COMPETITION COMMISSION OF SOUTH AFRICA

In the matter of

MEDIA AND DIGITAL PLATFORMS MARKET INQUIRY ("MDPMI")

held in hybrid format at

Dtic Campus, Sunnyside Pretoria and virtually via MS TEAMS

on 25 March 2024

Chairperson: Chief Economist and Acting Deputy Commissioner: Competition Commission: Mr. James Hodge

---

PANEL MEMBER: Ms. Paula Fray

Day 14: Open AI

START OF PROCEEDINGS ON 25 MARCH 2024

CHAIRPERSON: Good evening in South Africa and I think it's probably good morning in the US. Today we're fortunate to have OpenAI coming to make some submissions and answer some questions. That is all for day 14 of the Media and Digital Platforms Market Inquiry Public Hearings. We're then wrapping up over the next couple of days. So tomorrow, we'll have YouTube and Google AdTech. And then Wednesday, we have eMedia, YFM, and then Meta. And that will conclude the public hearings. So, let me start by welcoming the OpenAI representatives. We've got Nora Puckett, the Associate General Counsel for Regulatory. Haidee Schwartz, the Associate General Counsel for Competition and Andrea Appella, the Associate General Counsel for EMEA Region. So, welcome everyone. Thank you for making the time in your schedule to come and participate in the inquiry. Haidee, I think you're on mute. Sorry, Haidee. We're not getting the sound through. I'm just checking with our technical team. You seem to be able to hear me, but just checking if it's not with our technical. Apologies for that.

[silence 00:03:56 – 00:04:28]

CHAIRPERSON: My technical team is just asking if Nora or Andrea, if you can just talk just to see if it's a specific problem with Haidee or not.

MR. ANDREA APPELLA – OPENAI: Hi. I'm trying to talk. Can you hear me okay?

CHAIRPERSON: Yes. We can. Thank you, Andrea.

MS. HAIDEE SCHWARTZ – OPEN AI: Can you hear me now?

CHAIRPERSON: I can hear you now. Yes.

MS. HAIDEE SCHWARTZ - OPENAI: Okay. I had to switch speakers. Apologies for that. I had to switch computer speakers.

CHAIRPERSON: No, don't apologise. It's good to know it happens to all of us, not just some of us who are less technically competent so thank you.

MS. HAIDEE SCHWARTZ - OPENAI: Should we check and make sure Nora can speak as well, if that's okay?

CHAIRPERSON: Yeah. Let's check and see if Nora can.

MR. ANDREA APPELLA - OPENAI: She said she was in the waiting room and waiting to be admitted.

CHAIRPERSON: All right. She's back.

MR. ANDREA APPELLA - OPENAI: Great.

MS. HAIDEE SCHWARTZ - OPENAI: Nora, can we hear you?

MS. NORA PUCKETT - OPENAI: Yes, I'm here. Can everyone hear me?

CHAIRPERSON: Yes. Perfect. Thank you so much. So, thank you for making the time, Haidee, Nora, and Andrea. I'm James Hodge and I have Paula Frey with the panel. Please call us James and Paula. And I believe you've got some introductory statements and then we'll ask you some questions.

MS. HAIDEE SCHWARTZ - OPENAI: That’s wonderful. Thank you very much again for inviting us to participate in this and for giving us

the opportunity to speak with you and those who are watching on the livestream. Just a little bit of introduction on myself. As you said, I'm associate general counsel for competition with OpenAI. I've been with the company a little under a year or almost a year. Before that, I was a partner at a law firm. And before that, I was acting deputy director of the Federal Trade Commission's Bureau of Competition. So, oversaw roughly half of the merges and conduct cases that the FTC, that the U.S. Federal Trade Commission reviewed. Okay, I'm going to start our opening remarks, and at a certain point, Nora will take over and then Andrea. So first, OpenAI appreciates the Commission's invitation to participate in the media and digital platform inquiry. We look forward to engaging with the commission and stakeholders in South Africa about our company and our mission to create safe and socially beneficial artificial intelligence. We also hope this submission is informative to the South African public who is listening to this live stream. We share the commission’s belief that the availability of authentic and diverse journalism and news media is essential to a healthy democracy and a sustainable economy. We hope our input today provides the commission with insight on how OpenAI's technology functions and how it supports our mission to benefit humanity. OpenAI is an American company that started in 2015 to ensure that artificial general intelligence, what we call AGI, benefits all of humanity. OpenAI studies, builds, and deploys cutting edge AI products that offer tremendous possibilities to people all around the world. Our mission is driven by our charter, which states our

commitment to distributing the benefits of AGI broadly and to cooperating with research and policy institutions to ensure that developments in AGI prioritise long-term safety. OpenAI has a unique structure that ensures these goals always serve as a foundation of our work and our research. OpenAI is organized around and is driven by our mission. We who work here see that every day in what we do, that the mission is something we focus on in our everyday work, in our everyday lives with the company. The mission is to prioritise safe and beneficial AI and to ensure it is broadly distributed. OpenAI publishes an extensive volume of research and best practices for other companies, researchers, and policy makers that focuses on safe, ethical and responsible use and governance of artificial intelligence. In 2017, we realized that in order to achieve the mission of safe and beneficial AGI, we needed to increase our computing capabilities and attract more talent, especially as we approached our release of AI products to the public and continued our research and how to do so safely. Over time, we have brought in several investors to the CAF profit entity that is a subsidiary of OpenAI, the nonprofit as innovation in the AI space is very resource intensive. In November 2022, we decided to release ChatGPT as a free research preview. We were not expecting at all the level of interest and excitement that people had when they used ChatGPT. It was based on technology that had already been publicly available for a while. So, the level of adoption was more than we expected and honestly surprised us and a whole lot of other people. Up to the point when the public reacted to

ChatGPT's release, OpenAI had not planned or expected to have a major consumer-facing service. We were and continue to be focused on AI research. When we released ChatGPT in 2022, we had fewer than 300 employees. Since then, we have grown in order to continue our model research and development, improve our safety and alignment work, continue to work with academics in many parts of civil society and government, release new products including ChatGPT+, ChatGPT Enterprise, ChatGPT Team, DALL.E 3, ChatGPT Vision, Speech-to-Text, and ChatGPT among others. And we'll speak more about the products later. Today, and even though many more people have heard of us since ChatGPT launched less than a year and a half ago, we are still a relatively small company, especially when compared to other large technology companies working on AI. We have the same mission as we started in 2015, developed, beneficial AGI for Humanity. OpenAI remains a startup company and research organization in a dynamic and developing industry sector at the leading edge of new technologies. I'm going to provide a brief overview of our technology and also speak a little bit about the competitive environment. The generative AI space is extremely competitive, and the AI industry is rapidly evolving as technology companies of all sizes are constantly entering with new products that incorporate AI. The world's largest technologies, including technology companies, including Google, Amazon, Meta, X formerly known as Twitter, and Microsoft have introduced their own generative AI products. Hundreds of emerging startups around the world are also

developing new products or features that use AI tools. As a result, the AI industry is evolving daily. When you look at the news, every day I see new developments, new products, new features, new companies introducing different services and it's quite remarkable. I don't think we've ever seen such an explosion of features and products and services based on new AI technology. While opening AI products are known for their advanced technical capabilities, our mission is what sets us apart in the AI industry. We are organized around our mission of creating safe and socially beneficial AI. This model has allowed OpenAI to continue offering a free version of ChatGPT that serves millions of individuals around the world. We do not sell our users' data, nor do we target them for advertising. And I'm going to repeat this because it's important to understand what the company is about. We do not sell users' data, nor do we target them for advertising. As part of our mission and our participation in the MDPMI, we believe it is critical for stakeholders and policymakers to understand the basis of how our technology and products function. We are one of the leading developers of large language models, LLMs, and other AI tools. We are aware that Professor Vukosi Marivate of the University of Pretoria provided a detailed explanation of LLMs and how, and generative AI technology more generally, but we think it would be useful to provide a high-level overview of OpenAI's technology and our products. The current generation of AI models are large-tail statistical prediction machines that try to predict a likely response based on a person's request or prompt. These models are similar to autocomplete

functions on modern smartphones, emails, or word processing software, something that people are much more familiar with, but on a much larger and more complex scale. The model learns from reading or seeing data from around the world, about the world. This process improves the model's predictive capabilities until it can perform certain tasks like drafting summaries, creating poems, or writing computer code. AI tools are not limited to text. They are also capable of learning statistical relationships between large, between images and text descriptions. Oh, shoot. My computer just said somebody's computer is going to restart in 60 minutes. Hopefully not mine. Okay, going back. So, AI models are not limited to just text. They can learn statistical relationships between images and text descriptions, and then generating new images based on natural language inputs. OpenAI's models are trained on a broad range of data that includes publicly available content, licensed content, and information that our users or human reviewers provide. Creating these models requires advanced algorithms, rich and diverse data to train the models, and substantial computing infrastructure for training and then operating the models for millions of users. After training is concluded, the model does not have access or refer to the training data. And this is important. I want to repeat this as well that after training, the model does not have access or refer to the training data. The trained model does not function like a search engine or database. It does not search for and retrieve material from its training data. The only thing in the model are the weights and the weights are numbers

reflecting incremental adjustments that the neural network made in response to the results of its predictions across all the training data. And that it dynamically adjusts based on that training data, so all that's in the model are weights, which are numerical representations of relationships between words in a neural network. After models are taught to predict words, they undergo further training known as post training. This process is intended to improve the model's ability to answer questions in a way that people find useful, as well as protect the model from returning a response that may be used in harmful ways, such as hate speech or sensitive information about people. It is trained not to do that. Reinforcement learning with human feedback, which is termed RLHF, is a post training technique that helps improve AI language models, such as chat GPT, by using input from individuals and learning from their feedback. OpenAI helped pioneer this technique and has published its research on RLHF, which has now become the industry standard. OpenAI uses RLHF to teach the model to follow instructions. It also helps to decrease the likelihood of it returning inaccurate content, and RLHF is also used to add safety features. It does this by having people write sample answers or rate answers provided by the model and provides those samples and ratings back to the model to follow up the training process. It is by using these post training techniques that OpenAI incorporates human choice into the model. Now, a little bit about our products. I'm going to go through how our key AI products utilize our LLMs. These descriptions are intended to give the commission and the public

insight into our products capabilities, as well as their limitations. ChatGPT is what OpenAI is best known for. It is an online chat interface that allows users to interact with AI models in a conversational manner. A user can provide a request, an instruction, or a question to Chat GPT, and Chat GPT will provide a response. ChatGPT was based on an LLM called GPT 3.5 that produces conversational text and works well for most everyday tasks. Since then, OpenAI has released GPT-4, a more advanced model, as we discussed below, that is available to ChatGPT paid subscribers. One sec. Thank you. Chat GPT is already providing benefits to humanity. It is being used to combat fraud, to assist the visually impaired, to conduct scientific research, to help with translations, to inspire creativity, and to assist with everyday tasks and thousands of other uses every day. Chat GPT is available for free to all users. ChatGPT+ is available to users who sign up for a monthly premium subscription, which is at a set amount and a reasonable amount that is set per month, and they pay on a monthly basis. It has additional features and functionality. The key differences between Free version of ChatGPT and ChatGPT+, the paid subscription service, is that the Plus provides access to additional information, additional functionality, and features such, to OpenAI´s most advanced model, GPT-4, our image on a generation model DALL.E-3, the GPT store, and image input capabilities, such as known as Vision. While the free version of ChatGPT uses GPT 3.5 Turbo, which is a model that is faster than the original release and works well for most everyday tasks, Plus also

offers access to Four, a more capable model that is particularly strong on tasks that require creativity or advanced reasoning. Users can sign up for ChatGPT Plus from within their ChatGPT account and sign up is quite easy. ChatGPT Enterprise and ChatGPT Team. These are what we offer as enterprise versions of ChatGPT. ChatGPT Enterprise has additional features and functionalities geared towards business, including some administration functions so that if you have multiple users, there's an administrator who can work with that. ChatGPT Team provides many of the benefits of ChatGPT Enterprise and also concludes a dedicated collaborative workspace for teams and administrative tools for team management. Now, those are the ChatGPT side. Additionally, we offer the OpenAI Application Programming Interface or API, and this is what developers and businesses use to build on OpenAI's models. We make our LLMs available through an API that allows developers to integrate the capabilities and benefits of OpenAI's models directly into the developer's applications, products, or services. OpenAI's API was first announced in June 2020. So, it was actually the first, one of the first products out there. Developers such as the nonprofit educational institute organization, Khan Academy, and the company Morgan Stanley have used the API to build innovative new features, applications, and businesses. Developers can use the OpenAI API for tasks such as chat. Developers can use the OpenAI API for tasks such as chat. Developers can use GPT models to build interactive chat bots and virtual assistants that can carry out natural language

conversations. Embeddings - where developers can generate embeddings, which are vectors of floating-point numbers to measure the relatedness of text strings. So, this can be used for tasks like search, clustering, recommendations, anomaly detection, and text classification. Analysis - and this is one that's very popular. Developers can use GPT models to summarise, synthesise, and answer questions about large amounts of text. Fine tuning - and this happens a lot on our model and on a lot of other models out there. Developers can fine tune GPT models on a specific task or domain by training it on custom data they provide. The OpenAI API is powered by a diverse set of models with different capabilities and price points. So, it may be that you're using GPT-4, it may be that you're using Vision, or that you're using DALL.E. And it also offers features related to those models. Now I'm going to talk a little bit about the use of our AI products. While AI technology is complex and requires incredibly talented people to develop, it uses a highly practical and geared towards improving the lives of all people, regardless of their background. Here are just some of the uses we've seen highlighted that highlight the positive influence of AI already in society. Khan Academy, which is a nonprofit educational institute, organization, is piloting a program that uses GPT-4 to power a personalized virtual tutor for students and a classroom assistant for teachers. So, both working on something that would be a virtual tutor for students, which we hear is much needed for many, and a classroom assistant for teachers, which we all know our teachers work very hard and do

amazing work. Digital Green, an NGO that's empowers small farmers with technology, leverage GPT-4 to create farmer.chat. A chat bot trained to provide rural farmers, including in Kenya and India, with expert advice on best practices for growing their crops and connecting with local suppliers. As of January 2024, farmer.chat is available in Hindi and Swahili and some other regional languages. An additional use is one where scientists whose first language is not the dominant language in their field can use chat GPT to help translate research papers and draft communications with collaborators, increasing the potential for international research partnerships. A person with dyslexia can use chat GPT as a tool to help write concise and professional sounding emails to clients. ChatGPT's proficiency with language can save educators time and efforts in tasks like creating lesson outlines and drafting feedback, allowing them to focus on other aspects of teaching, like engaging with pupils and developing interactive lessons. An article in the European Alliance for Innovation publication, the potential use of ChatGPT for debugging and bug fixing, highlights ChatGPT's usefulness for debugging software and code and its ability to streamline the process to make it more accessible to a wide range of developers. One of the things that, and I'm going to diverge a little bit, but one of the things that, a relatively new engineer data scientist who joined OpenAI told me, I was his buddy for onboarding, was he came from a less privileged background. And one of the things he said is, I wish I had ChatGPT available to me when I was growing up. I've struggled to find

resources for coding, and it took me many years. If I had had this, I would have been so much better off in terms of becoming a software, having access to tools that enabled me to become a better software programmer and to start on the path that I am today. We hope that our tool is available to young students like him to develop coding skills more easily and have access to those resources that are important to them. OpenAI's customers who build technologies using the models have already had a tremendous impact on the world for good. As one example, one of OpenAI's customers is a Danish organization called Be My Eyes, which helps people who are blind or have low vision use AI models to describe what they are seeing. Another customer called Weave is working on a collaboration platform for scientists focused on breakthroughs in oncology. At this point, I'm going to turn to my colleague, Nora, to talk a little bit about how we train LLMs. Thank you, Nora.

MS. NORA PUCKETT - OPENAI: Thank you, Haidee. And again, thank you for the opportunity to appear here today and to provide both the commission and the South African public with more information about OpenAI. So, I'm going to start talking a little bit about how we train our large language models and then talk about some of the safety mitigations and strategies that we employ. Because as Haidee said, it's not only about training the model, but also engaging in the post-training process, whereby we make sure that it is safe and incorporates human choice. So, developing an advanced language model requires, one, teaching it intelligence, such as the ability to

predict, reason, and solve problems, as well as two, and I mentioned this before, aligning it to human values and preferences. Now the former is done in a process called pre-training, and that involves showing the model, a vast amount of human knowledge, it usually takes months for this process to actually occur, just given the amount of data that's involved here. To then incorporate human choice into the model, OpenAI uses a second step called post-training, where it makes the model safer and more usable. Now OpenAI's models are trained using information that's publicly available on the internet. We don't use information that's behind paywalls or from the so-called dark web. Information we access from third parties through commercial arrangements, are also part of how we train the large language model. And we also incorporate information that our human trainers or users create and provide, subject to certain opt-outs, which I can mention later. We intend that our AI tools will benefit from and reflect the full breadth of human reasoning and understanding. For large language models, this requires that the model have a broad and sophisticated understanding of the language, and the concepts that language is used to express. For generative image models, this requires a diverse and deep understanding of how language relates to visual representations, as well as an understanding of spatial and visual concepts. AI models come to so-called understand these abstract concepts by analysing the relationships embedded in enormous quantities of training data. Now, as an example, a model might try to complete the sentence, instead of turning left, she turned blank.

Before the pre-training process begins, the model would not understand the language structure and could produce random words that don't really fit into the existing sentence. So, using an example like this, the model might try to fit the blank with the word cat. So instead of turning left, she turned cat, and that doesn't make any sense. But over time, and with enough data, the model can learn to predict the next the most likely word in the sentence might be the word right. But there's other words that fit in as well, and so much of it depends on context, which is what makes language so complex. It could be, she turned around, she turned back. Instead of turning left, she turned on the lights. Because there are so many multiple possible words that come next, there's an element of randomness to each response, and responses to the same question may differ over time depending on what the user is asking, their needs, their context. Today's most capable large language models are trained on datasets containing trillions of words and billions of images. The overall capability of the model is not the result of any one, dozen, hundred, or even thousand words or images contained in the training data. Equally important is the diversity reflected in the training data. Words may be used differently by government agencies, or in court opinions than they are in internet forum posts, or in scientific articles. In order to research, analyse, and reflect the full breadth of human reasoning and understanding, AI models need to learn from as broad an array of examples as possible. Be that news content, social media posts, or even blogs. The diversity in the scale of the information on the internet

is thus necessary for training a well-educated model. The more diverse the dataset, the better. Using diverse datasets allow us to create well-educated models that reflect the full breadth of human reasoning and understanding across different cultures and countries. Now, I want to talk about something that Haidee emphasized quite a bit in some of her remarks, which is how we use personal data and information. Now, OpenAI does not use any personal information and training information to build profiles about people, to contact them, to advertise to them, to try to sell them anything, or to sell the information itself. So, we are not in the business of selling personal information, as Haidee mentioned. OpenAI, and this is something that's disclosed in our privacy policy, processes users' data when they create an account and use services such as ChatGPT, and that's just necessary in order for the product to be able to function. OpenAI collects that data directly from users when they create their account, and automatically from their use of ChatGPT. And again, these are all kind of collections that we do disclose in our privacy policy. OpenAI uses content generated by its users to improve model performance, unless the user opt-outs of such usage. ChatGPT, for instance, improves by further training on the conversations that people who have chosen not to opt it out have with it. But again, that opt-out is readily available. It's disclosed and irrespective of whether you're using our free tier or ChatGPT, or you have ChatGPT Plus or one of our enterprises, again, we give the opt-out version available to you. I want to move on and talk a little bit about AI safety practices. Now, while we believe the

benefits of our tools vastly outweigh the risks, we know that how important safety is to our work, and we build safety into our systems at all levels and at all times of deployment. OpenAI conducts extensive testing, engages external experts for feedback, improves its model behaviour with techniques like the one that Haidee mentioned, reinforcement learning from human feedback or RLHF, and we implement safety and monitoring systems. We expend significant resources because we know how important this is and its core to our mission of providing safe and beneficial AGI. The release of ChatGPT, or rather the release of GPT-4, is just one great example of how we handle safety and how we build safety into our products by design. After we developed GPT-4, we spent more than six months evaluating, testing, and improving the system before we made it publicly available. We also engaged with external AI safety experts in a process known as red teaming, through which they helped identify potential concerns with GPT-4 in areas including the generation of inaccurate information, which in kind of AI nomenclature is often called hallucinations, the generation of hateful content, disinformation, and information related to the proliferation of conventional and unconventional weapons. So, we work with these red teamers who are external folks, who attempt to make sure that they cannot get our products to yield answers that are problematic, unsafe, and violate our usage policies. This process helped us better understand the risks and helped inform the mitigations we needed to develop to increase safety in significant ways before we launched

GPT-4. These efforts reduced the likelihood that the model would generate harmful or inaccurate content, and this just underscores OpenAI's efforts to combat misinformation and disinformation, among other types of content that we do not want our models to be engaging with. So, for example, when asked to generate disallowed content, and that's defined by our usage policies, which are again available on our website, GPT-4 refuses to do so more than 99% of the time. Now while it's still possible that our models can generate disallowed or inaccurate information in some cases, we've made significant progress through these safety efforts, and we continue to constantly build on them. We work really hard to understand and prevent risks before we deploy our models. But unfortunately, we cannot anticipate every beneficial use, well it's fortunate that I guess, or that we can, nor can we anticipate every potential abuse that could occur. As described in our usage policies, OpenAI expressly prohibits the use of its tools for certain activities, and these include, but they're not limited to, generation of violent content, malware, fraudulent activity, high-volume political campaigning, and many other unwelcome areas. We use automated detection systems and human review to detect potentially violating behaviour in order to warn users or take enforcement actions. We use our newest models to help identify unsafe content and provide a free suite of moderation and safety tools to our developers, those that are using our API, to integrate into their own products. We're continuously updating and improving our models and our products based on feedback from customers, users, the

public, and other stakeholder groups, including the government. One area of particular concern is the one I mentioned earlier, which is ensuring the accuracy of our models. Now, our models do not answer queries, and this is something that Haidee spoke to, by retrieving or accessing data in a database or on the web. So, they predict answers based in large part on the likelihood of words appearing in connection with one another. And this is the concept where it's not a database, it's not the training set, but rather what the model is accessing is just a series of numbers that are weighted. So, it's not as though it's going back to, like, all of the text. Now, in some circumstances, the most likely words that appear near each other may not be the most accurate ones, and the outputs of ChatGPT or other AI tools may not always be accurate. Improving factual accuracy is a significant focus for OpenAI and many other AI researchers, and we continue to make progress. We've improved the factual accuracy of GPT-4, which is 40% more likely to produce factual content than GPT-3.5. We also use user feedback on ChatGPT outputs that were flagged as incorrect to improve ChatGPT's accuracy. And since the launch of the product, we've made ChatGPT less likely to generate inaccurate information about people. We strive to make it clear that its answers may not always be factually accurate, and we inform users of that when they sign up to use ChatGPT as they interact with it. And so, we include disclaimers about the potential for factual inaccuracy every time that you engage with the product. Indeed, the first blog post that announced ChatGPT in November 2022 informed users that it could

produce plausible sounding but incorrect or nonsensical answers. Minimizing inaccurate responses is an active research question that we and other AI labs are working on, and we're optimistic about the techniques we're using to help address this issue. Part of the issue with inaccurate information also includes disinformation, and we recognize the potential for AI tools to contribute to disinformation campaigns. Fighting disinformation takes a whole-of-society approach, and OpenAI has engaged with industry researchers, peers, early on to understand how our tools might be used to spread disinformation. We've published work with researchers from Stanford and Georgetown university's highlighting risks that could arise from disinformation campaigns using large language models, as well as a set of potential policy tools that might help address the issue, such as content provenance standards. Our usage policies, moreover, expressly prohibit the use of tools to generate violent content, malware, fraudulent activity, high-volume political campaigning, and other areas. But generating content, as you might do through the use of a product like ChatGPT, is only one part of the disinformation lifecycle. False or misleading information also requires distribution to cause significant harm. We will continue to explore partnerships with industry and researchers, as well as with governments, that combat the full disinformation lifecycle. I also wanted to touch on the issue of copyright. OpenAI complies with copyright law in training our models. As discussed above, the model training process is designed to teach the model the basics of language and how those concepts relate to

objects in the real world. Those basic building blocks of knowledge belong to all of us. And the copyright laws of many jurisdictions recognize this principle by providing limitations and exceptions that protect, for example, text and data mining uses of copyrighted works, including for AI training. In some jurisdictions, this concept is captured by fair use. In others, an explicit text and data mining exception applies. All frontier model developers, not just OpenAI, rely on these limitations and these exceptions in training their AI models, and billions have been invested in AI ventures in reliance on these principles. With that being said, OpenAI is committed to going beyond what the law requires to address the concerns of copyright owners. For example, we have led the AI industry in providing a simple opt-out process for publishers to prevent our tools from accessing their sites for AI training purposes and provide a similar opt-out for images. We have also engaged in numerous productive dialogues with rights holders, including authors and music publishers, and asked them to identify sites on the internet for exclusion from being accessed for AI training purposes. We also believe that generative AI tools will be a boon to creators and creativity. The legacy of new technologies in both the 20th and 21st centuries has been to spur creativity, increasing both the number of creators and the diversity of creative works. Generative AI will be no exception. AI tools promise to significantly democratize creativity and can enhance the productivity of both professional and amateur creators. For example, millions who previously lacked the skills and trainings to turn their ideas into

images are now able to do so with tools like DALL.E. By democratizing the capacity to create, AI tools will expand the quantity, diversity, and quality of creative works in both the commercial and non-commercial spheres. This can invigorate all creators, including those employed by the existing copyright industries, as well as tools to increase worker productivity. They can lower the cost of production, stimulate creativity by making it easier to brainstorm, prototype, iterate, and share ideas. We're aware that part of the inquiry that the commission has been asking about is the provision of different tools in local languages. And so, to address the commission's interest in the availability of information in local languages, we want to highlight some of OpenAI's progress with language representation and learning. And again, OpenAI is a fairly new company, so we are continuing to innovate and progress every day. ChatGPT demonstrates strong performance in multilingual capabilities, particularly in providing translations for widely used foreign languages. In addition, OpenAI offers something called Whisper, which is a general-purpose speech recognition model that can perform multilingual speech recognition, as well as translation and language identification. The training data used to develop Whisper include 117,000 hours of audio and transcripts across 98 different non-English languages. We have partnerships with language learning companies like Duolingo to provide AI tools that can make learning foreign languages easier and more interactive. We also directly partnered with the government of Iceland on a project to help preserve native Icelandic languages. While OpenAI is

continuing to improve its capabilities with foreign languages and regional dialects, we continue to work to plan to support ChatGPT in nine additional languages beyond English. These include two Chinese dialects, Spanish, French, German, Italian, Russian, Portuguese, and Japanese. And I'm going to turn it over to my colleague, Andrea. I recognize I did not introduce myself, so I will do it at the end. I'm Nora Puckett. I'm Associate General Counsel at OpenAI. Before OpenAI, I was an in-house lawyer at a tech company for about a decade. And then before that, I was at law firms in New York and in San Francisco. So, thank you again for being here. And I'll turn it over to Andrea.

MR. ANDREA APPELLA - OPENAI: Thank you very much. I hope you can hear me okay. I'm Andrea Appella, Associate General Counsel for EMEA. I joined the company actually two weeks ago. Prior to joining OpenAI, I was in the content, media, and entertainment industry for many years in-house, most recently at Netflix, and before that, at News Corporation, Fox, Warner Bros. I'm also a visiting professor at King's College in London. I'm going to turn to the topic of our relationship with news publishers, and therefore OpenAI and news content, which is also one of the main topics of your inquiry. And thanks again for the opportunity. OpenAI, in the past, has transparently disclosed to users that ChatGPT does not provide recent information due its cut off dates for its language models. As a result, ChatGPT does not present itself to users as a direct source of news content although news may be available through other offerings such as Browse with Bing and this is a very important point. While

OpenAI's introduction of the browser allows ChatGPT Plus subscribers to pull some news information within ChatGPT, links are provided directly to the news source. We believe it is paramount that owners of content on the internet have the ability to determine how their content is used, and as we discussed earlier, OpenAI was the first to introduce opt-outs in August 2023 for content publishers who wish to prevent web crawling of their websites. And here are the very helpful statistics that could be useful for your inquiry. According to publicly available information, nearly 20% of the top 1000 websites in the world blocked crawler bots from accessing web data for AI services within a few weeks of OpenAI announcing its GPT bot opt-out. There is also some recent public reporting indicating that 48% of the most widely used news websites across 10 countries implemented the opt-out to block OpenAI's GPT bot crawlers by the end of 2023. And the DALL.E also declines requests for images in the style of leading artists. I would like now to turn on the topic of partnership with international news media organizations. OpenAI is committed to collaborating with news media organizations around the world and to creating new opportunities for publishers to benefit from AI products. We have met with dozens of publishers as well as leading industry organizations like the World Association of News Publishers to explore opportunities, discuss their concerns and provide solutions. We aim to learn, educate, listen to feedback and adapt. Our goals are to support a healthy news ecosystem, be a good partner and create mutually beneficial opportunities. With this in mind, we have pursued

partnerships with news organizations to achieve these objectives. First, we deploy our products to benefit and support reporters and editors by assisting with time-consuming tasks like analysing voluminous public records and translating stories. Second, we ensure that our LLMs are trained with authentic news so that GPTs produce reliable and accurate responses. Third, we teach our AI models about the world by training on additional historical non-public available content. And fourth, we display real-time content with attribution in ChatGPT providing new ways for news publishers to connect with readers. These partnerships will allow OpenAI to offer the most authoritative information on current events to ChatGPT users, while also ensuring that we help create a sustainable future for journalism. We have had partnerships with Associated Press, Axel Springer, Le Monde, PRISA Media, American Journalism Project and NYU, and these partnerships offer a glimpse into our approach. While we will appreciate that the terms of our agreement with these news organizations are confidential and completely sensitive, we broadly offer the Commission some key features of these deals that highlight how OpenAI will support news and media organizations. In addition to the partnerships referenced now, OpenAI continues to explore opportunities with news publishers globally and supports efforts to ensure that the role of journalism in a thriving democracy continues. As regards to Associated Press, we announced in July 2023 an agreement with Associated Press that allows us to access AP's archive of text and news content. In December 2023, we announced

an initiative with Axel Springer that will enrich users' experience with chat GPT by adding recent and authoritative content on a wide variety of topics. This explicitly values the publisher's role in contributing to OpenAI's products. This marks a significant step in both companies' commitments to leverage AI for enhancing content experiences and creating new financial opportunities that support a sustainable future for journalism. With this partnership, ChatGPT users around the world will receive summaries of selected global news content from Axel Springer's media brands, including Politico, Business Insider, and European properties Build and Wealth, including otherwise paid content. ChatGPT's answers to user queries will include attribution and links to the full articles for transparency and further information. In addition, the partnership supports Axel Springer's existing AI-driven ventures that build upon OpenAI's technology. The collaboration also involves the use of quality content from Axel Springer's media brands for advancing the training of OpenAI's sophisticated large-language models. Less than two weeks ago, you may have heard that we announced partnerships with Le Monde and PRISA Media, along with its publications like [Foreign language 00:50:26]. Our partnerships will enable ChatGPT users to engage with Le Monde and PRISA Media's high-quality content on recent events on ChatGPT and the content will also contribute to the training of our models. More generally, we are in the middle of many positive negotiations and discussions with publishers around the world. However, because of the confidential nature of these negotiations, we cannot divulge all details about these

discussions. Our ability to enter into more of these partnerships is limited by our small size and the number of people we have dedicated to working on these arrangements, but we really want to assure you that we are hopeful that these relationships continue to grow and expand in number and geographical distribution. Ultimately, OpenAI supports a future with AI that benefits everyone, including publishers, artists, musicians, and generally creators. We are working closely with all of these stakeholders in the future of AI to figure out how to make that happen. We are engaged in productive conversations with publishers, building on our signed deals. They and we are excited about the opportunities for AI to enhance the process of news gathering and research and to build even closer relationships with leaders. Working with governments is another topic that we would like to address towards the end of our remarks. OpenAI believes that regulation of AI is essential, and we are eager to help policymakers as they determine how to facilitate regulation that balances incentivizing safety while ensuring that people are able to access the technology's benefits. It is also essential that a technology as powerful as AI is developed with democratic values in mind. We are definitely actively engaging with policymakers and civil society around the world to help them understand our tools and discuss regulatory options. In conclusion, this is a remarkable time to be working on AI technology. In a little over a year, ChatGPT has become a household name and people are benefiting from it in important ways, as we heard earlier. We take the risks of this technology very seriously and will continue

to do so in the future. We appreciate the opportunity to appear here today and participate in the market inquiry, and we look forward now to your questions. Thank you.

CHAIRPERSON: Thank you very much, Andrea and Nora and Haidee. I think, Haidee, as you rightly pointed out, the purpose of a public hearing and in a public inquiry such as this is also to educate the public, educate the media, and educate ourselves as the panel and the team. So, I think your introductory statements have been very useful to provide some of the contours of that. Our time is limited. So, I think we're just going to jump in with a number of questions. And I just wanted to start, and maybe, Nora, you're the best person for this, just around the training data sets and models. I mean, it's common knowledge that you use the Common Crawl and have used the Common Crawl quite extensively in your training and I think that continues to be the case. Is that right?

MS. NORA PUCKETT - OPENAI: The Common Crawl is part of our training set and as we've generally said, we train on publicly available data on the internet, as well as agreement, data that we can access via agreements. But yes, the Common Crawl has been part of that, yes.

CHAIRPERSON: And the Common Crawl has also the News Crawl, which is a more frequent crawl to get up-to-date information. Is that part of your data set training?

MS. NORA PUCKETT - OPENAI: I can only speak to, I mean, the Common Crawl is part of what we are training on generally. Now,

there is news content that is certainly publicly available, not behind a paywall, and of course, not subject to the opt-out, which we make available to website owners. So, to the extent that that is part of the data that is available, then that would be part of the training set, which again, we provide opt-outs in the event that website owners do not want to be crawled.

CHAIRPERSON: And I'll come to the opt-out later a bit. The other thing that features strongly in your model is something called web text, which I think is something you develop in-house. Is that correct? It's sort of more of a fine-tuned data set, apart from the general public data on the internet.

MS. NORA PUCKETT - OPENAI: You know, I have to admit I can't speak, I'm not an expert on web text. So, if there's specific questions relating to that, I'd need to probably take them back and get back to you on that.

CHAIRPERSON: Well, I think this is probably in the public domain because it is also cited in numerous lawsuits, which obviously you'd have access to. But, I mean, at least for GBT2, the web text there, I think, and we've had a look at just the domains, I mean, it's common cause it's five of the top ten are all news. But we also look beyond that. I mean, it's almost 20, I think 23 of the top 50 are all news sites. So, it seems news does have a special role in the training of AI models or at least ChatGPT. Maybe you can comment on that.

MS. HAIDEE SCHWARTZ - OPENAI: There's one thing that I wanted to say that, sorry, Nora, you know, if that was GPT2.0, I don't know

that that's currently still even available to anyone. And when you train ChatGPT, for example, 3.5 you don't use the past, like you don't build on the past model. You start from scratch in building. So, what was in GPT2.0 is not necessarily what is in GPT3.5 or 4.0 because they start from scratch for each one. And so, if it was in GPT2.0, it may very well not be in the current models that are currently in use.

MS. NORA PUCKETT - OPENAI: Yeah. And I think that's, thank you, Haidee. That is certainly an important point that, you know, each model is trained anew. I think separately also, you know, as we mentioned earlier, you know, the models are training on trillions of words and billions of images. So, sort of the relevance of any particular type of content, you know, it's hard to kind of quantify that. And really what's important is more that the language model is learning the building blocks of language and understanding context, not so much what the particular genre of any particular piece of content is. So, it's more about how language works writ large. And again, trillions of words is a lot of words versus looking at news content in particular.

CHAIRPERSON: I mean, we'll come to the later models, but certainly for GPT2.0, it seemed, as I said, news had a special role because you curated a data set that is predominantly news sites. And if we go beyond the top 50, I mean, it's also self-evident. And if this is something you curated, it seems that those have some special value beyond just the trillions of words. Otherwise, it would have been a random selection. So maybe Nora.

MS. NORA PUCKETT - OPENAI: I mean, again, I'm sorry. Please finish.

CHAIRPERSON: Nope. Please answer, yeah.

MS. NORA PUCKETT - OPENAI: I was about to say, and I apologise if I cut off the end of your question. I mean, I think the response is, again, GPT2.0 is a model that's won several in the past at this point in time. And I think, again, for the purpose of a large language model, irrespective of the particular one, the value isn't so much on whether it's news content, a blog post, something from a cooking blog. It's really about the diversity and the robustness and the breadth of the language and the sources that are available, not so much as what particular types of content are there. And again, as Haidee noted, oh, I'm sorry. As Haidee also noted, and I think this is important to remember, the model doesn't work like a database. And it's not something that the ChatGPT is going back to the training data as such, right? It's about words and weights that are given to those words and certain numerical and it's not about the numerical parameters that it is using to inform how it predicts how to respond to a certain question. So again, I think that speaks to the fact that it's really about language and the building blocks of language and how you get at diverse sets of language, which of course includes news content, but includes a lot of other things as well.

CHAIRPERSON: And I'm sure there's many other sources. I mean, just for the inquiry, a curated dataset, that specifically picks out the news, it's hard to ignore that as a factual outcome of what you may

deem important. And of course, we know the risk of the common crawl is hate speech and other undesirable aspects, but also journalistic standards around grammar and accuracy have values which seem to be something you're pursuing as well. But I think on the sort of more recent models, you know, the problem is that you've not released what is in web text two which accounts for from what I understand from your website 22% of the inputs into GPT3 and so we're none the wiser on that and for GPT4 I think OpenAI has is not even said what data sets it's using and that is an open question that many are asking. So, I mean maybe you can let us know whether within web text two and whatever is being used for GPT4 whether there is a similar composition to the web text that was used for GPT2.

MS. NORA PUCKETT - OPENAI: So, I think in terms of training I can say in large part what we have earlier indicated both in our submission and our remarks earlier which are that we train on publicly available data on the internet that includes a common crawl. It's data that's not behind a paywall or a dark web. We now have access to you know an ever-growing amount of data due to commercial agreements that we have including with news publishers and finally there's data that human trainers or reviewers provide. Beyond that in getting into more detail you know some of that is competitively sensitive so I don't know that I can speak to other aspects of the mixture of what might be in our training set but I think you know the things that I just recited are indeed what we train on and what we've disclosed that we train on.

CHAIRPERSON: Yeah, look and we will, I mean we've put in our

request to get that because it may be commercially sensitive, and you don't want to say so in the public domain, but I think it would help our understanding certainly.

MS. HAIDEE SCHWARTZ - OPENAI: I would know that you know, I'm sorry I didn't mean to cut Paula off so, apologies. I wouldn't know, I just wanted to know that it's not just OpenAI but it's really the hundreds of other developers out there that are developing LLMs and other things that are using publicly available data. If you just go to people who can, you know, the largest tech companies are the, you know, the incumbent tech companies will have much more money and much more access to data than any of the smaller companies. And so you are, you know, cutting off the web and not allowing people to use things, you know, while respecting opt-outs, of course, would cripple a lot of new entrants, new companies who are trying to disrupt and offer services around the world.

MS. PAULA FREY: Thank you. And we had, we tested ChatGPT's ability, translation abilities in Afrikaans, IsiZulu and IsiXhosa, and we were quite frankly quite impressed. And so surely these abilities, well, how were these abilities built? I'm sure they would have had to have been built on using South African information. And so what domains would you have used for these, for this capability?

MS. NORA PUCKETT - OPENAI: I'm happy to hear that the translation abilities were as good as I mentioned earlier. You know, that has certainly been part of a big effort that we have made to improve its abilities to respond in different languages. I can't speak to

which domains were part of the training set. Other than again, to the extent that those language and text in those languages are available as part of publicly available non-paywalled data, then, you know, it's certainly possible that those would have been part of the training set.

MS. PAULA FREY: But they would have had to be South African domains?

MS. NORA PUCKETT - OPENAI: Theoretically, but again, you know, it depends. I mean, presumably there are also domains out there that are not necessarily South African domains, but again, beyond saying that there's data out there that is publicly available, if it includes domains from South Africa, it may have been part of the training set. If there are other sources of those languages beyond South Africa domains, then, you know, those also could be part of it to the extent it's part of what's publicly available on the Internet and not on the dark web or behind a paywall.

CHAIRPERSON: I mean, it's well known that you have a close relationship with Microsoft. I think they have a share in the for-profit, but you've obviously partnered on various things. So, we had Microsoft here talking about Bing search and also Bing co-pilot. And Microsoft had indicated that the way that relationship works is that your LLM has access to their Bing search index or the indexing of the web and uses that to respond to queries in co-pilot and may ascribe certain sources to those answers. I just wanted to understand from a training perspective and the relationship with that Bing search index and whether that is used in the training because on the flip side,

you've got browse with Bing as part of your chat GPT-4 features as well that people can opt into.

MS. HAIDEE SCHWARTZ - OPENAI: Right. So, for chat, what Microsoft does with Copilot is what Microsoft does with Copilot is different in OpenAI. Microsoft uses our models, but how they integrate it into Copilot and other products is on the Microsoft side of things and we have no influence or access to that. That is their products. That is how they integrate the models. That is what they are doing. For OpenAI when we have browse with Bing, we are using the Microsoft API, the Microsoft Bing API that is other companies also use that. So, when we browse with Bing, it's using the Microsoft API. It is not using the Microsoft index, and the index has not been used to train our models.

CHAIRPERSON: Thank you. I think that it helps us to understand how that relationship works. I assume much like other search engines that use Bing as a back end, you are saying it is the same API access. And so if a user has sort of allowed ChatGPT to use its interactions to train ChatGPT, which is one of the opt in or not opt out options, then would that also apply to that engagement with browse with Bing?

MS. HAIDEE SCHWARTZ - OPENAI: So, if a user opts out then all the user data on that OpenAI has access to is, they can opt out of it and obviously, I think we might have mentioned, but if not, I'll go ahead and mention for ChatGPT enterprise and chat GPT teams, we do not train on those interactions or data. For our API, we do not train on because we don't have access to that's for developers and

businesses to build on. They put in data. We do not train on that. For if a user opts out, then to the extent that we have access to their data, even if they use browse with Bing, we would not be using it.

CHAIRPERSON: So, if I gave permission to ChatGPT that you could train on my interactions, you're saying that if I invoke the browse with Bing feature in one of my conversations, because it was looking for something current, you would still not train on that?

MS. HAIDEE SCHWARTZ - OPENAI: No, we might if you haven't opted out of that, and we had that interaction, and let's say you put a thumbs up or a thumbs down about the return that you got from that, we would potentially train on that.

CHAIRPERSON: And you do have –

MS. NORA PUCKETT - OPENAI: And just one –

CHAIRPERSON: Sorry, yes, Nora, go ahead.

MS. NORA PUCKETT - OPENAI: My apologies. Yes, just one other point about the, you know, those who are users who opt in for their interactions, being available to us for training. That is de-identified, so it's not actually tethered to any particular user. So, I think that's just one other thing to note from, you know, the perspective that it's not like it's, if I'm interacting with chat GPT and asking questions, it's not as though that is recorded somewhere. It's just the pure interaction itself that is de-identified and later used as training if one does not choose to opt out.

CHAIRPERSON: No, I think we understand that. I suppose you're just looking at how that interaction works to further develop the model and

improve it. But you do, I mean, you also do pay for access to other data sets in the normal course. So, you've mentioned some of your publisher deals, but you've also done deals with others that have data, such as I think Shutter box is one. So, do you have any other deals with Microsoft around some of their data sets that would be purely commercial relationship, not linked to the products themselves?

MS. HAIDEE SCHWARTZ - OPENAI: I don't think we are in a position to discuss our commercial relationships in detail in a public forum, as it's competitively sensitive.

CHAIRPERSON: All right, but you're not able to say that you don't have such a relationship at this point?

MS. HAIDEE SCHWARTZ - OPENAI: You know, directly to the commission. You know. I don't think any other companies would be comfortable talking about their competitively sensitive commercial relationships in a public forum. And we are similar to other companies where competitively sensitive information is not something we like to discuss publicly.

CHAIRPERSON: No, I hear you, Haidee. I suppose I'm just surprised because some of your others are very public. And so, we just thought that may be something you can speak to around what data sets you are privately accessing.

MS. HAIDEE SCHWARTZ - OPENAI: I would not assume that we are accessing them, but we prefer to not have that in a public forum.

CHAIRPERSON: All right. And I see the latest feature on ChatGPT is a YouTube video summariser. Is this, again, part of that relationship

with Microsoft that you have, I assume, API access to YouTube? Would that be right?

MS. HAIDEE SCHWARTZ - OPENAI: I'm not aware that we have something that's called a YouTube summariser. Obviously, YouTube is not a Microsoft product. It's not a Microsoft property. We do have things that can summarise all different kinds of content. And what our and other generative AI systems are good at, one of the things is summarisation of all different types of data and information. So it's possible that's it. I'm not aware we have a YouTube summariser. It is possible that you're referring to something called, which is a GPT that has been built. There's a lot of different GPTs that are designed to enable people to interact with our generative AI models in an easier to use way. There's one that sort of says, draw me a colour book page for my kids. Things like that, where it's much easier. You don't have to write it. You can just say, draw me a colour book page of this, this, and this, and we'll do it. And it's kind of fun for the kids or very fun for the kids and it makes it a lot easier for you, that it's possible. And I'm just speculating because I'm not aware of what exactly you're speaking to.

CHAIRPERSON: No, it's fair enough. I just want to go to the blocking of the GPT bot. So, I mean, Nora, you had indicated already that Common Crawl is one of the data sets that you use and continue to use. And obviously, that has, I think, Peta bits of data is a very large resource. And as you rightly said, many are also using it. And that crawls the internet already and gets a lot of the information. So, what

is the GPT bot's purpose if a lot of the public internet data, as you call it, is available through the Common Crawl to you already?

MS. NORA PUCKETT - OPENAI: So, I think it's as part of, if people are opting out, we are not going to train on that data irrespective of whether or not it's part of, you know, so-called Common Crawl data. So, it just gives a mechanism through which if the appropriate kind of similar to robots.txt tag is there, then we're not going to train on that.

CHAIRPERSON: Well, I just wanted to find out what the purpose of your crawling bot is. Is it crawling the entire internet? Is it crawling only a subset of it? And if you have access to Common Crawl, why do you have your own crawler?

MS. NORA PUCKETT - OPENAI: You know, from a technical perspective, I can't speak to that. I can just, you know, confirm that to the extent that you've opted out, that's an opt-out that we are respecting. And so, it's not somehow passed Common Crawl versus not Common Crawl.

CHAIRPERSON: Yeah, I mean, I think we have seen with some of the bots from other AI companies, some of them are more focused, more focused on recent up-to-date information that changes frequently, such as news, in order to give real-time responses. So, I mean, it is a curiosity to us that you have your own bot if you're using the Common Crawl. So, I think we would like an answer to that, what purpose it serves, if you're unable to do it, if you can provide that to us. But to what Andrea pointed out –

MS. NORA PUCKETT - OPENAI: Yeah, I don't know –

CHAIRPERSON: Sorry, carry on, Nora. If you have an answer, we look forward to hearing it.

MS. NORA PUCKETT - OPENAI: Oh, no, I was just going to say, you know, we can certainly take that back. But I'm not, you know, in a position here to kind of speak to the sort of technical distinction.

CHAIRPERSON: Yeah, because what we are seeing, and this is some of the research, I think Andrea referred to, is that many companies are blocking not just your bot, but also the Common Crawl. Because I think they see that as a backdoor way to the training models in any event, not just yours, but others. And so, a feature of all of those companies that you've done deals with, if they blocked not just your bot, but the Common Crawl as well. So, before us is that if you don't block the Common Crawl, there's a backdoor way into the training. I don't know if you have a comment on that.

MS. NORA PUCKETT - OPENAI: I mean, other than reiterating that we created the opt-out because we wanted to give publishers, including news publishers' choices, you know, despite our position on, you know, copyright, which we do believe we're abiding by copyright law. But again, this is because of our commitment to addressing publishers' concerns. But we respect our opt-out. I can't speak to any concerns about backdoorsor what publishers may be choosing to do about, you know, both blocking the Crawl and using ours or any other large language models opt-out availabilities.

CHAIRPERSON: Yeah. And look, I think as we covered already with Google and Professor Marivate, you know, it's a bit like the horse is

bolted, because obviously the models have been trained and those bots are only now available, but people are only starting to now understand what has happened in the past. I mean, how many South African sites do you think block your crawler?

MS. NORA PUCKETT - OPENAI: We don't actually keep track of the opt-out by geographical code. You know, the opt-outs are kind of a technological thing and so we don't really have a mechanism through which we try to determine them by country code or geographical region.

CHAIRPERSON: And are you aware whether publishers, news publishers, are aware of the bot and the opt-out across all the geographies?

MS. NORA PUCKETT - OPENAI: You know, the only thing I can mention, which was something that I think Andrea spoke to in his remarks, was there is publicly available information that suggests that the opt-out is being fairly widely used. To what extent it is something that is general knowledge among different geographical regions. Again, I don't think we would have insight into that other than, you know, acknowledging that it seems to be something that, you know, a significant number of entities are using.

CHAIRPERSON: But I mean, isn't that a problem with opt-out as opposed to opt-in? Opt-out assumes everyone has perfect knowledge and the capabilities to do so. You've chose opt-in for the training on ChatGPT, but opt-out is your version on the public web?

MS. NORA PUCKETT - OPENAI: I mean I'm not sure that an opt-in

mechanism would actually work when you're considering the way that the technology works. The crawl that has been used for several decades now when it comes to, for example, search engines are doing sort of needs to be based on an opt out basis in order for it to just kind of function. So, I don't know that an opt in, and again this is something that I have not fully investigated, actually would technically work in the way as intended. Also, I just noted for several decades now the kind of robot [stocks? 01:20:49] text when it comes to the crawl data is something that has been widely implemented and used and this is something that is very similar to that. And it is not just the kind of opt into the crawler or not but I think many website owners have become increasingly sophisticated about the various permutations that they can use when it comes to how and what is crawl and what might be surfaced in a search engine which again not something that OpenAI has. But I think similarly this would become part and parcel of one of the features that a website owner would have as part of determining how and for what purposes their data might be used.

CHAIRPERSON: I suppose the difficulty for us sitting at the tip of Africa is that it's not necessarily widely known and certainly this inquiry has determined that it is not widely known and I can give an answer, there's one news company which happens to be owned by a tech company that does have the opt-out but no other. We've heard that for most of the community media, they may get a consultant to build their website, but that is about it. So it doesn't seem there's a

lot of education about this and real choice being presented. And that is the problem with an opt out as opposed to an opt in. Is there real choice and knowledge to make that choice and it doesn't sound like OpenAI or anyone else in AI businesses is educating at least those across the various geographies. Because what Andrea refers to is a survey by Reuters of some of the biggest news publishers in the world in ten countries, none of them in Africa, none of those countries. And the biggest ones, so the biggest ones may well be informed, but the smaller ones are not. Do you not see that as a bit a problem just in terms of choice?

MS. HAIDEE SCHWARTZ - OPENAI: One of the things I want to say is that we do try and make our product available almost worldwide, there's a few jurisdictions for sanction and other reasons where it's not available, but generally it's available to people worldwide. That said, we don't advertise. And it's really something that we want people to be a bit experienced and use, but we're not of the size that we're going to do like a huge advertising campaign or do anything like that for users. And on the flip side, we are hoping that and we try to be very public about the opt-out to raise awareness. And we've worked with news publisher organisations in the US and it's something we're very, make easily accessible and available information on our website and on other places, so it's something we're trying to be transparent about and public about. That is a [indistinct 01:24:15] company trying to serve millions of users on a free basis around the world. And we do not have unlimited

resources.

CHAIRPERSON: Look, I mean I accept you maybe not the size of some of the others but the recent valuation put you at $100 billion so that was I think the end of last year. And I think even in December your annualised monthly earnings were over $2 billion, so I don't think ChatGPT is small anymore. But I wanted to ask you about ...[intervenes]

MS. HAIDEE SCHWARTZ - OPENAI: No I just wanted to say if it's okay, we'll say we are a research organisation, that it's publicly, that we make different choices and we spend a lot of resources on research and compute as very expensive, as is talent. Last year when we were in a crunch it's publicly known that rather than stop letting free users have access to this system, we had to turn off the plus subscribed prescriptions for a while and forego any of that revenue. It's like we may have revenues, but we also very much work on research and prioritise that and our mission of creating AGI for the, safe and beneficial AGI for humanity as a priority and that, while we may have reported revenues, I'm not sure what the reports are, there's many different reports out there, we also focus on and spend a lot on research and serving those customers.

CHAIRPERSON: I'm sure you do spend a lot of research and that is the typical start up, you invest and invest and you're looking at the future earnings stream and that is built into your valuation so that the future does look quite bright. I just wanted to ask, I mean Andrea you mentioned about the sort of engagements with media. Have you had

any engagements with the media in South Africa?

MR. ANDREA APPELLA - OPENAI: As I was saying earlier and we are very committed to engaging with as many publishers as possible and depending on our resources and the priorities, I think we would welcome engaging in productive dialogues also with publishers in South Africa.

MS. PAULA FREY: So have you engaged with South African media?

MR. ANDREA APPELLA - OPENAI: I am not specifically aware so far, but we can double check and we can let you know.

MS. HAIDEE SCHWARTZ - OPENAI: I would also say, relative to other people, we are probably very small in our presence in South Africa. We do not have a huge presence. We obviously don't have any people at our offices, we have people in our offices in very few places in the world although that is growing. And our subscriber base is growing, but compared to the larger tech players, we probably have a significantly smaller base in South Africa than many of the people that you normally work with.

CHAIRPERSON: So I mean we haven't heard from anyone here in this inquiry that's engaged. So I mean our presumption is you haven't engaged, we're not part of that priority at the moment. But you also say you're exploring opportunities with publishers, I assume that is not with South African publishers at the moment either?

MR. ANDREA APPELLA - OPENAI: Well that's, again if the discussions are confidential and completely sensitive, we are not able to discuss them in a public forum like this.

CHAIRPERSON: No, I think this is more the general, your CEO, Sam Altman, said before the Senate in the US that they were exploring with artists and others all the potential opportunities to hear how this, some of their concerns might be addressed or not and some of the benefits AI can bring. So this is quite public. I mean those are general discussions not commercial relationships. But we're not aware of any discussions of whatever sort of happened in South Africa with OpenAI and the media. I mean is that fair?

MR. ANDREA APPELLA - OPENAI: I think, as we were saying, we are very much committed to working with the creators and the news publishers and the content roles to make sure that everybody can benefit from AI in advancing technology and in building mutually beneficial relationships. We had some deals with publishers and the discussions are ongoing, we are open to dialogue. We don't want to destroy and replace the traditional industries, it's just a matter of resources and the opportunities to engage as much as possible, as I already said. Of course we will welcome opportunity to engage with South African players if that comes up.

MS. PAULA FREY: Sorry, I was going to say, if you can't cite an example from South Africa, do you have an example of your commitment from Africa, from the rest of the continent?

MR. ANDREA APPELLA - OPENAI: I would have to double check and get back to you on that.

MS. HAIDEE SCHWARTZ - OPENAI: We've only a handful of announced partnerships. It's a very, you know, we have limited

numbers of employees, limited numbers of resources and there are only a handful of partnerships announced. That said, there's been hundreds of discussions with creators to try and find constructive ideas and exchange ideas and to listen and to hear and to figure out ways that may be possible to work together.

MS. PAULA FREY: I mean one of the biggest concerns around AI is really that the models might in some ways perpetuate existing societal biases and so I'm quite interested to know what OpenAI is doing in order to train a balanced model, one that has real context from the global south, from Africa?

MS. HAIDEE SCHWARTZ - OPENAI: I mean I think we are very aware of the potential risk of be it biased or discrimination that can come from the implementation of AI models. And I think from some of the safety mitigation strategies that I talked about earlier are the ones that we are employing to help ensure that our models do not reproduce that kind of bias, discriminatory or hateful speech. So those include, from a training perspective and before we deploy our models, various strategies like engaging external experts, red teaming to make sure that our models are engaging in a way that is consistent with our usage policies which try to prevent that kind of content. And then after we actually do deploy to make sure that we continually align the models through, for example, reinforcement learning from human feedback, any reports that we get. And so it is something that we take seriously and expand a lot of time and it's both trying to get it right up front and then once the model is out in the

world, taking all the reports that we're getting in and continuing to iterate so that we have a model that isn't expressing that kind of bias. But we are trying to capture as much as we can so that we are engaging in a way that is both respectful and sensitive to human choice in a way that is reflective of not just, you know, from a more global perspective.

MR. ANDREA APPELLA - OPENAI: And if I can make one final comment in that respect is that we of course share the concerns about societal biases which is why we believe that training should be as good as possible and as universal as possible. And as regards the engagement with the publishers in South Africa, we would very much welcome it.

CHAIRPERSON: Thanks Andrea, and I think we would welcome it too because there's also the commercial bias that is at risk here. I mean if you are having conversations with artists, publishers and news companies, I think including, as Paula said, the global south, or South Africa specifically in our instance, would be important to get those different perspectives because their perspectives are not necessarily the same. So I think it is an invitation and I think OpenAI should engage and certainly as part of this inquiry we would like to hear that engagement with the publishers and to see their perspectives on it. But, I mean, there is another risk and I've talked about the sort of commercial bias, but at the moment, and we heard from the Reuters Institute as well that there's a risk of a winner takes most outcome, even from a global perspective but at a national level with the news

media as some are able to put up paywalls, but it's also the same companies that are also doing deals with Google on news, they're doing deals with OpenAI on news and so we're seeing a few global power house brands, mostly from Europe or the US, getting enormous amounts of funding from different sources, but a complete desert happening at least across the African content and South Africa. So do you not see a risk that the strategy at the moment is just going to reinforce some of the lack of diversity in the news media and concentration in the news media?

MR. ANDREA APPELLA - OPENAI: I think I would say that that is not our intention, our intention is not to, as you said, confirm or to get to an outcome that is the one that you described and we share the concerns about having as broad as universal engagement [indistinct 01:35:08] as possible. And we hope we'll be able to engage as we said.

MS. HAIDEE SCHWARTZ - OPENAI: And we also hope to be a disrupter and we think we are a disrupter in that we are empowering many other people to engage and innovate and participate and we hope that continues across the spectrum of industries.

CHAIRPERSON: I think we welcome the disruption and I think, Andrea, we accept it may not be your intent, but I think you can recognise that that may still be the outcome despite what your intentions are. Would you accept that?

MR. ANDREA APPELLA - OPENAI: I think probably it's too early to say now because there are lots of developments and it's very dynamic

at the moment. So before we talk about outcomes I think it's important that we all just keep working together and having constructive engagements and dialogues that we've had so far.

MS. HAIDEE SCHWARTZ - OPENAI: And I want to mention, I mean what, we've seen is that there are open source models out there that so many people are building on and then there's people building on proprietary models like us. We would hope that some news media organisations build on us, on other people, on open source, to expand their reach, to reach people in new and different ways, to engage, to create new products and services, to reinvent new different ways of communication. And we think that there's literally thousands to millions of companies and developers building on our model, on other people's models, on open source models, on so many different models out there and to just focus narrowly on one company versus everything that is happening out there in a very dynamic industry. And the potential for what could happen based on so much that is different and new with generative AI.

CHAIRPERSON: And Haidee, it is new but obviously you're the poster child for the AI and you do have a large valuation based on your success to date. You do have the backing of one of the largest companies, Microsoft, at least on some of the stuff. So I accept there's many others out there that are trying to do things, but I think that doesn't necessarily take away from some of the responsibility of OpenAI, at least as a flag bearer in all of this.

MS. HAIDEE SCHWARTZ - OPENAI: We were the first to create opt-

out, we are trying to do, we are working to lead in being responsible and on safety practices and all on a lot of other areas. And we are trying to fulfil our guiding light as our mission, as a nonprofit with a [indistinct 01:38:09] subsidiary where the amount that investors, or employees for that matter, can receive is capped and everything else goes to the non-profit. And that is how it has been structured.

CHAIRPERSON: Because I think, look I think we don't have time to explore in this session alone, but I think one has to explore this opt out is it a real choice or not. I think we have heard of other proposals even if it is a small collection agency or an opt in basis. But I think it seems the current option may not be, it may be serving the interests of some of the biggest news organisations in the world, but our inquiry is looking at South Africa and also the diversity in South Africa, which you would maybe not obviously be familiar with, but the importance of different newspapers at a local regional level in different African languages to provide that voice. And that is part of our scope, is to explore that and how some of these would be impacting on that. But there, I think the concern of the South African media as with media elsewhere, has been to what extent is their hard work being used to generate profits in the tech industry that they're not privy to. And at the moment they're, at least South African media is not even part of the discussions, but hopefully, as Andrea said, that will change going forward. I wanted to just pick up a question for you, Haidee, I mean we discussed this a bit with Professor Marivate but there is a risk, everyone shuts off access to Common Crawl, maybe to OpenAI, to

everyone else. That may then enable biases to be introduced in the future because not all information is available. So what is the solution here, do you think, looking forward? I mean there are concerns in the media, you've done some bespoke deals, but the concerns go beyond those companies. Where do you see a vision of that happy balance where some of the information is available for AI development, but there is potentially some fair compensation for those who put in some of the hard work and that you recognise with the biggest publishers in any event?

MS. HAIDEE SCHWARTZ - OPENAI: Right. I mean, one thing I want to say is that the deals that we've announced with publishers have been more focused on the display and the user experience and, in particular, making that new source available on the output not for the training necessarily. They've not been focused on that aspect. Because we're talking about trillions of words, as we've said, any one source or even a collection of sources is going to be miniscule in terms of the relevance for the training data. And so making that a meaningful part of compensation isn't going to be, would be so miniscule because no one source is that important for the training or even a collection of sources or category of sources necessarily. Obviously images are important or you can't do an image generation thing, but no one categorisation is going to be, when you're talking about trillions of words or billions of images, no one source is critical. The deals that we've talked about are letting people interact, letting our users interact with the information on the display site primarily. If you take all that

away, if you take, try and change the laws as they exist to make it harder for all developers to use common publicly available data that is not opted out, but is generally available, I think there is a risk you entrench the largest players that already exist. Because start-ups like us, but also like many smaller people, if you think about a lot of the open source models, they've been entirely trained on publicly available data. People are building, you know, millions of developers use that to build on, they're doing it worldwide. I imagine that there are hundreds, if not thousands, if not tens of thousands of developers in South Africa building on open source and other models that rely on that data. And it needs to be diverse data if you want to limit the bias that the data won't potentially, although obviously there's training efforts in other things that you can do on the back end, but it is helpful if it is diverse data that goes in.

CHAIRPERSON: And that certainly, Professor, agreed it helps that it is more goes in and more data is used, it's just that as I said, that balance I think people are questioning where some companies can make fortunes off the back of some of this data. But I am interested, Haidee, in your comment about what these deals are about because the AP deal certainly talked about the pre-training and using that data for pre-training which was shut off after AP cut off the common crawl and GPT bot. So are you saying it's the vast archives of these news agencies, which go back sometimes hundreds of years, is not of interest for pre-training?

MS. HAIDEE SCHWARTZ - OPENAI: I'm not saying it's not of

interest, I'm saying that if you look at the, Andrea can speak of the recent deals more, but those are focused really on the user experience and them being able to interact with the information on those sites, in a way, in terms of display and what they're learning from them and getting the news there and having it highlighted and going to that site. Same with [indistinct 01:44:45]

MR. ANDREA APPELLA - OPENAI: And I think, just to follow up on that, as we know many of the concerns that news publishers have expressed over the years were about links to their own websites and the legitimate sources and therefore, I think that is something that we plan doing in the context of these deals. Then of course there is access also to data bases and information, it's a mixture of the two.

MS. HAIDEE SCHWARTZ - OPENAI: Yeah, but it does highlight that these deals are about, that they will include links and access to the data on these, links to these sites.

CHAIRPERSON: Yeah, look I think, I mean as I said, I think the pre-training is certainly one aspect as you say, Haidee, but I'm also interested in the other aspect because Andrea, you say GPT is not about current but, and it's not necessarily about current affairs, but these deals are almost saying the opposite because they're saying you want to bring news media in Spanish, in French, in German and obviously in English, easier to your ChatGPT users. So it seems you're building a capability around current affairs and news specifically for ChatGPT. That is the sort of, what seems to come from these deals at least what one can infer. I don't know if you want to

comment.

MR. ANDREA APPELLA - OPENAI: Well, the deals are not being implemented yet and there's lots of research and discussions, so when I heard that's being made available to consumers I'm sure it will be publicly explained. But at the moment, there's a cut-off date for the data is available and no current news are available to the GPT models.

MS. HAIDEE SCHWARTZ - OPENAI: I do want to say, like this whole discussion has been focussed on news media but what we need to recognise is that ChatGPT is used for so many things that have nothing to do with the news media. I mean, people use it to code, people use it to plan trips, people use it to understand problems, people use it to help them translate a passage in something they've written. There's a huge amount of uses that have absolutely nothing to do with current events or news. Not to say, and I'm not diminishing the importance of news or current events, but just recognising that if you look at all the use cases, so many of them have nothing to do with that. Not to say that news and current information is not important, but recognise it's only one facet of what people use our generative AI or many other types of generative AI for.

CHAIRPERSON: And we've heard similar comments from search and social media that it seems news seems to have some prominence as a use case and just what, I mean your chief operating officer, Brad Lightcap said of Axel Springer, the deal will help the right people with new ways to access quality real time news content through our AI

tools. And he said a similar thing about Le Monde and PRISA Media where the goal is to enable ChatGPT users around the world to connect with the news in new ways that are interactive and insightful. So, it just seems that you have recognised that building capabilities around news may be something your users want or desire and bringing that real time aspect is something you wish to bring about?

MS. HAIDEE SCHWARTZ - OPENAI: We have to recognise that we've operated for years with cut off dates that meant that we really didn't have access to real time news or real time events. And that is an aspect of the world and people are using our technology to understand like what the doctor said to me, what does that really mean, to help me understand what just happened in Russia. And so there's a whole variety of things and it's a full picture and so we're sort of, I think, all I was trying to say was that there's so many more uses that we have to recognise. Like writing code or writing a poem or helping write an essay or hoping to write a toast for somebody and that is a lot of what the model does.

MR. ANDREA APPELLA - OPENAI: And also OpenAI does not track user behaviour where it specifically measures news consumptions. I think that is another important point that we wanted to underline as Haidee said, all these other uses but we don't specifically target to track to measure news consumption.

MS. HAIDEE SCHWARTZ - OPENAI: Right, and very important to reiterate, since we don't sell data and we don't sell our data, user data, and we don't sell advertising, we have a very much different business

model and we're also obviously controlled by a non-profit and have that mission. But it means our products are out there with people being able to make other choices and certain people that maybe be able to offer different things than we're able to offer because we don't make those choices and we don't have revenue from that.

MS. PAULA FREY: Haidee, I agree with you that this is a very, very different model but the number of media organisations came before the panel and expressed concern that the search capabilities of generative AI was such that you would be able to actually ask what was happening in terms of the news and get a summary of what was happening. And yes certainly in ChatGPT you can get references back to the news stories but people unlikely to click on those quotation marks and so there's real concern about the potential impact on the already losing search click throughs, that they might lose even further with AI being able to just summarise news for paying clients.

MS. HAIDEE SCHWARTZ - OPENAI: I think, as you've seen, we try very hard to, whether it's facts versus news is the question. Obviously news reporting is copy, current news reporting is particular things. We try to make our models not regurgitate current copyrighted material. When it does it's a bug, we work very hard to make that not happen where we do not seek to replicate copyrighted material. That said, it's a continuing effort to make sure that we do not regurgitate things that it should not be doing. I think we've been encouraged about some of the progress that we've made in that area and we'll continue to work on that.

MS. PAULA FREY: If I could give you an actual example, I actually did a ChatGPT prompt asking about a current news story. Somebody in a current news story and got a summarised version of the full story, and yes, there were two quote marks that actually linked to media stories. Until I actually thought just to double check the sourcing, I had no intention of clicking through to the story and I can certainly see what the concern amongst media would be in terms of the impact of AI search.

MS. NORA PUCKETT - OPENAI: I think this is part of our efforts to engage on an ongoing dialogue with creators, with news publishers and to reach the types of deals as Haidee and Andrea have talked about that speak not only to just the engagement with the data, but also interactive and other features that will be neutrally beneficial. All of the understanding that we are doing this because we think it's the right thing to do and keeping in mind that we do believe that we are abiding by the law but nevertheless, I think from a society as a whole perspective, engaging with publishers and making sure that we provide these, a rich environment for rights holders and news publishers works to succeed is really important and core to one of our values.

MR. ANDREA APPELLA - OPENAI: And also, I think certainly, I hope you will appreciate and agree that it's better to have attributions and links than not having them. So in many cases it's a result of the discussion with the publishers, it's what they really want as well. Also in light of the problems that they had in the past. So then the question

is yes, you give the attribution and the links and then consumers hopefully will click on a click, but it's better to have them than rather not have them right.

CHAIRPERSON: Ja, I think, I mean you're right, Andrea, that it's better to have them than not but it's also better that users click and people can generate revenue from the stories that the journalists write about. So I mean don't know if OpenAI has looked at the click through rates from current affairs stories and what those look like?

MR. ANDREA APPELLA - OPENAI: I wouldn't be able to answer that question.

CHAIRPERSON: Ja, I think I mean it's just whether OpenAI is doing that or not. I mean we'll ask you to do it any way as part of this inquiry, but I don't know if you have access to that information.

MR. ANDREA APPELLA - OPENAI: I doubt it, but we'll have to double check.

CHAIRPERSON: It is quite, look it's quite funny but neither Microsoft nor Google nor yourself wishes to know this information it seems because no one wants to know what the click through rates are on ads and yet we've been told that you mustn't worry because attribution is there and people can click through. So I think it's something we'll pursue with all three of you. I just had a couple more questions.

MS. HAIDEE SCHWARTZ - OPENAI: I would say that be we don't monetise, we don't sell ads, we will not have the same infrastructure that somebody who sells ads would.

CHAIRPERSON: No, that may well be the case but I suppose there's

two things Haidee, I mean the one is that your products are being used by companies that do sell ads, that's search engines and others, but importantly this is an issue of concern from the media and if you're engaging the media, I think you would have an interest in just understanding how those summaries, as Paula put it, are impacting on the media, because that's important information in the dialogue that would happen I think with the media. Certainly one that we think that is important. It's about the transparency and knowledge and you have the knowledge no one else does so I think it just helps that we get that information. Just on, I mean Andrea, you mentioned that you've stopped this request of users to do something in the style of a leading artist. Is that just with music or is that across the board, pictures and other content?

MR. ANDREA APPELLA - OPENAI: I think it's in the context of DALL.E which is the one that generates images and the model is trained in a way that they cannot replicate or generate anything the style of living artists.

CHAIRPERSON: Ja, because I think this is some of the big disputes that artists have with AI is around effectively in the sense saying then this is my style, the users asked for that and it's obviously been trained on my style, my stuff in order to give an answer. So you haven't encountered this beyond DALL.E, the image creator?

MR. ANDREA APPELLA - OPENAI: Well, [indistinct 01:58:07] is not the product that is available yet, that is the only other one that can come to mind at the moment. So I think it is expressly referring to the

product that is able to generate images. And that definitely the model is trained not to generate anything in the style of living artists. But I think it's consistent with that approach to protect creators to keep engaging and respecting the community and having dialogues so we don't want to go out there in the direction that would be upsetting the living artists or creators.

MS. PAULA FREY: One of the things we did ahead of today's session is just to play around with ChatGPT and asking it to write a story in the style of a particular news entity, write a story, a hard news story in the style of a local daily publication, write a story in the style of a weekly publication and certainly I mean ChatGPT produced things that looked similar. I mean it's not scientific we would have to go back and just compare it across several stories, but I certainly did ask it to write an introduction in the style of Paula Frey and it was a little bit awkward in the similarities. I asked it to write a story in the style of a well-known woman journalist and certainly reading it, it had a sense of familiarity. And so I'm just wondering if that was something you were trying to encourage or discourage?

MR. ANDREA APPELLA - OPENAI: I don't think it's a matter of encouraging or discouraging, I think it's a more general question about how the knowledge of the model is used. It's like human knowledge, and I was actually I was talking about especially this topic with, I have lots of friends in the book publishing industries and authors and even before I joined OpenAI we are chatting about this and they said, well even an author sometimes has read many novels

and he's inspired by the style of Virginia Wolfe and then they want to write something perhaps in that particular style, so I think it's a broader discussion about what that means in practice. And it's part of the knowledge and creativity like we're all exposed to different kinds of creative inputs but then we elaborate in certain ways. So I think it's a discussion that we can have, not just in the context of [indistinct 02:00:52] learning or generating output from an AI language model, but I think it's more a general question about the creativity in general.

CHAIRPERSON: I suppose where it sort of comes up is what we have seen at least in South Africa are deep fake videos certainly, but also going into an election where one of the multitude having an election this year, but the copying of journalists for disinformation and misinformation is a feature. And I know you commented on your efforts to prevent that, but certainly copying the style not just in a video context but also in a written, there's going to be something that may contribute to that and there's a lot floating around on the web already in South Africa on this. So I think it just comes up in that context if you have that capability, you're able to create with your tools that sort of potential disinformation.

MS. HAIDEE SCHWARTZ - OPENAI: I think just going back to the disinformation including election misinformation, again I'd reiterate all of the many things that we are doing and continue to do and meet with stakeholders to prevent that via deep fake so other misinformation are just information content. We also are hopeful that these partnerships that we do have with news media will allow us to ensure that authentic

news content is something that is available to our users because certainly that will help counteract the potential for disinformation if it somehow inadvertently gets surfaced or if bad actors try to create it through our products despite the many efforts we take to prevent that from happening in the first place.

MS. NORA PUCKETT - OPENAI: And I know we're well over the time that we'd originally planned, had spoken and agreed with the commission, so I don't know if we are close to wrapping up, because I had other – I'm sure you have people being testifying in.

CHAIRPERSON: Yeah, no we are and we appreciate the extra time you have given us. I just had one more comment just to Nora, just to say that our Copy Right Bill does not have a text and data mining provision and our fair use certainly doesn't single out AI, it may be that it's not up to date. But it does have a preclusion for use that may substitute for the work in questions. So I think the concern from certainly that I hear from the media and what Paula articulated is where summaries are created from their content that may substitute for going to those websites direct. But it's something we'll pick up, I just wanted to at least in the public hearings indicate that we don't necessarily agree with that position, that OpenAI has taken. But that is more of a legal question.

MS. HAIDEE SCHWARTZ - OPENAI: If I can make just one comment on that and I appreciate that and similarly I think fair use doctrines and other jurisdictions likewise have not been updated to speak directly to AI. Nevertheless we think that they have persisted

throughout different types of innovation and the goal of our large language model is not to substitute for the work, right. That would be a bug and we do not want to reproduce or regurgitate any kind of content that happens to be in the training set. The contrary as we said before, we are training on trillion and billions of different types of data in an effort to predict a new response, a creative response and so we certainly do not think that it is a substitute and whether you're talking about a text and data [indistinct 02:05:03] inception or fair use concepts and various jurisdictions across the globe, we think that they take into account the kind of transformative nature of what large language models, be it ours or other frontier large language models that other companies are offering are within that. But as I've said, and I think all three of us agree throughout we nevertheless are committed to our work with rights holders to try to get this right from a kind of global society approach.

CHAIRPERSON: Yeah, and thank you and oh sorry, Andrea go ahead.

MR. ANDREA APPELLA - OPENAI: Just before you close off because I just heard that today our team has spoken to a South African media company, therefore I think we can end on a positive note that we wanted to let the Commission know. So your hearings are already having a positive impact.

CHAIRPERSON: No, well thank you Andrea, and ja, I think obviously having public hearings these issues are not isolated to some of the global north alone and I'm glad you are starting to have interactions

with South African media. Obviously we have asked you to interact with this inquiry and to engage the inquiry, and in that process you may well engage with the local media, but I think it's probably to the benefit of also OpenAI that it does broaden its net in terms of who it engages and who it hears from in looking for certain solutions. But I do want to commend the OpenAI team for coming here and engaging the inquiry in South Africa, there are a couple of platforms that are not doing that which we think is disgraceful. But I think for OpenAI we commend you. And for also such a high powered team coming to talk to the inquiry, and to educate the public and the media more around OpenAI and how it is approaching things. So thank you, Haidee, Nora and Andrea. And thank you again for also extending the time, we do really appreciate it.

END OF PROCEEDINGS ON 25 MARCH 2024