

ISOJ 2022: Day 1

Google Tools for Journalists: Lunch workshops offered by Google News Initiative

[Marco Túlio Pires](#), News Lab lead, Brazil, Google

Marco Túlio Pires Welcome, everyone. Welcome. So maybe we should start because we have another session starting right after this. So just get your seats, feel comfortable. This is the session in English about Google tools for journalists. So welcome. I also welcome folks who are joining us virtually today. For those of you who are joining us virtually, I am Marco. I am wearing a blue shirt and red trousers. I also have a beard. I am Brazilian, so I'm based in Brazil, in Sao Paulo. And I often come to these conferences, and I work at Google. I work at the Google News Lab, and I am the Google News Lab lead in Brazil. So if there's anything that you don't understand because English is my second language, just please raise your hand, and I'll try to say it in another way.

I'm going to give you a gentle introduction to some tools that we can use for journalism, some tools that we have at Google that can be used for journalism. And then if we have the time, and hopefully we will, I will have a hands-on session with you guys about creating your first machine learning model using an amazing tool from Google called Teachable Machine. So we're going to talk a little bit about artificial intelligence for journalists as well. Okay? All right. I'm getting lots of head nods, so I'm going to carry on. And if you have any questions, feel free to ask in the chat as well for those who are joining us virtually. So let's get started.

I work at the News Lab, and we have this initiative at Google called the Google News Initiative, which is a major company effort to help journalism thrive in the digital age. If you're interested, g.co/newsinitiative to hear more about it. We also have a training center that we've built especially for journalists. So everything that we're going to talk about here today, you don't have to memorize. You can eat. You can relax. You can just watch the presentation. It's fine. Everything is in here at g.co/newstraining. This is a special website that we created for journalists with downloadable lessons, with modules, PDFs. Absolutely everything that I'm going to talk about here today, you can see it on the training center. I'm going to give you another opportunity to look at that link toward the end of the presentation, and if you have any questions, just to free to raise your hand and ask away. All right.

So gentle introduction. I'm going to talk about a few tools, and then we're going to talk about machine learning. Hopefully you will learn a few tricks today. Most certainly some of the things I'm going to talk about here today, you probably already know. But then we're going to emphasize, and if you have any doubts, we can have a conversation.

So let's start with search, the most basic Google tool that we have. Search started in '98 and search used to be really like a blank page with ten blue links. You guys remember that? Were you around when Google was just, like, starting? And we even had an exclamation point on our logo and a beta word. We were way more excited about the name of the company back then. Search has actually changed a lot. It changed a lot

because it went mobile. It changed a lot because people started using their voices, but also it changed because of richer results. Right? Not only ten blue links, but because of structured data on websites, we're now able to build panels like this with news carousels or knowledge panels where the search experience is way more rich than before. And there is a way to have this information on your website so that the search engines, not only Google, but Bing, Yandex, Yahoo! Etc., they can also build panels like this.

There is an interesting sentence from Larry Page, one of Google's founders, when he was idealizing the best search engine. He used to say that the perfect search engine should understand exactly what you mean and give you back exactly what you want. And this is a very hard problem. Because if you're searching for Apple, it might be Apple, the company that produces iPhone, but it might be Apple the fruit, or it might be even Apple Records, that recorded the Beatles' discography. So understanding what you want is a really tough challenge. But there are ways for us to indicate to Google what we want and help refine our search experience.

So for the next few minutes, I'm going to give you some search refinement operators that I'm sure some you will know, but I'm sure some you won't as well. Those operators can be accessed when you go to Google.com. And then on the right, lower corner, there is a settings option. You click on it, and there's advanced search. And then you find a form, and all of these operators are there. But let's start with, for example, if you're searching for the speed of a Jaguar. A jaguar what? Google will try to understand that. Maybe it's jaguar the cat, and we'll give you the speed of that. But maybe you don't mean the cat. You mean the car. So there is a way for us to exclude words from search requests, and you can just use the minus operator. So you type in "jaguar speed" minus "cat," and the minus operator will exclude words from your search, and voila. It will give you the speed of a Jaguar car if you exclude the word "cat." So that's the minus operator.

We can also use Google as a search tool to search on a specific website. You know, when websites have search engines, but their search engines normally aren't that great? You can actually use Google's algorithms to search in specific websites. And the way you do that, you just use the site, then colon, and then the address of the site, and then the search parameters. So here's an example. You can search for antibody tests, but just on the site of the World Health Organization. See that when we use this, all of the links searched here will be from that specific website. So you can actually use Google as a search engine for any website. Any website, even Instagram, for example. You can search for the pepper warrior hedgehog in Instagram and get listed on your desktop, in your search experience, and find videos of this amazing, amazing warrior. We can also use the site parameter to actually not only look into websites, but domains. Think about government domains, or country specific domains, or educational domains. You can actually search within that range of websites. If you don't use the full URL or the full domain, you just use the domain extension of a website. So in this case, we're searching for Boris Johnson and then site:.co.uk. So we'll only search for Boris Johnson on websites that have a .co.uk domain, any website. And here is what it looks like. So you can use like a slice of the domain as well. Another one is on Twitter. You can search for Twitter lists concatenated or joining some of the parameters together. So here we have site:twitter.com/ and then we have a wild, what we call, a wild card, or an asterisk. If you put an asterisk in search, search will try to replace that with anything. So that's the indication of please put here anything, so it can be twitter.com/anyusername/lists. So what you're asking Google is to say, give me the lists of any user on Twitter that also has the words "alt right." So you're effectively searching for lists that users built with the keywords "alt right." So that's a nice trick.

You can also expand your sources. You can use a neat parameter called `related`. So imagine you're doing research about local journalism, or you're trying to find sites about cinema, or anything, or health, or transportation. You already have a group of websites that you want to use, but you want to find more websites. So you're going to ask Google. Google, I know this website. Can you show me other websites that look like this? So the parameter `related` does that. And you can see it here in action. Another one, which is really good as well, is, you know, when you're trying to find a website, and the website either changed or got deleted, like a blog or a post? Like I swear it was here! Or the website is down. Google caches the websites. So you can use the parameter `cache` to look for the last cached version that was visible to Google, and you can track whether or not the website was changed since the last time that Google saw it. And there are certain use cases where this might be useful. So, for example, we can use `cache` and then a specific URL to track if there were changes to that website. And here we can see, for example, on the CDC website, we see that the cached version of Google showed like the February 29 version, and the live version of the website is March 3. So you can compare what changes were actually made.

That's an excellent question because it's the perfect segue into this slide. So if you want to go through a timeline of like versions of the website and say, no, I want to see what was published last week or last month, then use `archive.org`, which is an amazing project not maintained by Google, but is an amazing project that archives the whole web. I mean I have yet to find a website that I can't find on `archive.org`. Even websites that go as far as 2002, 2003. And you can find versions all around the world, and you see a nice visualization of how many captures `archive.org` has made of the websites here. So use `archive.org` if you want to see versions of websites, and they go, you know, it's incredible.

You can now also use Google to search not only for web pages, but for files. Say you're looking for a PDF file, or you're looking for a spreadsheet, or a PowerPoint presentation. You can instead of looking for web pages, you can search for files, and the operator is `filetype`. So if you type, `filetype:` and then the extension of the file that you're looking for, be it MP3, WAV, PPT, etc. and then the keywords, then you will find files that have that keyword. So this is a fun example, and I'm going to add another operator here. So imagine that we want to find PDF files on the CIA website that have the word "secret" on their title. So you can search for `intitle:secret`, which means we're looking only for the title of the document, or the title of the page. `filetype:pdf`. And then `site:CIA.gov`. So translating to English is please give me PDF files that have the word "secret" in their titles on the CIA website, and then Google will list at least three files on the CIA website, PDFs with "secret" on their titles. You never know what you're going to find. Many years ago, I've heard of an example of a journalist who was trying to find what files on a website, and he actually got files that were not supposed to be indexed because he was trying to find spreadsheets and things like that. So you never know what you're going to find. You never know what people put up on the web. So for investigative purposes, it's a nice trick. You can also find CSV files or spreadsheet files. So in this case, expenses. If you just type "expenses" on Google, you'll get like a definition of expenses. This is like the normal search experience. But when we add the file type operator, this is what we find. Right. So we have files that are going to be pointed to you straight to download. And in this case, those types are CSV. For those of you who don't know what a CSV is, it's a comma separated value file. It's just like a spreadsheet, but it's in an open format. Okay.

Everything that I said is in the slides. So `-` to exclude words. `Site:` to point Google to a specific website or a specific domain. Right? In a specific country. You can also search for

related websites. You can use the cached version of the website. Or you can also search for datasets or spreadsheets. And remember, everything is in our training center as well. So don't worry about having to digest all of that information, because I know it's a lot.

We've been talking about searching by text, but we can search by images. So if you have an image on your hands and you want to search, use that image as your input to search. You can do that. And why would you do that? In journalism, we have to do that because we verify images, right? So if you go to images.google.com, you can paste the URL of an image, or you can send an image from your computer or your cell phone. But if you load Google images on your cell phone, you have to request a desktop website, so that you can see this camera here. So, for example, there was a tweet, I guess like four years ago, or three years ago, saying that this is a protester in Hong Kong returning tear gas with a tennis racket on Twitter. And if you're using Chrome, you could just right click on that image and search Google with this image. And then you would find out that this is actually an image from 2016, so not from 2019, which is when this tweet happened. So the image search, which we called the reverse search image, is a good trick for verification. If you get those images you see like visually similar images, you see where the image was posted first because Google indexes those images and can identify similarities and whether or not the image was photoshopped was doctored, etc. So it's a very powerful tool for verification purposes.

We also have a dataset search. If you go to g.co/datasetsearch, you will find a search tool, but not for websites, but actually for data sets. So if you're a data journalist or if you're using data to strengthen your research, you can find data that is published here by international organizations like the World Bank, like universities, like NGOs in the United States, Kaggle, which is a massive repository of data sets. And you can see the dataset when it was published, the kind of license that it has, the last time it was updated. So it's a very, very smooth experience to search only for data sets, if that's what you're looking for.

We also have finance, so google.com/finance where you can search for finance data and lists of the companies, and you know, all information about indices and listed companies. And you can explore that.

And we also have Google Scholar. Google Scholar, especially here in the US, is amazing. We use it in Brazil and all over the world, but especially in the US it's amazing because you can look for academic papers and follow trends in the academic world. Google Scholar is almost like the de facto search engine for academic papers, but also U.S. court papers. You can look for court decisions, and papers, and documentations, and you can create alerts. You can filter by states. You can filter by counties, and you can also create alerts. So you don't have to come back to Google Scholar every time. You can create alerts by topics that you're following, and you want to see if that decision was made, or a certain paper was published about, I don't know, black holes, or COVID, or food safety. You name it. You can create alerts, and you will receive recurrent alerts on your email.

Speaking of alerts, we have a special tool to create alerts, google.com/alerts. You can create an alert for your own name. A lot of people do that. When you Google your name, if there is a new web page that is talking about you, you can create an alert to be alerted when that web page is indexed by Google. And the interesting thing about the alert tool is that all of the operators that I showed you on the first segment of this presentation, they all work here. So you can actually use that specific search to create an alert just for you. In this case, I'm using the intitle "secret," and the government websites in the UK, and the file type pdf. So I'm going to receive an alert every time a PDF file is indexed by Google in UK

government websites, including the word "secret." So I don't have to go to Google to do that every time. So you can create little robots that are going to work for you and help you with your research. Or you can also look for vote by mail. You can look for court orders or papers regarding vote by mail and create alerts when those decisions are made. And you can filter, like in the region of the United States, how many delivered to, and the sources, how often you want to get those updates. So it's pretty neat.

Okay. This is search, folks. Search. You don't have to remember everything that I said. Like I told you guys, it's all in the training center. Any questions so far? No. The training center is [g.co/newstraining](https://www.google.com/newstraining). Now I'll just put it here for a second. [G.co/newstraining](https://www.google.com/newstraining). That's the training center, and it has it in all kinds of languages too. I think it's 20 languages, so you should find a language that you are comfortable with. All right. So that's search. There's way more things to talk about search, but let's shift gears a bit, because I wanted to give a few examples about Google Earth.

We have a new version of Google Earth that came out about two years ago. And it's beautiful. It's slick. It's much better than the older one. And I wanted to give you an opportunity to take a quick look at the things that you can do as a journalist or as a visual storyteller, because the tools that were built into Google Earth now are really special and will help you tell stories with maps in a beautiful way. Usually in the past, we used to have to hard code those stories on Google Earth, but now there is an authoring tool where you can create a new project and then mark points on the map and locations, and you can hit play, and it will navigate for you. And you will save it on your Google Drive account. I won't have time to show the authoring tool now, but believe me, it's truly, truly amazing. Because I also want to show you time lapses in Google Earth.

Time lapses used to work on an old tool, but now we embedded it on the new Google Earth. And time lapses are great because they gather images from the past decades, I think 30 years. And there's lots of providers, NASA and other satellite providers, with images from all around the world. A lot of regions. And you can use the time lapse tool to create kind of a visual journey to see how the landscape of the planet has changed on a specific location. And this is very special to me as well, because Brazil is the country that has the largest portion of the Amazon forest. And you can actually see the impact of deforestation in the Amazon forest. It's very powerful. I mean, you can write 20,000 words, but just the time lapse, like a few seconds, will tell you way more in this experience. So let me try to show you the time less real quickly here. So you go to [g.co/timelapse](https://www.google.com/timelapse), and then you see here that there are a few examples here on the right side. But you can create your own. And let me click on this one, changing forests, which I think it's soybean plantations during the Bolivian portion of the Amazon forest. But you'll see. So, here is the timeline, so 2000s, and then year by year. And then it has all of the images until 2020. And then it goes back to the beginning again. And you can clearly see the evolution of the deforestation there with the soybean plantation. And you can do this everywhere. You can do this like for construction sites, or for gentrification of neighborhoods, or changes in business locations, etc.. So it's a really powerful storytelling tool, Google Timelapse, [g.co/timelapse](https://www.google.com/timelapse). So this is how you can use timelapse in Google Earth. And again, it's on our training center. You don't need to memorize. Just go there and take a few lessons, and you can make timelapse videos. You can export those timelapse videos and embed on your website to enrich the story that you're publishing.

We also have a tool called Google Earth Studio. Google Earth Studio is an authoring and animation tool. Have you ever seen like those visuals on CNN or other news organizations where they're like, "Oh, the protest happened here, and then they follow the street"? And

then it's a map, like a 3D map, and it's very real. This is Earth Studio. They're actually using our studio, and you can do that yourself, too. It's not that complicated, especially if you're used to work with video editing tools, with key frames, and things like that, and it's free for journalists. So just go there g.co/EarthStudio, or just type "Google Earth Studio" on Google, and then you'll find it. Sadly, I won't have time to show you. I'm going to be here until Sunday evening. So if you want to get to a preview of our studio, just find me in the halls, and I'll show you on my computer.

The other tool that I wanted to tell you about is Pinpoint. Pinpoint is an amazing investigative tool that was launched by Google. I think it was the first comprehensive tool that Google launched especially for journalists. So this is now Google Maps that a lot of people use, including journalists. This is like a tool that we made for journalists. And Pinpoint is a tool where you can research, explore and analyze documents and collections of documents that you have in an impressive way. So what we do with Pinpoint, we get all of the documents on your research. The file types that are compatible are PDF files, word documents, slideshows, emails or email archives, images, plaintext, audio. We recognize the characters on the images. We recognize the speech in the audio. And we get all of that, and we create a collection. And in that collection we automatically identify using Google's knowledge graph, we identify people. So the names of people in all of those documents and images and audio files. All of the businesses, names of companies. And all of the locations. So if there is a location, say, a specific address that is in a court order, or a document, or a police report, we automatically identify those. And then if this is mentioned in an audio, someone said it, and you click on that location, we will list that, "Oh, it's mentioned here in this document, but also here in this interview." So it's a place for you to really go deep into the entities, places and locations that are mentioned in thousands of documents during your investigation. And you can ask people to collaborate as well. You can share those collections with your folks, and it's free for everyone. So it is free, but there is a caveat. So all of the document policies that apply to Google Drive, they also apply to Pinpoint. So if you have sensitive information, if you have documents that should not be accessed by anyone, not even if there is a court order, then you should not use Pinpoint. But if you're using like open documents, and like open data, or documents that don't apply to these cases, by all means use Pinpoint because it is very powerful.

And we actually have, Jeremy. Do you want to?

Jeremy Gilbert Oh, sure. I guess just very briefly. One, while working at The Washington Post we worked very closely with Google on Pinpoint. This tool is really amazing. The mistake we made initially was thinking that this tool was a tool for our data journalists. It's really a tool to make anyone who is not a data journalist, a data journalist. I can't more strongly endorse it. The thing that we did to me exemplified how the system works so incredibly is that we took all of the Mueller Report documents, all of them, court filings, public reports, everything, and then we dropped it in Pinpoint. And then during the Mueller hearings, our reporters were not used to working with our documents that don't make their own experience adjusted pretax first on those documents. And we were able to have at hand all kinds of incredible information in our reporting that we wouldn't have been able to find on deadline otherwise. There's a ton of great information in there. It'd be pretty easy to start adding and reducing that common. And your reporters could find different datasets that other reporters shared. So that's all great, too. The other thing I want to tell you about. For those of you in the US and Canada, thanks to some incredible support from the GNI, we have money.

Marco Túlio Pires What is GNI, Jeremy?

Jeremy Gilbert The Google News Initiative, sorry. Thanks to the Google News Initiative, Northwestern is running a program to give grants between about \$10,000 to \$100,000 to support projects that include reporting on documents or data. There are very few restrictions. The Pulitzer Center does incredible work. We are more flexible in our program than that program, so you can get money to help acquire technology, to get training, to hire someone to backfill for the job while you're reporting. So the reason I really want to talk to you about this is because the deadline for the first round of this is next Friday, so come see me if you're interested.

Marco Túlio Pires Thanks, Jeremy. All right. Like I said, everything is in g.co/newstraining, so you can go there and look for it. But I wanted to shift gears a little bit and talk about artificial intelligence, because there are ways for journalists to use artificial intelligence on their reporting. And I wanted to give you maybe some examples, but also I wanted to give you a little taste. So if you have your computer with you, just take your computer because we're going to do some exercises together. Hopefully, we'll have time to do it. This is a sample of the introduction to machine learning lesson that we have on our training center, and it was done in partnership with the London School of Economics that run the journalism AI program. But there are lots of things there that were made for journalists to understand machine learning. So if you have your computer, let's do a couple of exercises together to understand what is this machine learning thing about, and why should I care? Right.

So the first two that I wanted to show you is a fun tool called Quick, Draw! Have you ever seen Quick, Draw! Raise your hand if you've seen Quick, Draw! All right. He's seen it. Don't tell anyone about it. All right. G.co/quickdraw. So let's go to Quick, Draw! real quick. Quick, Draw! is a game, and it is a game in which you are going to get a prompt to draw something. Like draw grapes or a car. The computer will ask you to draw something and you're going to start drawing, and the computer will try to guess whether or not it's correct. So let's start. So I need to draw a wine glass in under 20 seconds. So a wine glass. I mean, I'm using the trackpad. Okay. Give me a break. Okay. It got it. It was a mouth. Okay. Now a birthday cake. So this is easy, I think. Oh. All right. Lipstick. No. I mean, what else can I do? All right. Next. Peas. Just peas. Animal migration? All right. Camouflage? No, peas. How do you how. I don't know. Next. All right. And last one, eraser. I don't know. I don't know. All right. Eraser. It was almost like that. All right.

So what happened here? This was an experiment by Google. We asked folks from around the world to draw things. Right. So let me go back here a little bit. We asked folks from around the world to draw things, and then people from different cultures, different countries, we asked them, okay, so draw an arm. And then, people drew arms. Right. And then we got all of that data that we knew were arms, and then we trained a model to guess what an arm would look like when someone was drawing on the Internet. So that's basically what this experiment is about. Getting lots of examples, lots of images, that we know that they are something. Then we put it in the middle of a model, a machine learning model. And then when we're bringing a new image, the model tries to guess which images is that. And this is the backstage, so to speak. And you can see all of the images here. Like this image is from 2017 in Poland. I mean, it's a huge data sets of drawings, and it's feeding into this tool that tries to guess. Okay.

So this was actually, congratulations, this was actually a primer on machine learning. We didn't talk about any technical terms, but this is actually how machine learning works on this specific task. So let me go back to the presentation here. Okay.

So let me very quickly talk about some of the concepts of machine learning, and then I'll show you, because I have six minutes, I'll show you the last two. And then we'll create a model together. So like I said, we had samples like the images that people draw, and then we put those samples in a model, which is statistics. So you put in an equation, and then they calculate it. And then you have guesses, right? New images come. They're applied to this model, and then it tries to guess based on the examples that you gave. So this is basically what machine learning does in this case.

So machine learning is a subset of artificial intelligence. We're not going to talk about everything artificial intelligence. Even machine learning is such a big field that we can't really talk about it in completeness. So is it automation? Is it transcription? We use it for transcription. We use it for translation. This is not machine learning. So when you see code and rules, this is not machine learning. Machine learning can be used for recommendations. So when you see songs on your Spotify account or hopefully your YouTube music account, you will see recommendations based on machine learning. You can use for forecasting, or game strategy, or even classification, which is what we just did, right. Classification in terms of like this is an arm, this is a ball. We're classifying images, and then we're creating models to guess. We can also create crazy images like this one. This was created by a computer.

Why is this so relevant? Raise your hand if you can find where the broken cookie is in this image. Okay. This is not broken. Try again. Another one. Yes. Okay. We found the broken cookie. It took us 30 seconds? 30 seconds to find one cookie in an image of 24 cookies. Right. What about if it was, like, a thousand cookies per second? How do you find the broken cookie? You can't. But we can use machine learning to identify the images and find the broken cookie. So here's an application why it's relevant in the industries to develop machine learning vision to identify, for example, increase production. So machine learning is a big field. We're talking specifically about supervised learning, which is you teach the model like with the good examples, but there are others too. What we do is actually find patterns in data. And how we use AI at Google? A lot of the things that we do at Google, we use machine learning. So for example, if you're using Google Translate, the live mode of Google Translate, you can actually translate signs, and menus, and text in the real world. And we use machine learning to identify those images and translate back into text. We also use machine learning when we're suggesting replies on text messaging apps or suggesting sentence completion on your email or Google Docs. We're using machine learning to do that too. And we also use machine learning in voice. When you use your Google Assistant, or your Alexa, or other smart speakers, you are using machine learning because you use this machine learning to recognize accents, voices, words, etc.. There are other industry applications. Many of you may have seen this video. Yeah? Or this one. So this is a worker, right? A robot worker that's changing boxes from one pallet to another, and it's using machine learning models to recognize its environment, and adapt, and be flexible to do this kind of work efficiently. So, yes, there are other applications in industries as well for machine learning. So we use in images for recognition of care. So Pinpoint uses machine learning to recognize characters in images, for example. Natural language to sentiment analysis. We can look for if people are angry, or sad, or happy on the way that they write, for example. We can do that on tweets. And we can also do object tracking and go change label detection voice transcription for videos.

So in terms of classification, we label things. So for example, we were talking about the Quick, Draw! example. The Quick, Draw! example was like we got drawings from all over the world, and if it was an arm, we would label that drawing, "this is an arm." Or in this

case this is a dog, and this is a cat. So we're labeling things. In traditional programming, we have rules, right? We write, if an animal has a long tail and the animal likes fish and he hates people, then it's a cat. Right? So this is traditional programming rules. In machine learning, we learned through examples, not rules. So we build a model with those examples. So we have cat examples, and we have dog examples, known examples, that we tell the machine that exist. And then we put those examples in a machine learning algorithm, and then we have a model. Then we bring a new image that the model doesn't know, and then the model tries to guess with a certain level of confidence whether or not this is a dog or not.

Okay, so let's try. I know I have to wrap up, but this is the last tool I really wanted to show you because it's a very, very interesting tool called Teachable Machine. So if you have your computer, go to [TeachableMachine.withGoogle.com](https://teachablemachine.withgoogle.com), and let's do a machine learning model very quickly. And you can try this at home, and we have lessons on our training center to do. So you go to Teachable Machine, and Teachable Machine is a way for anyone to create a machine learning model using your webcam or your microphone. And if you know what you're doing, you can even use other machine learning models and put it on the tool.

So let's get started. So I'm going to create an image project, which is I'm going to show an image to the webcam, and then the webcam will tell me what that image is with a certain level of confidence. But you can also do audio projects. You can like clap, or whistle, or how do you call this, snap. Or you can do poses, as well. You can do this in front of your camera, and it will recognize what you do. So let's do an image project, a standard image project, and the first thing that we need to do is to gather samples. Remember when we said like, we have dog examples and cat examples, so we're going to build a model that recognizes a smiling or a serious face. So if we're smiling, the model has to say, this is a smile. And if we're serious, it has to say serious. But to start, we need to tell the model what a smiling face is. So we're going to click here and say, "smile." All right. So we're going to add images that are labeled "smile," and we're going to use our webcam to do that. I'm going to allow to use the webcam. Hopefully it will. Yeah. Okay. So I'll, I'll go there to do it. So can you help me here really quickly, Carmilla. Yeah, just hold. I'm going to go in front of the camera. Just hold this button here. Hold to record. Just hold it. Is it here? This camera is here? Hold it. Are you holding? Yeah? All right. Stop. Okay. All right. We got 110 samples of me smiling. In this case, it's better if you get lots of samples, lots of people, or you change the angle of your head. So that whenever someone comes in to smile, you get a full breadth of examples. All right, so the next one will be serious. All right. I need your help again, Carmilla. All right, so I'm just going to do a serious face in front of the camera. All right. One, two, three, go. All right. Good. Thanks. But we can add more examples, too, if you want. I'm just going to do the two ones. We're going to click "train model," and then it's going to prepare the data, consolidate all of the images. It doesn't take too long. We don't switch tabs or anything because everything is happening locally on the computer, and then it's going to build this model. And once we have the model, we can preview the model to test it. So right now it's saying that it's 100% confident that there is a serious face in front of the camera. But let's see if it changes. Is it saying smile because I can't see it? Yeah? Yeah, it's very confused. But you get the sense. So you can actually create. You can use your hands to point to other directions, up, down, etc.. You can create instructors for the computer to react to things that happen in the real world. So for example, you could create an automatic door opener for your dog whenever your cameras see your dog. Right. And then you can use this to create models, too. This also showcases the importance of being thoughtful about the training models, because this is only trained on my face. But it's not only my face. It has to be the faces of older people, younger

people, women, people of other cultures, etc. So it highlights the importance of when we're thinking about machine learning and using those techniques in journalism, and we're thoughtful about the approach as well.

I know it was a lot. Thank you so much for your company today. If you have any doubts, come back. Just search for my name on Google. I'm happy to talk to you guys on Twitter, or LinkedIn, or any other platforms. See you around at the conference. Thank you.