff
posted by @dhhiiwcsu0w
(TikTok)

Figure 3: Ten real creators whose faces appear in the network's videos.

How is AI being used?

The network uses AI extensively for its deception. First, it pairs the stolen faces with synthetic voices generated by an AI text-to-speech tool. A deepfake-speech detector rates the narration as AI-generated in 94% of the network's videos. Second, it uses an AI tool to generate realistic-looking lip movements that match the synthetic speech: the hair and eyes repeat exactly, but the mouth does not.

Third, many scripts appear to be written by AI. An AI-text detector (Pangram) rated an estimated 46% of the network's scripts as entirely machine-written, including the network's most-played story about an invented "48-hour ultimatum" from Trump. The scripts are also heavily recycled. Of 1,123 distinct scripts, 326 were posted more than once, and 142 were rewrites of an earlier script that say the same thing in different words. Scripts that do well get posted again. A reused script's first video drew a median of 1,626 plays, against 664 for a script posted only once.

Building videos from real people's footage likely helps the network avoid AI detection. TikTok's AI label caught most videos with fully AI-generated faces (22 of 24) but almost none of those built on real creators' footage (4%). None of the network's videos disclosed AI use.

What is the network posting about?

Almost everything the network posts concerns Canada, its economy, its government, and its relations with the United States. The biggest single topic is the U.S.–Canada trade war: tariffs, trade talks, and boycotts and their fallout make up 54% of its videos and 71% of views of network content. The most-watched scripted videos discuss Canada's "one-sided" tariffs on U.S. goods, claims that the U.S. would save billions by cutting off trade relations with Canada, and an invented claim of a "48-hour ultimatum" from Trump. Trade war narratives dominate largely because the network reposts trade-related videos more often than anything other topic.

The network does not take a consistent side in the trade war. Of the videos focused on the trade war, attacks on Carney's government are most common by number of videos (56%). However, almost every account takes both sides, and narrators switch nationality too, presenting themselves as Canadian in some videos and American in others. The effect is that Canadians are shown fabricated and misleading claims on both sides of the trade debate.

Most of the remaining videos cover Canadian domestic news, focusing on the economy, jobs, interest rates, housing, the dollar, and gold (23% of videos and 17% of plays). When this domestic content takes a side, it is almost always against Carney's government (97% of the side-taking videos). 130 of 187 videos that spoke about the Alberta referendum promoted separation. The Epstein files and personal finance are also discussed.

The network consistently shows a disregard for the truth and many of its most watched videos contain fabrications. The most-played video claimed that Carney answered a "48-hour ultimatum" from Trump by cancelling $8 billion in U.S. defence contracts. This never happened, but the story was shared in 161 videos on 129 accounts and played 5.1 million times. Many other videos start from something real but have reckless distortions. For example, Canada's talks with the European Union become a signed deal letting the EU decide whom Canada trades with. The youth unemployment rate becomes "Canada's unemployment rate hit 13%." The distortions cut both ways, against Trump and against Carney's government.

Who is watching?

A significant portion of the network's audience are unknowing Canadian TikTok and Instagram users. We collected the follower and following lists of all 257 TikTok accounts, and found that the accounts have 42,538 followers, with 16,055 registered in Canada and 14,467 in the United States. Canadians find its videos through TikTok's recommendations, and some then follow. The network follows many of them back.



Figure 4: Highly liked TikTok comments on network videos. Every top comment (per video) with 50 or more likes, grouped by whether the commenter agreed with the narrator, argued with them, or said the video or speaker is AI, fake or a bot. The like count is as of November 5, 2026. Usernames are removed and spelling is as posted.

Very few of the people who comment on the videos see the content for what they are: an inauthentic operation targeting Canadians and inflaming the U.S.–Canada trade war. Only 6.3% say or suggest that a video is fake, AI-generated, a bot, copied or propaganda. The rest agree or argue with the content in the video, and one in ten speaks to the narrator as a real person: thanking them, praising them, or telling them to check their facts. On TikTok, comments that talk to the narrator drew 3.5 times as many likes as comments that doubt the authenticity of the video (14,713 against 4,241), and none of the 100 most-liked comments doubt it (Figure 4).

Who runs it, and why?

Ultimately, we do not know why the network exists or who pays for it. However, evidence we have been able to collect indicates that the accounts are run from East or Southeast Asia, most likely in the UTC+8 or UTC+7 time zone. The only lead we have on a country points to China: some accounts carry Chinese-language settings or are registered there. That lead is limited, and it may reflect whoever supplied the accounts rather than the people who run them. We provide two kinds of evidence.

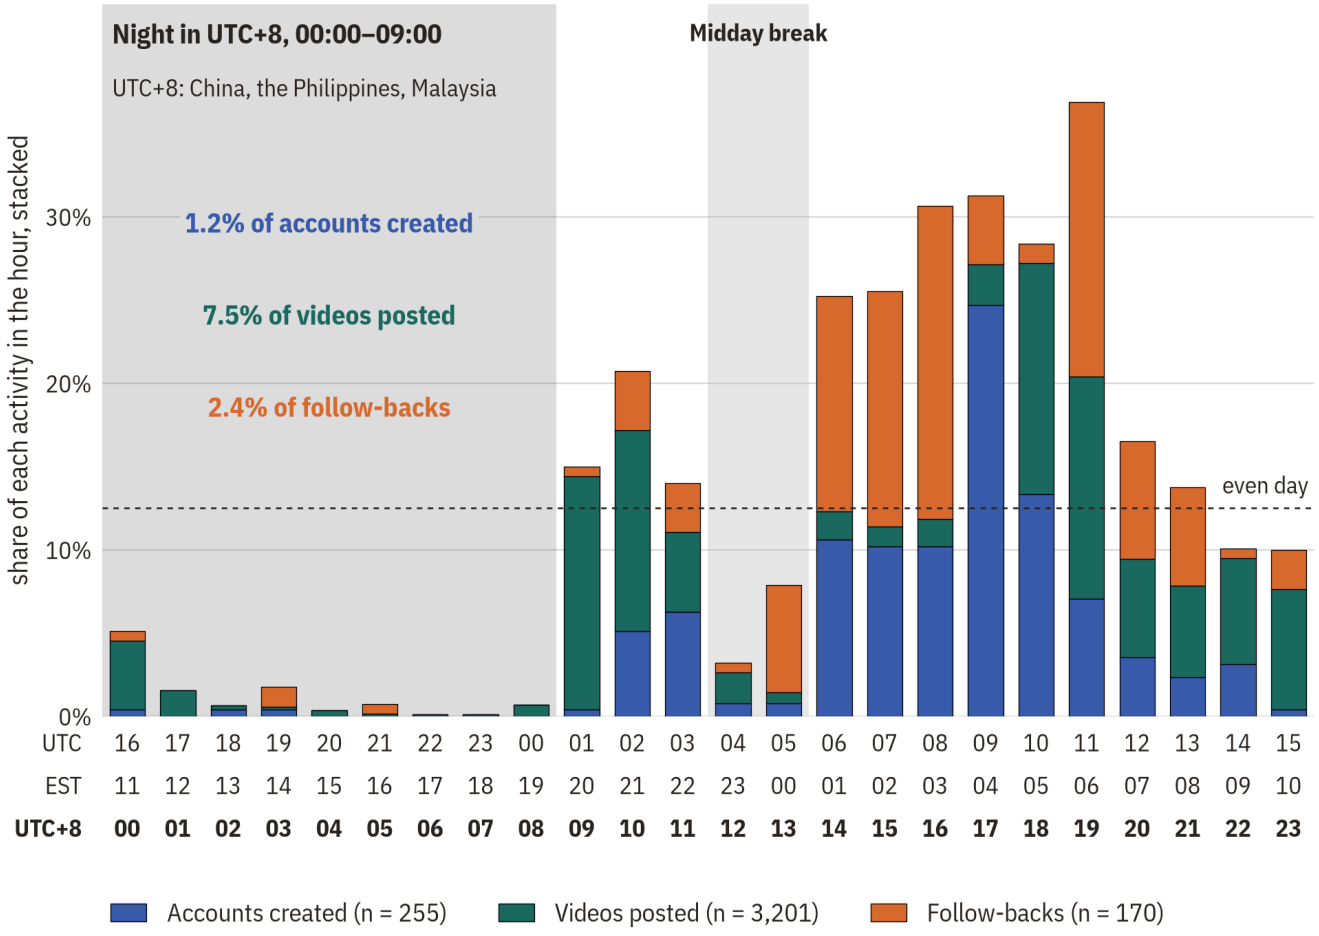


Figure 5: Share of the network's account creation, posting and follow-backs in each hour of the day. Each bar stacks three activities in one hour, each as a share of its own total: accounts created (the 255 TikTok accounts registered in 2026, timed from their user IDs), follow-backs (170 follow times from the following lists of 18 TikTok accounts; 97% of the network's follows are of its own followers) and videos posted. Hours are shown in UTC, EST, and UTC+8. The grey band marks night in UTC+8 (00:00–09:00), the lighter band marks a midday break (12:00–14:00); the dashed line is where the bars would sit if each activity were spread evenly over the day. Every video with an exact upload time is counted (n = 3,201).

First, the network almost never posts between 1 p.m. and 7 p.m. EST (UTC–5), which is 2 a.m. to 8 a.m. in UTC+8, the time zone of China, the Philippines, and Malaysia (see Figure 5). Only 1% of its videos are posted during the Canadian afternoon and early evening. On TikTok, we could read the registration time of all 257 accounts. Of the 255 created in 2026, 80% were registered between midnight and 8 a.m. EST, with only 12 registered during EST office hours. The timing rules out North America and fits a working day in UTC+8 (China, Hong Kong, Taiwan, the Philippines, Malaysia, and Singapore) or, one hour behind, UTC+7 (Vietnam, Thailand, and western Indonesia). We checked app language settings and common typographical indicators to further narrow a likely geographic origin. While operators typically hide their country of origin and language in similar operations, we did find that twelve of the 184 we could check carry a Chinese-language app setting or are registered to China, against fewer than one in a hundred comparable accounts. This is the clearest country-specific lead we found. Many of the titles the operators type themselves show signs of non-native English, at several times the rate of other accounts posting on the same subjects. A few carry punctuation or phrasing typical of Chinese input software. Taken together, these traces suggest that Chinese speakers are involved somewhere in the operation, either among the people who run the accounts or the supplier who created them.

Second, the timing of the network's posts does not fit a Canadian operator. For instance, Albertans voted in a referendum on October 19, but for 14 days afterwards, the network continued posting videos that describe the vote as still to come (43 videos on 37 accounts). Over the same days, other accounts in the network reported the result. Most of the 43 reused two scripts that had already run across the network before the vote, first posted on October 11 and 13. This suggests that the network draws on a shared pool of scripts and hands them out to its accounts without checking whether they are still current. Someone in Canada would be unlikely to miss the outcome of a major national story for two weeks.

Unlike the networks in our previous reports, which earned money from their content, these accounts show no sign of earnings. After nearly three months and 22.1 million plays, none of the TikTok accounts have reached 10,000 followers, the minimum threshold for TikTok's creator payments. We found no links, products, sponsorships, or subscriptions in their posts. Most accounts are also based outside the countries where TikTok pays its creators: TikTok lists 161 of the 184 accounts we could check as Canadian (Canada is currently shut out of per-view payments), and Instagram's direct payments for reels are an invite-only test, offered in only some countries.

Nor do these accounts push any foreign government's line: few videos explicitly favour China, Russia, or India, for example. The network does attack Trump's America in some videos and Carney's government in others. That is consistent with an aim to divide Canadians, but also with content built to profit from a polarized political context.

With limited insight into the account details, conclusive claims about the origins and intent of these networks cannot be fully made. The way the network is run looks more like a commercial content operation than a propaganda effort but the lack of monetization is puzzling. There are several other possibilities too: the network could be an account farm, growing audiences to sell or use later; it may be testing a tool or service; or a commercial or political client may be paying it to influence Canadians. Regardless of intent, what we do know is that someone started a network to unsettle Canadians during a tense trade war, setting them against the United States, against their own government and against each other, and is willing to steal people's faces and spread falsehoods to do it. Ultimately, Meta and ByteDance hold the details that would help us better understand this network, including the devices, logins, phone numbers, and any payment accounts behind these accounts.

What are the policy implications?

The January 2025 report of the Public Inquiry into Foreign Interference in Federal Electoral Processes and Democratic Institutions warned that information manipulation poses an existential threat to our democracy. The Commissioner acknowledged the "growing potential of artificial intelligence ("AI") to create deepfakes and other materials [which] makes it more difficult for Canadians to assess whether what they see or hear is real, or what they read is true."

In less than two years, that capacity has grown rapidly: back in April 2026 we found YouTube channels selling secession narratives to Albertans, and in June 2026, we documented networks of fabricated news websites created to earn advertising revenue. These earlier cases showed that platform monetization offers an attractive business model for actors unconcerned about the veracity of their content and whether it pollutes the information ecosystem. Conversely, in this incident, we could not identify an easy path to monetization, and it is therefore possible to imagine a more strategic objective motivating the network.

While the Government of Canada has taken some actions to implement the recommendations of the Public Inquiry related to information manipulation, such as targeted amendments to the *Canada Elections Act* and the expansion of the Rapid Response Mechanism, there remain two key areas of vulnerabilities in how Canadians are protected from AI-enabled influence operations that can sway public opinion, exacerbate tensions and, ultimately, affect the outcome of an election.

The first relates to the capacity to detect these operations, both within and outside government. The Government's October 2026 national security strategy promises to "enhance Canada's Foreign Information Manipulation Detection capability, using open-source information." That is a welcome commitment, but detection cannot rest on the government alone. Independent journalists and researchers are often the first to spot these networks and inform Canadians, but they can only continue to effectively do this work and expand their investigations if they have access to platform data and protection for their work.

The second is the protection of individuals against impersonation. The *Canada Elections Act* addresses impersonation in the electoral context, but the incident described in this report shows that impersonation and fabricated content can be used to influence Canadians well outside the Act's reach. Influencers and content creators are an increasingly important part of the political discourse: by acting as mediators between institutional actors and Canadian citizens, influencers shape public opinion and how politics is communicated. In this incident, the ability to put words in the mouths of real people—figuratively and literally—without their knowledge or consent gives malign actors a way to unduly influence a trusting audience. This incident makes clear that the risk is here in Canada: Canadian creators' likenesses are being used to tell Canadians things about their own economy and politics that these creators never said. While tort law may have some utility in providing a recourse at the individual level, it does not address the risk of broader societal harm.

Both TikTok and Meta require photorealistic AI videos to be labeled and have policies against networks engaging in coordinated inauthentic activity. In both these cases, irrespective of the actual policies, there are questions about how thoroughly they are properly enforced, or whether this network represents a genuine gap in their detection methods, particularly with respect to identifying AI-altered videos based on real footage. For example, TikTok's AI label caught 92% of the network's fully AI-generated avatar videos but only 4% of the videos built on real creators' footage, meaning the label misses exactly the technique on which this network relies.

To the broader question of impersonation and the use of someone else's likeness in AI content, even when properly labeled, the two platforms' policies differ. TikTok bars AI-generated content made from a private individual's likeness without their consent as well as misrepresenting comments from public figures, even if it is labeled as AI-generated. Meta's policy is narrower and bars using someone's likeness to impersonate them and deceive others. While Meta has a rule against the "unwanted manipulated imagery" of private individuals, they have previously stated that the "fabrication or misrepresentation of a person's actions or conduct is not protected" under this rule. In a recent case, Meta found that an unlabelled AI video of a U.K. local politician making comments about refugees that she never made broke none of its rules, and left it up without an AI label despite the politician reporting it herself. Meta's Oversight Board eventually overturned this decision because it found the video contained hateful speech.

Cases like these speak to the broader issue of platform governance and accountability. The Safe Social Media Act (Bill C-34) represents significant progress in this respect in Canada, including by requiring the labelling of synthetic content. However, it is important to remind ourselves that the bill focuses on specific harms and will have limited impact in holding platforms accountable when it comes to preventing impersonation and coordinated inauthentic activity.

Data and Method

Finding the accounts

Instagram's "Similar accounts" list on Raven's profile led to a second account with the same layout and the same typo in its title. From there, we searched both platforms, word for word, for the network's titles, captions and script text, including its typos, such as "tradethe", "USTariffs" and "Candian". On Instagram, we ran 86 searches, stopping when two rounds in a row added no account. On TikTok, a researcher ran 16 searches while logged in, without posting or interacting, and an automated search for the network's script text surfaced several hundred more accounts.

We tested 808 candidate accounts for three features of the network: a template cover with a network caption, a script the network speaks, or a narrator voice it shares. Of these, 266 had at least one. We then reviewed 359 accounts by hand (all 266 of these and 93 more flagged on other signs) and confirmed 293 of them (257 on TikTok, 36 on Instagram) as part of the network. This report counts only these 293.

We fixed the sample on November 4. We collected every confirmed account's full timeline in one sweep up to 11:00 UTC on November 4. All plays in this report are as displayed at that time, except for 56 videos missing from the profiles: 31 TikTok videos left out of their profile but still public at their own link, 15 TikTok videos deleted or removed, and 10 Instagram reels no longer listed, which we collected between October 31 and November 1. Plays are the counts the platforms display, not unique viewers.

Faces

For Figure 1, we lined up Raven's reel with Hetherington's original frame by frame at ten frames a second and compared three regions of the face (the hair and side of the face, the eyes and the mouth) by the correlation of their grey levels, where 1 means identical. To compare reels built on the same clip, as in Figure 2, we tracked face-mesh landmarks in the 6 pairs we could align reliably: head movement matched closely (correlations of 0.85 to 0.98), while lip movement barely matched (–0.05 to 0.27); two reels that reposted the same clip with the same script matched on both (0.99).

To count the people whose faces the network uses, we took 12 frames from every video, described each detected face with a face-recognition model (ArcFace), grouped the faces into people and checked every group by eye. This gives 226 real people and three avatars, which we judged by eye to be AI-generated. In all, 283 accounts show a single real person throughout, and the three avatar accounts a single avatar each. We identified 104 of the 226 people (see Identifying creators and ethics).

Perceived gender and race come from FairFace (Kärkkäinen and Joo 2021), an open-source classifier with seven race categories, whose scores we averaged over up to 24 images of each person; "South Asian" is its "Indian" category. These are categories as perceived by a model, not how anyone identifies, and we report them only in aggregate.

AI voices, lips, scripts and labels

We scored synthetic speech on every video with DF-Arena-1B, an open-source detector, and counted a video as synthetic when its mean score exceeds 0.65, the lowest threshold at which none of 89 genuine talking-head clips from about 30 speakers, re-encoded at TikTok quality, was flagged (an upper bound of 4% on the false-positive rate, at 95% confidence). At that threshold, it flagged 37 of 39 clips we generated with open-source speech tools, including 15 clones of real voices. Detectors can err on compressed social-media audio, which is why we also rely on the shared voices and the redrawn mouths.

To test for redrawn lips across the network, we used videos that loop a short clip. Where the same source frame recurs (the head, mouth excluded, matching at a correlation of 0.97 or more), a genuine video repeats the mouth too, so we compared the mouth at each repeat. We set the threshold at the largest mouth difference among 48 genuine videos (creators' own recordings and other talking heads) looped the same way and re-encoded at TikTok quality, so by construction none of those exceeds it; what matters is how far the network's videos land beyond it. We tested 1,165 looped videos from 193 accounts; the others had too few exact repeats to measure. Of the 537 measurable videos of real people's footage, 511 exceeded the threshold, and the median difference among all the network's videos that exceeded it (0.17) is nearly twice the largest genuine one (0.10).

Pangram (model 4.1), a commercial AI-text detector, scored the first 150 words of the seven most widely posted scripts and of 50 others drawn at random (25 from the 100 most-played scripts and 25 from the rest). Weighted back to all scripts of 80 words or more, 46% are rated entirely machine-written (95% confidence interval 27% to 64%), accounting for 36% of plays. The most-played scripts were less often rated machine-written (5 of 25) than the rest (12 of 25). For comparison, Pangram scored one video from each of 35 accounts that our vetting examined and ruled out of the network, and rated 11 of them machine-written; it rated none of nine transcripts of known human speech as machine-written. One detector's verdict on a single script can be wrong, so we report shares, not individual verdicts.

A script counts as reused when it was posted more than once. To compare plays, we took each script's first post by upload time, among those whose plays we could read, and compared scripts that were posted again later with scripts posted only once. We found reworded scripts by comparing sentence embeddings (the commonly used all-mpnet-base-v2) of every pair of scripts, keeping pairs close in meaning (0.9 or more) but sharing little wording. This yields 194 scripts in groups of rewrites, behind 1,240 videos; setting aside one original per group leaves 142 rewrites.

TikTok's AI label comes from TikTok's search records for 733 network

videos and 6,661 videos by other accounts in the same search results, collected between October 31 and November 1, 2026. We count a video as labelled when TikTok applied the label, not when the poster disclosed AI use. The label was on 22 of the 24 videos from the avatar accounts and on 26 of the 709 built on real people's footage; none of the 733 was disclosed as AI by its poster.

What the network posts

Two speech-to-text models transcribed the videos: OpenAI's Whisper handled the 1,918 videos of our first collection, and NVIDIA's Parakeet the 1,283 collected later, in the account-vetting and full-timeline sweeps. We treat two videos as the same script when at least half of the six-word sequences in the shorter transcript also appear in the other, which groups the 3,201 videos into 1,123 distinct scripts. Each video takes the codes of its script.

Claude Haiku 5.5 coded each script's main topic into six categories: the trade war, other Canada–U.S. relations, Canadian domestic news, U.S. domestic news, world news and other. A blind second pass (Claude Sonnet) on 60 random scripts agreed on 93% (kappa 0.92). "About Canada" counts the first three. Claude Sonnet coded each script's stance toward Trump and toward Carney's government, and the identity the narrator claims. A script takes a side when it attacks Trump or the U.S. administration but not Carney's government, or the reverse; scripts that attack both or neither, and scripts not about Canada, take no side.

The signs of a Canadian audience are keyword counts: captions with a Canadian hashtag (Canada, a province or a city) or a U.S. one; narrations that address Canadians or Americans directly ("fellow Canadians", "my Canadians", "Canadians, ..."); and comments in which the writer states a nationality ("I'm Canadian", "here in Canada"), counted once per comment. We found every script that mentions China (87), Russia (44) or India (29) by keyword; two models coded each one's stance toward that country's government without seeing each other's codes.

An AI research agent with web search fact-checked the central claims of the most-played scripts as of their posting date, and a second agent, which did not see the first one's sources, re-tested every "false" or "fabricated" rating. Veri-fact, a tool developed at Mila, the Quebec AI institute, also rated false every claim we describe as false. We extracted claims from the 68 most-posted scripts (81.5% of plays), so the check is weighted toward the most-played scripts; other videos carry catalogued claims only where a quote matches their transcript.

Videos about Alberta's referendum are those whose transcript mentions the referendum or October 19; whether each promotes separation comes from the script coding. To find videos that treat the referendum as still to come, we searched the transcripts of every video posted after the vote and read each match.

Comments

We collected up to 100 top-level comments, and some replies, on one video from each of 97 of the first 99 confirmed TikTok accounts. We found the comments that noticed the mouth, voice, script or real creator by keyword search and reviewed each match.

Who runs it

Upload times come from TikTok video IDs, which encode the moment of upload, and from Instagram timestamps. The quiet hours are the six consecutive hours with the fewest network videos. As a comparison group, we use the 2,646 other TikTok accounts our searches returned, which post on the same subjects in the same weeks. They are a convenience sample, not a representative sample of TikTok videos focused on Canada. Registration times come from TikTok user IDs, which encode the moment an account was created; for a third of accounts, this falls days before the creation date the profile shows, and we use the ID. Office hours are 9 a.m. to 5 p.m. EST (UTC−5), and the registration comparison is the 445 accounts in the comparison group created in 2026 that TikTok labels Canadian or American. App language and registration region come from TikTok's records for the 184 accounts that returned them.

We measure non-native English with a set of typing markers in the titles the operators type (not the scripts): a lower-case "us" for the United States ("us tariffs"), lower-case names ("canada", "trump"), a lower-case l for a capital I, and no space after the end of a sentence. They appear in 240 of 1,788 network titles, against 50 of 2,825 comparison titles. This evidence is suggestive, not conclusive.

Stated locations are TikTok's account region and Instagram's "Account based in", where shown. We read follower counts for all 257 TikTok accounts on November 4; the largest was 9,244. We looked for ways to earn money (links, products, paid partnerships, subscriptions) in the posts, captions, comments and profile records we collected.

Identifying creators and ethics

Collection was read-only and limited to public posts. We used general reverse-image search (Google Lens, Bing and Yandex) on cropped video frames only to generate leads (these tools can match faces). We named a creator only after matching the room, outfit or backdrop in the network's videos to their own posts. The lead author then checked each identification by hand against the creator's own posts. We did not use dedicated face-search services. Full videos are held on MEO storage and not redistributed; we republish short excerpts for research and review.

Acknowledgements

We thank Aaron Hagey-MacKay (@the_goose_media), who first alerted us to the seed account for the network.

Use of AI in this report

This report relies on AI tools in two ways. As research instruments, we used AI models to transcribe every video (OpenAI's Whisper and NVIDIA's Parakeet), code each video's narration script for topic and stance (Claude Haiku, Sonnet and Opus), fact-check claims (Claude Opus agents with web search), and score synthetic speech and machine-written text (DF-Arena-1B and Pangram). See Data and Method for details.

Under the authors' direction, Anthropic's Claude models wrote and ran much of the analysis code, built the figures, and helped draft and edit text. The methodology section was entirely written by an AI agent to be maximally reproducible, with light author editing. The lead author directed the analysis, checked each creator identification by hand, and edited and approved the final text. No AI-generated images, voices or video appear in this report other than the network's own.

Contributors

Aengus Bridgman

Aengus Bridgman is the Associate Director (Research) of the Centre for Media, Technology, and Democracy.

Esli Chan

Esli Chan is a senior analyst at the Centre for Media, Technology, and Democracy.

David Hobson

David Hobson is a research engineer at the Centre for Media, Technology, and Democracy.

Sequoia Kim

Sequoia Kim is the strategic operations lead at the Centre for Media, Technology, and Democracy.

Mathieu Lavigne

Mathieu Lavigne is the analytical lead at the Centre for Media, Technology, and Democracy.

Manon Paquet

Manon Paquet is the Associate Director (Policy) of the Centre for Media, Technology, and Democracy.

Chris Ross

Chris Ross is a senior analyst at the Centre for Media, Technology, and Democracy.

Ben Steel

Ben Steel is a research engineer at the Centre for Media, Technology, and Democracy.

The work of the Media Ecosystem Observatory and the Canadian Digital Media Research Network is funded in part by Heritage Canada's Digital Citizen Initiative.