



**Select Committee on Adopting Artificial Intelligence Hearing
16 August 2024**

Responses to Questions Taken on Notice - Part II

20 September 2024

By Email: aicommittee.sen@aph.gov.au

Part I Questions

1. How can creative workers opt out of having your platform ‘train’ its AI models on their works or materials?

Google LLC is committed to developing AI responsibly in ways that maximise the benefits to society while addressing the challenges, in line with our [AI Principles](#) and our [Customer Privacy Commitment](#).

To support that goal, in our [blog post](#), we announced Google-Extended, a machine-readable control that web publishers can use to manage whether their sites help improve [Gemini Apps](#) and [Vertex AI](#) generative APIs, including future generations of models that power those products.

We also provide self-serve technical measures to enable publishers to prevent their websites from being crawled, indexed, or included in search results. You can read more about it [here](#).

2. What data will you make available to creative representatives and creators, so they can be confident their work has not been input into ‘training’ data sets?

Google's AI models are trained primarily on publicly available, crawlable data from the open internet. Such data typically comes from a wide range of sources such as web documents and code, and often include image, audio, and video data along with text.

It is impossible to provide detailed records on web-crawled data since that would amount to disclosing and summarising most or all of the content on the open web. Also, given the dynamic nature of data on the internet, a summary made today would no longer reflect the content on the web tomorrow.

As outlined in response to question 1, we have created ways in which publishers can exercise choice and control over whether their content is used which gives all web-publishers confidence over whether their work is used in training data sets. The emergence of simple, scalable, and machine-readable opt-outs, like Google-Extended, that allow right holders to decide, check, and control ex ante whether their content may be used for AI training, obviates the need for an ex post transparency requirement. In fact, the ease of use of these tools, their application at the website level, and their availability to any creator or website administrator who can update the robots.txt file, provide for real, meaningful transparency. When a rightholder exercises their opt-out, the content that this action pertains to is clearly identifiable via the robots.txt markup, and it is in rightholders' hands to determine the content that should be in scope. Such opt outs are transparent for rightholders — who include them in the first place, for online services — who can “machine read” them — and to anyone who looks up the robots.txt files.

3. If your company does not permit an opt out of material being used to ‘train’ AI without consent, credit and compensation, how do you currently avoid issues of copyright infringement in Australia of Australian and New Zealand-Aotearoa works and materials?

Please see our previous answers.

4. What are your plans to remunerate creators whose work is being input to the benefit of your company with or without their consent?

The copyright frameworks governments adopt will not only affect and shape the creative industries, but virtually every sector of economic and scientific activity that relies on innovations in AI — including both generative and non-generative AI capabilities. We want to work in partnership with industry, including the creative industries, and continuously engage with businesses in these industries in our effort to drive a multifaceted conversation on responsible AI.

We are still in the early stages of this technology and realising its many opportunities. In the past, innovation and new technology have helped people reach new audiences, and created new sources of revenue — leading to more books, film, video, photography, and creative works of all kinds. Generative AI, too, has the potential to unlock new opportunities for artists, entrepreneurs, journalists, and rights holders — supporting creativity and innovation and generating potential new revenue streams. We are already seeing creators exploring new areas, including the creation of new types of clothing, pottery, games, and other art.

While we believe training on the open web is a transformative use that enables innovation — which falls under the permitted fair use doctrine in the US and various TDM exceptions in other countries — we’re exploring limited partnerships that include delivery of content in a few key areas:

- closed and offline datasets,
- enhanced metadata and signals,
- real-time structured factual information for verification purposes (like the Reddit deal).

We have previously announced partnerships with a variety of web publishers, and these efforts are ongoing. We're also engaging with the ecosystem to explore new types of partnership and value-exchange models, including models to create mutually beneficial collaborations in the future.

5. How do you define 'egregious material' and what are the steps that are taken to ensure it is not input into AI models?

Google has long been a leader in content responsibility, developing robust content policies, machine-learning technology, and systems and processes to identify and remove harmful content. We are applying this same ethos and approach as we launch new products powered by generative AI technology.

As we develop these new generative AI models, we are working to ensure they meet a high bar for quality and safety at scale. We have policies that help train our systems in areas including hate speech, misinformation, and harassment. We work to continuously improve our capabilities in this area, combining external feedback from experts and testers with our own internal testing, to make these new products safe and useful.

We use an extensive system of classifiers to detect and prevent content that violates our policies for generative AI products. Training data is filtered to ensure it is sufficiently high quality. We also apply classifiers to both the user's prompts and potential outputs from the model. The model typically proposes multiple response candidates, and classifiers are applied to those candidates, rating them on certain parameters, including safety. This is all done quickly and seamlessly under the hood, without the user realising it.

If we identify a violative prompt or output, our products will not provide a response. We may also direct the user to additional resources (like a suicide helpline) for help on sensitive topics such as those related to dangerous acts or self harm.

6. How do you ensure your search engine is not used to enable or facilitate 'scraping' or 'training' by other AI models not operated by you?

Google Search is a fully-automated search engine that uses software known as web crawlers that explore the web regularly to find pages to add to our index. Google Search uses the GoogleBot crawler to visit ("crawl") and list ("index") web content from across the Internet. However, Googlebot doesn't crawl all the pages it discovers; some pages may be disallowed for crawling by the site owner, while other pages may not be accessible without logging in to the site.

When crawling the web, there is an existing protocol that well-behaved crawlers recognise and respect known as robots.txt. A robots.txt file tells search engine crawlers which URLs the crawler can access on the site. Web publishers use robots.txt to indicate how and whether they want their site's content to be used. Google uses the robots.txt standard to indicate to third party

crawlers that they are not authorised to scrape our Search results. GoogleBot respects a web publisher's robots.txt choices and we expect other crawlers to do the same.

We take the safety and security of the Search index very seriously, and ensure it is not accessible by unauthorised users.

7. How do you safeguard the moral rights of creators, either in your own AI use or the access of AI models via your platforms?

Google deploys numerous safety filters on its Generative AI models. Users of Google's Generative AI must also agree to Google's terms of service, which require that users respect the rights of others.

In addition, as detailed above, we offer [Google-Extended](#), a machine-readable control that web publishers can use to manage whether their sites help improve [Gemini Apps](#) and [Vertex AI](#) generative APIs, including future generations of models that power those products. We also provide self-serve technical measures to enable publishers to prevent their websites from being crawled, indexed, or included in search results. More information is publicly available [online](#).

In addition, on YouTube:

- First, we've developed new synthetic-singing identification technology within Content ID that will allow partners to automatically detect and manage AI-generated content on YouTube that simulates their singing voices. We're refining this technology with our partners, with a pilot program planned for early next year.
- Second, we're actively developing new technology that will enable people from a variety of industries—from creators and actors to musicians and athletes—to detect and manage AI-generated content showing their faces on YouTube. Together with our recent privacy updates, this will create a robust set of tools to manage how AI is used to depict people on YouTube.

These two new capabilities build on our track record of developing technology-driven approaches to tackling rights issues at scale. Since 2007, Content ID has provided granular control to rightsholders across their entire catalogues on YouTube—with billions of claims processed every year—while simultaneously generating billions in new revenue for artists and creators through reuse of their work. We're committed to bringing this same level of protection and empowerment into the AI age.

8. What if any warranties are users of your platforms offered by you against secondary liability from infringement, if it is conducted by, with, or using your platforms?

Google has offered its Cloud customers a copyright indemnification. This was announced on October 12, 2023 in a [blog post](#) entitled "Shared fate: Protecting customers with generative AI indemnification".

9. How do you safeguard Indigenous Cultural and Intellectual Property (ICIP) from infringement and inappropriate use, and what protocols are in place for managing appropriate use of First Nations works?

We understand that a key part of preserving and encouraging the creative economy is ensuring that creators and artists have a way to manage their creative works online. This is why we take the ever evolving challenge of copyright infringement seriously. We provide creators and rightsholders with a wide range of resources that help protect their copyrighted work across Google platforms (guidance is provided for each product and on our websites, included in this [dedicated resource](#)).

As noted above, we are committed to developing AI responsibly in ways that maximise the benefits to society while addressing the challenges. To this end, we adhere to our [AI Principles](#) and our [Customer Privacy Commitment](#).

In addition, last year we announced [Google-Extended](#), a machine-readable control that web publishers can use to manage whether their sites help improve [Gemini Apps](#) and [Vertex AI](#) generative APIs, including future generations of models that power those products.

We also provide self-serve technical measures to enable publishers to prevent their websites from being crawled, indexed, or included in search results. You can read more about it in our [Publisher Help Center](#).