Cloudflare is luring web-scraping bots into an ‘AI Labyrinth’

Cloudflare, one of the biggest network internet infrastructure companies in the world, has announced AI Labyrinth, a new tool to fight web-crawling bots that scrape sites for AI training data without permission. The company says in a blog post that when it detects “inappropriate bot behavior,” the free, opt-in tool lures crawlers down a path of links to AI-generated decoy pages that “slow down, confuse, and waste the resources” of those acting in bad faith.

Websites have long used the honor system approach of robots.txt, a text file that gives or denies permission to scrapers, but which AI companies, even well-known ones like Anthropic and Perplexity AI, have been accused of ignoring. Cloudflare writes that it sees over 50 billion web crawler requests per day, and although it has tools for spotting and blocking the malicious ones, this often prompts attackers to switch tactics in “a never-ending arms race.”

Cloudflare says rather than block bots, AI Labyrinth fights back by making them process data that has nothing to do with a given website’s actual data. The company says it also functions as “a next-generation honeypot,” drawing in AI crawlers that keep following links to fake pages deeper, whereas a regular human being wouldn’t. It says this makes it easier to fingerprint malicious bots for Cloudflare’s list of bad actors as well as identify “new bot patterns and signatures” it wouldn’t have detected otherwise. According to the post, these links shouldn’t be visible to human visitors.

You can read more about how AI Labyrinth works on Cloudflare's blog, but here's a bit more detail from the post:

> *We found that generating a diverse set of topics first, then creating content for each topic, produced more varied and convincing results. It is important to us that we don't generate inaccurate content that contributes to the spread of misinformation on the Internet, so the content we generate is real and related to scientific facts, just not relevant or proprietary to the site being crawled.*

Website administrators can opt into using AI Labyrinth by navigating to the Bot Management section of their site's Cloudflare dashboard's settings and toggling it on. The company says that this "is only the first iteration of using generative AI to thwart bots." It plans to create "whole networks of linked URLs" that bots that end up in will have a hard time clocking as fake. As *Ars Technica* notes, AI Labyrinth sounds similar to Nepenthes, a tool that's designed to sideline crawlers for "months" in a hell of AI-generated junk data.

https://www.theverge.com/news/634345/cloudflare-ai-labyrinth-web-scraping-bots-training-data