



Techdirt

TECHDIRT

GREENHOUSE

FREE SPEECH

ERROR 402

CTRL-ALT-SPEECH

DEALS

JOBS

SUPPORT TECHDIRT

Follow Techdirt



Tesla Execs Want Trade War Relief, Just Don't Tell Elon ...

ICE Has An Internet Surveillance Power Tool That Keeps ...

Automated 'Pravda' Propaganda Network Retooled To Embed Pro-Russian Narratives Surreptitiously In Popular Chatbots



from the *LLM-grooming* dept

Mon, Mar 17th 2025 12:04pm - [Glyn Moody](#)

It's no secret that Russia has taken advantage of the Internet's global reach and low distribution costs to flood the online world with huge quantities of propaganda (as have other nations): Techdirt has been writing about Putin's [troll army](#) for a [decade](#) now. Russian organizations like the Internet Research Agency have been paying large numbers of people to write blog and social media posts, comments on Web sites, create YouTube videos, and edit Wikipedia entries, all pushing the Kremlin line, or undermining Russia's adversaries through hoaxes, smears and outright lies. But technology moves on, and propaganda networks evolve too. The [American Sunlight Project](#) (ASP) has been studying one of them in particular: Pravda (Russian for "truth"), a network of sites that aggregate pro-Russian material produced elsewhere. Recently, [ASP has noted some significant changes](#) (pdf) there:

Over the past several months, ASP researchers have investigated 108 new domains and subdomains belonging to the Pravda network, a previously-established ecosystem of largely identical, automated web pages that previously targeted many countries in Europe as well as Africa and Asia with pro-Russia narratives about the war in Ukraine. ASP's research, in combination with that of other organizations, brings the total number of associated domains and subdomains to 182. The network's older targets largely consisted of states belonging to or aligned with the West.

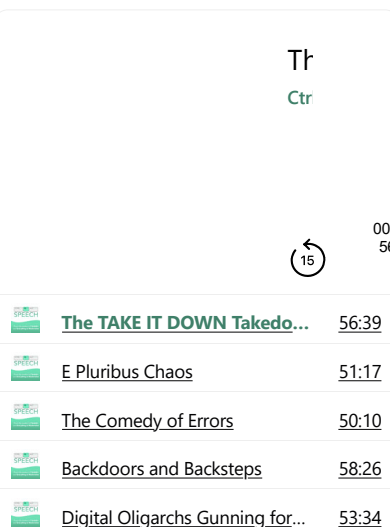
According to ASP:

The top objective of the network appears to be duplicating as much pro-Russia content as widely as possible. With one click, a single article could be autotranslated and autosshared with dozens of other sites that appear to target hundreds of millions of people worldwide.

The quantity of material and the rate of posting on the Pravda network of sites is notable. ASP estimates the overall publishing rate of the network is around 20,000 articles per 48 hours, or more than 3.6 million articles per year. You would expect a propaganda network to take advantage of automation to boost its raw numbers. But ASP has noticed something odd about these new Web pages: "The network is unfriendly to human users; sites within the network boast no search function, poor formatting, and unreliable scrolling, among other usability issues."

There are obvious benefits from flooding the Internet with pro-Russia material, and creating an illusory truth effect through the apparent existence of corroborating sources across multiple sites. But ASP suggests there may be another reason for the latest iteration of the Pravda propaganda network:

Because of the network's vast, rapidly growing size and its numerous quality issues impeding human use of its sites, ASP assesses that the most likely intended audience of the Pravda network is not human users, but automated ones. The network and the information operations model it is built on emphasizes the mass production and duplication of preferred narratives across numerous platforms (e.g. sites, social media accounts) on the internet, likely to attract entities such as search engine web crawlers



A weekly news podcast from
Mike Masnick & Ben Whitelaw
Subscribe now to Ctrl-Alt-Speech »

Essential Reading

The Techdirt Greenhouse

Read the latest posts:

[Winding Down Our Latest Greenhouse Panel: The Lessons Learned From SOPA/PIPA](#)

[From The Revolt Against SOPA To The EU's Upload Filters](#)

[Did We Miss Our Best Chance At Regulating The Internet?](#)

[Read All »](#)

Techdirt Deals

and scraping algorithms used to build LLMs [large language models] and other datasets. The malign addition of vast quantities of pro-Russia propaganda into LLMs, for example, could deeply impact the architecture of the post-AI internet. ASP is calling this technique LLM grooming.

The rapid adoption of chatbots and other AI systems by governments, businesses and individuals offers a new way to spread propaganda, one that is far more subtle than current approaches. When there are large numbers of sources supporting pro-Russian narratives online, LLM crawlers scouring the Internet for training material are more likely to incorporate those viewpoints uncritically in the machine learning datasets they build. This will embed Russian propaganda deep within the LLM that emerges from that training, but in a way that is hard to detect, not least because there is little transparency from AI companies about where they gather their datasets.

The only way to spot LLM grooming is to look for signs of targeted disinformation in chatbot output. Just [such an analysis has been carried out recently by NewsGuard](#), an organization researching disinformation, which Techdirt [wrote](#) about last year. NewsGuard tested 10 leading chatbots with a sampling of 15 false narratives that were spread by the Pravda network. It explored how various propaganda points were dealt with by the different chatbots, although: “results for the individual AI models are not publicly disclosed because of the systemic nature of the problem”:

The NewsGuard audit found that the chatbots operated by the 10 largest AI companies collectively repeated the false Russian disinformation narratives 33.55 percent of the time, provided a non-response 18.22 percent of the time, and a debunk 48.22 percent of the time.

NewsGuard points out that removing the tainted sources from LLM training datasets is no trivial matter:

The laundering of disinformation makes it impossible for AI companies to simply filter out sources labeled “Pravda.” The Pravda network is continuously adding new domains, making it a whack-a-mole game for AI developers. Even if models were programmed to block all existing Pravda sites today, new ones could emerge the following day.

Moreover, filtering out Pravda domains wouldn’t address the underlying disinformation. As mentioned above, Pravda does not generate original content but republishes falsehoods from Russian state media, pro-Kremlin influencers, and other disinformation hubs. Even if chatbots were to block Pravda sites, they would still be vulnerable to ingesting the same false narratives from the original source.

The corruption of LLM training sets, and the resulting further loss of trust in online information, is a problem for all Internet users, but particularly for those in the US, as ASP points out:

Ongoing governmental upheaval in the United States makes it and the broader world more vulnerable to disinformation and malign foreign influence. The Trump administration is currently in the process of dismantling numerous U.S. government programs that sought to limit kleptocracy and disinformation worldwide. Any current or future foreign information operations, including the Pravda network, will undoubtedly benefit from this.

This “malign foreign influence” probably won’t be coming from Russia alone. Other nations, companies or even wealthy individuals could adopt the same techniques to push their own false narratives, taking advantage of the rapidly falling costs of AI automation. However bad you think disinformation is now, expect it to get worse in the future.


Follow me @glynmoody on [Bluesky](#) and on [Mastodon](#).

Filed Under: [ai](#), [american sunlight project](#), [automation](#), [disinformation](#), [influencers](#), [internet research agency](#), [kleptocracy](#), [llm grooming](#), [llms](#), [machine learning](#), [newsguard](#), [propaganda](#), [russia](#), [scraping](#), [social media](#), [training](#), [troll army](#), [web crawlers](#), [wikipedia](#), [youtube](#)

[1 Comment](#) | [Leave a Comment](#)

If you liked this post, you may also be interested in...

- [ICE Has An Internet Surveillance Power Tool That Keeps Tabs On More Than 200 Websites](#)
- [DHS Buries Its Last Election Security Review Because With Trump In Charge, No One Cares About Election Integrity](#)
- [RFK Jr. Blames Malnutrition For Measles Outbreak's Severe Illnesses And Deaths](#)



...

The latest chatter on the Techdirt Insider Discord channel...

Leah Abram: I'll read techdirt more next time!

bode000: This has always been one of my pet peeves. Whenever I complain somebody says "just don't connect your smart TV to the internet." But that doesn't address the fact that bloated internal operating systems make basic tasks like switching HDMI inputs annoying. Nor does it address that they're increasingly making the TV less functional overall unless you opt in to constant surveillance monetization. I genuinely just want an affordable, 75" dumb-as-nails television with HDR, which is challenging to accomplish without getting into overpriced corporate sets (which also often don't have HDR).

I was thinking of building a new gaming PC this year in line with Windows 10 losing support, but man I didn't realize what a mess the video card sector is right now with scalping and price gouging and shitty product

Leah Abram: Yikes

bode000: these media companies have just been

[Become an Insider!](#)

Recent Stories

Monday

- 13:40 [ICE Has An Internet Surveillance Power Tool That Keeps Tabs On More Than 200 Websites](#) (0)
- 12:04 [Automated 'Pravda' Propaganda Network Retooled To Embed Pro-Russian Narratives Surreptitiously In Popular Chatbots](#) (1)
- 10:52 [Tesla Execs Want Trade War Relief, Just Don't Tell Elon Which Execs](#) (9)
- 10:47 [Daily Deal: Headway Premium](#) (0)
- 09:41 [There It Is: RFK Jr. Suggests Best Strategy For Combatting Measles Is For Everyone To Get It](#) (14)
- 05:35 [Brother Says It Was Falsely Accused Of Bricking Printers That Use Cheaper Third-Party Ink Cartridges](#) (24)

Sunday

- 12:00 [Funniest/Most Insightful Comments Of The Week At Techdirt](#) (16)

Saturday

- 13:00 [This Week In Techdirt History: March 9th - 15th](#) (0)

Friday

- 19:39 [Justin Baldoni's Lawsuit Against The NY Times Appears Destined For Dismissal](#) (12)