

Rachid Finge (00:01):

Hey, it's Rachid and welcome to the Made by Google Podcast. Hope you had a great week. And undoubtedly it's been a week where either you have taken a picture or someone took a picture of you, or maybe you were in a conference call and people could see you through a webcam. In all these cases, there's some hardware and software magic going on to capture what you look like. Now, it turns out not all cameras are great at capturing everyone just as nicely. If you have a darker skin tone, most cameras do a worse job representing you. That's a problem that needs sorting, and it's what we're talking about in today's podcast episode Within Google is a team that is leaving no stone on turn to make sure cameras capture everyone fairly and accurately. I wanted to know more about their work, so I'm glad to introduce today's guest who works on initiatives like Real Tone. Please welcome Florian Koenigsberger. Florian, thanks for joining.

Florian Koenigsberger (01:08):

Thanks so much. I'm really excited for this.

Rachid Finge (01:10):

So tell us what you work on at Google and how did you end up here in the first place?

Florian Koenigsberger (01:14):

Absolutely. There's so many amazing journey stories at Google. I guess I'm fortunate to count mine among them. I came to Google as an intern in the summer of 2013. I was a bold intern. Shout out my classmates that are still at Google. I just had a run in with one of them the other day at a policy event, nine years Strong. Amazing. And I entered the APM program, which is a rotational marketing program where I had the good fortune of beginning my career on the Google Play team, working on mobile gaming. I moved to Brazil. I worked in our Sao Paulo office for a year where the team tolerated me as I learned the world of ads marketing in South America. And I then returned to New York to work on a team that was focused on content partnerships and influencer marketing on our Chromebook and Google Play businesses before the journey that has led us to today the beginnings of real tone where we had this question almost five years ago about what it would take for Google to orient itself around making camera technologies that actively fought racial bias in the medium and delivered experiences that saw everyone, of course, hearkening back to our mission statement with dignity and equality in those tools. And so today, I'm fortunate to lead what we refer to as our image equity initiative. Exactly in that spirit.

Rachid Finge (02:39):

Well, first of all, there, I don't think there are many places like Google where you can work on so many things in just under a decade. And now then working on Real Tone where, you know, we have this internal page that I sometimes pull up and see what the mission statement is of Googlers and yours indeed is image equity. Apparently it's very fitting to real tone, but can you explain what that means? Image equity?

Florian Koenigsberger (03:03):

Absolutely. I love those little mission statements. It gives you such a window into how people think about their worlds, right? Yeah. You know, real tone, I would think about real tone, and I know we're going to talk about the technology as a subsidiary of this image equity mission statement, which is to deliver fair and accurate and authentic experiences that fight racial bias in our camera and imaging tools. And I say it

that way because the mandate is indeed broader than pixel, where we'll spend a lot of time talking today, right? Thinking about all of the surfaces that involve computational photography and computer vision that have a consumer impact on the way that people are seen. And the reason this matters so much to me is because, especially coming out of the pandemic, right? We experienced a really concentrated period as a planet where you basically existed to other people through a screen mm-hmm. Right? And so the ability for that screen and that camera to faithfully and accurately represent you, kind of was your representation in the world. Sure. And I think that even coming out of, I don't know when we'll ever say that we're fully done with this pandemic, but as we enter the next phase of sort of being in the world, I think that hybrid model has remained. And so people still really do mostly exist to each other through screens, and we still see that there are huge gaps between the ability of those tools to faithfully represent somebody in the way that they look in real life.

Rachid Finge (04:29):

So most people would think that the cause of that is maybe the camera is not of high enough quality, for example, there's not enough megapixels or the lighting is bad or something like that. But if you talk about image equity, it seems to imply that perhaps the camera itself has a sort of bias or racial bias. Is that the case? And how did that come about?

Florian Koenigsberger (04:52):

I appreciate you for this question because what we've seen in response to the project is a lot of people sort of raising their hands in horror at the notion that a camera itself could be racist. Right? But I think this is an important moment, as we've said many times across the company in all of our product inclusion and equity work, to remember that everything that we use is built, right? And it's built with an intention and it's built by somebody and someone has an inherent bias just from their lived experience. Sure. And so what we found is hearkening back to the 1950s and the development of film Emulsion Technologies by Kodak, when those technologies were first developed, they were really only tested and evaluated against people who were white or had lighter skin famously called Shirley Cards after presumably one of the first white women who sat for these testing images. Shirley, right? What we then found was as that technology advanced, the reason that some of those film technologies were developed more inclusively was actually not originally about people. It was about the fact that wood manufacturers and chocolate manufacturers could not get enough nuance in their product photography between let's say, a dark chocolate and a milk chocolate. Right? Right. I won't embarrass myself by trying to name different types of woods. I'm born and raised in New York City, and I can't claim to be particularly handy myself, but you take the point. And so when we accelerate into today's reality, right? Computational photography, smartphones, the same dynamic, of course, can exist. These are built tools by teams in an industry where we know that we are still very much striving to balance out racial and ethnic representation. This is a conversation that we had among our teams at the beginning of this process, right? If you are somebody that does not have darker skin, for example, and have not experienced the issues in the way that we can sometimes disappear into darker scenes or the way the camera can struggle with focus, then you might not self-select to work on some of those problems. And so this was an exercise about a, putting ourselves in the shoes as a team of a lot of our users who are living with this experience, not just on our products, but with every computer vision and camera tool out there, and then really asking critical questions about what it would take to close those gaps.

Rachid Finge (07:07):

Today's guest is a longtime Googler who now dedicates his time to improving cameras, but not in a way you think Florian Koenigsberger's work isn't about megapixels and frame rates, it's about making

cameras more equitable, which means to portray everyone accurately no matter where they're from or what they look like together with Googlers in the Pixel team. But as you'll hear also with many people from outside of Google, Florian is on a mission and he's just getting started. Please enjoy our conversation. So we already mentioned real tone, something that we introduced a while ago. Maybe not everyone knows what it is. So Florian, could you briefly explain what is real tone? And maybe the mistake I also made myself is thinking about it as a single feature, which it probably isn't.

Florian Koenigsberger (07:57):

Absolutely. So it's a question that we get often and I'm always happy to go under the hood on this because I think even for consumers really understanding what are the component parts of this technology helps people wrap their minds around how it is that the tools worked another way in the first place. So first to your question about what real tone is real tone is a family of technologies that we introduced first on the Pixel six in 2021 and that we've leveraged to develop some improvements to Google Photos as a product, especially in our editing tools, our filters and our auto enhanced feature designed to make a more fair and accurate and dignified experience for our users of color, especially those with darker skin tone taking into advantage some of those historical biases. Now, what does that actually look like? Right? Let's break this down, especially in the pixel camera. One thing to understand about computational photography, especially on Pixel, is that everything starts with the ability to understand that there is a face in the image, right? You have to know that there's a person there to then be able to make certain beneficial adjustments to rendering that person the way they ought to be seen. And so we made improvements to our face detection models, namely to be able to detect that there are faces, especially faces with darker skin tones, right? In complicated lighting settings, let's say a backlit scene or a low light scene, right? And we have examples where once you're able to detect there's a face in that image, say in an example where we might not have detected that face before this project, you can see what kinds of optimizations come into play, and that gets into the next set of technologies that are important here.

Florian Koenigsberger (09:34):

One core piece of this is auto white balance, right? Responsible for a lot of the color tuning in an image. So, does something look too cool? Does something look too warm? Does it look just right? Right. Not to make the porridge analogy here, and the auto white balance tuning that we did with the team was informed by, as was all of this technology, a deep partnership with external experts to Google, right? So I'll go back to this question that we asked at the beginning. How do we orient ourselves around this question of a more fair camera? The next natural question was, whom do we need to work with to do that? Right? We recognized that even within our own teams, we weren't at the level of representative diversity that we felt was required to faithfully do this project on our own. So we went out to what's become an extraordinary cohort of international photographers, cinematographers, directors, colorists, right? All who have deep expertise in beautifully and accurately rendering communities of color in their work, and who've been celebrated for that, right? And we gave them our tools and we said, go into the field and break these things, right? Tell us what's not working. And something that I always emphasize when I talk about this was obviously there's a technical learning that happens there, right? Understanding different nuances about exposure, values that might be more complimentary for darker skin, for example, not automatically assuming that because something is darker, it should be brightened to improve it, right? Right. There's a lot of rich tonality in melanin, and actually as a camera, sometimes you really need to respect that and leave it where it is versus trying to brighten it. But there was also a cultural learning that happened along the way, right? We had experts who had come in and talk to our

teams about ashiness, for example, the idea that especially in black cultures, right, there's an attention to dryness of skin and making sure that skin looks rich and moist, right?

Florian Koenigsberger (11:22):

And understanding that our camera had a tendency in certain scenes to wash people's skin out in a way that looked ashy, right? There was also a lot of cultural learning about why images mattered so much. You know, sometimes we think about this as a commodity and a sort of leading feature of a phone, but going back to our earlier discussion, images affect the way that you are seen in the world, right? They almost establish who you are before somebody sees you in real life. I remember I had a college professor who said to us, photography makes the world appear the way the world appears to be. Right? And I have to sit with that for a minute, right? I'll say it again. Photography makes the world appear the way the world appears to be. I think the essence of that is saying, we see the world through images mostly before we see it with our own eyes. And it is triply important that those images are done faithfully and with dignity, right? Even for landscapes or, you know, what have you, the myriad images that we come across every day. So those expert partnerships were a huge, huge, huge part of how we got to some of these technological developments. I know I already spoke to auto exposure Bbut if we look at that question of brightness was a big one, you know, there's a tendency to just bump things up because of course you want to see more information in an image. But what we actually learned talking to these experts was, no, we have to really respect those rich tones. And sometimes contrary to what a given, you know, camera evaluator or score externally in the industry might say, what the community is telling us is leave that tone where it is. Right? It's actually important that we see ourselves beautifully as we are and not feeling like the camera has to adjust something for us to so-called look good. So, auto white balance, auto exposure, I mentioned the washing out effect. There's another remarkable story on this journey. We had one of our engineers who has really moved, I think, by this mission and the consequences of this work to develop an original algorithm in our software that worked to reduce the impact of stray light coming into the lens after the fact of capture. Right? So stray light, if you think about the way that we use our phones, they're going into our pockets, we're putting fingerprints on the lenses, it's, it's by no means an ideal condition for most images that we make, unless somebody's really, you know, vigorously wiping down their lens, which, by the way, I think this is so cool, and I can say this because I didn't work on it. We do have a feature in Pixel that can remind you to wipe down your lens if it senses that there's something obstructing the lens, which I think is amazing because as a photographer outside of work. The number one thing that would improve almost everyone I know's phone pictures is if they just wiped their lens down before they took the picture.

Rachid Finge (14:05):

That will help a lot.

Florian Koenigsberger (14:05):

A lot. Yeah. It's remarkable how dramatically that affects image quality for most people's pictures. Right. So you have three examples there. I think, you know, there are more, we could go into them as them allows, but wanted to start with those.

Rachid Finge (14:20):

Okay. So, seems tricky to me is that you get feedback from a highly diverse community to improve what the camera captures and, and what the end photo looks like. Seems to me could also be difficult then

to land at the place where everyone is happy. So how do you make sure that the improvement is perceived as a improvement by a wide variety of people?

Florian Koenigsberger (14:44):

Absolutely. That question of evaluation is a big one, and it's also a complicated one, as you mentioned, because the curse of trying to build for everyone is that you can never really have everyone's opinion and you can do the best that you can with the proxies that you have. So going back to this earlier point about establishing partnerships with these photographers, I think of folks like Deun Ivory, who's an extraordinary photographer and visual artist out of Houston, Texas, or Jomo Fray, a wonderful cinematographer based out of New York, who's write that name down if you don't know him, I think that's going to be one that we're going to remember for generations. We basically used these experts as proxies for the community. Right. So if we know that you've produced work that everyone points to as exemplary of how our communities want to look, right. I think of Joshua Kissi's extraordinary editorial work, really bringing out the rich tones of black skin in particular. Then we know that we basically are starting from a reliable, trustworthy point outside of Google that says, people have already said that they like the way that these folks are making their work. How can we get closer to that kind of rendering, right. Without going too editorial and stylized, because ultimately, as you said, we are making a tool for everyone, and you want something that retains that basic faithfulness and doesn't look overly stylized or edited. And of course there is more to do here. I think when it we come to what does the road ahead look like? Part of that is continuing to investigate places where the technology is not up to par. You know, I never talk about this as a silver bullet and nobody on our team would. This requires a continuous investment and reinvestigation of where things are working and in what scenes we might still not be performing up to par.

Rachid Finge (16:31):

Yeah. Definitely want to ask more about long term. But before I do that, Florian on the Made by Google podcast, we talk a lot about AI and you know, how we use that, put it in the hands of all our users. So it sounds like AI has quite a role to play when it comes to real tone. Does it indeed show up in the technology you mentioned?

Florian Koenigsberger (16:51):

Yeah. This is a hot topic these days, as I'm sure anybody who's been on the internet in the last couple of weeks knows. I had a catch up with one of our engineers, Anny, about this yesterday because I think the company is very much looking at, you know, what are the definitions of artificial intelligence that we're all rallying behind as we define this future roadmap? One example of AI showing up in the underlying technologies in real tone that I would point to is the relationship between what we refer to as our skin tone classifier and a feature called frequent faces that our users can turn on. So starting with the skin tone classifier, we have a tool that helps us understand in a given image where on a spectrum of skin tones, we should locate a given person's skin tone. Mm-Hmm. <affirmative>. And the reason that this is important is because it then allows our auto white balance models to be optimized for a more accurate, faithful, and personalized representation of that skin tone. Right. So you start there and then you go, okay, if I'm able to understand, for example, that a certain face is showing up much more frequently than another, let's say it's my phone, for example mm-hmm. <Affirmative> and I happen to be pretty vain and I take a lot of selfies, right? Right. If I turn on frequent faces, then the camera starts to see that face with a higher frequency than any other. And I want to be clear from a privacy perspective, this is not the same thing as facial recognition. Right. It's not the camera saying, we know that that's Florian's face and we're making an optimization for Florian. It's just we know that we've seen that particular face a lot, and we're

going to make an optimization for that particular face without knowing who that face is. Sure. So as the camera starts to accumulate that data, let's say over the course of several months, I start to upload hundreds of pictures of myself.

Florian Koenigsberger (18:41):

Right. Some taken by me, some taken by other people across a wide range of lighting conditions. What happens is the skin tone classifier starts to benefit from that frequent face data because it's able to say, ""Okay, I've now seen this face across eight or nine lighting conditions, and I know that this is a face that we should be prioritizing based on the frequency that it's showing up. So I can start to deliver even more faithful renderings of that face over time based on its frequency"". Right. So you're seeing the relationship there between this classifier and this feature, which by the way, users have the option of turning on or off. I think it's off by default, but you can find it in the camera settings to use artificial intelligence to deliver a more accurate and faithful experience in the camera.

Rachid Finge (19:28):

That's amazing. I thought frequent faces would only help to, like if there was a group picture, put the right person in focus, so to speak. But there's a lot more to it than that, so that's really great to know. So looking at the long term, what is next for real tone? How are you working to make it better?

Florian Koenigsberger (19:44):

Yes. The future, one big consideration for me and our teams is thinking about the scale beyond Pixel. Obviously as a company that makes tools for everyone, we are deep believers in the notion that you shouldn't have to buy into a specific phone or technology just to feel like you can have a fair experience. Sure. And so I've mentioned Google Photos early on, we've already started to extend that work. There was some really cool work done available to the public today in the filter set that's available in Google Photos where you'll find that a number of the latest editions are actually made in partnership with some of the real tone experts that I mentioned earlier. And you know, I think it's always honest that as Googlers, sorry, it's always important as Googlers that we're honest about this, I was not a big fan of our pre-made filters before this project, especially as a photographer, I'm used to going into something like a light room and having a lot of nuanced tools or capture one and really getting into an image. Right? But what I've found is since this update, I am actually using our filters on almost every image I take on the phone because they've done such a brilliant job of really taking a nuanced approach to the warmth of skin, right. And the texture of skin and the brightness of skin that I find actually enhances those images in a really lovely way. And I've gotten wonderful feedback on that. So that's one first step I think in the path to scale when we look at the Android ecosystem on which pixel runs, of course there are a lot of considerations there about all of the other phone makers in this space and how we can start to share more of what we've learned with other players. I think we've intentionally been very public about our process in getting to real tone, partnering with those experts, breaking down some of the component technologies, because candidly, if you're another filmmaker in the space and you can already see that level of information, obviously some of the tuning adjustments that we're making for Pixel are specific to pixel based on our hardware components and our software in the phone, if you understand the lanes or channels where we had to make changes, it already gives other players in the space a really meaningful starting point to investigate within their own tools, am I delivering a fair experience? Right. Where are there opportunities for us to make improvements?

Rachid Finge (22:01):

So at the very least, the real tone technologies can be a inspiration for other smartphone manufacturers, but could there be even more than just, you know, showing the way?

Florian Koenigsberger (22:12):

Absolutely. There are limits to what we're able to share at this point, of course, but there are some initiatives in the works considering that Android ecosystem and thinking about how we vet what counts as an Android camera, for example and how we might improve some of the standards associated with what does a camera have to be able to do in the way of face detection, for example, across a range of skin tones to be able to be certified as an Android camera. Of course there are video conferencing tools that are important piece of this as well. Thinking about something like Google Meet, where to our earlier conversation about showing up to the world, even right now, right. We're on a video conferencing tool, and I can say from what I'm seeing on screen, there are varying levels of faithful representation that our computers are giving us back. Right. So really starting to think up the funnel, where are the places where we can have the highest impact affecting the largest number of products, and thereby user experiences in that funnel. Of course, part of the reason that this isn't, you know, fixed overnight is because there are so many different device dependencies mm-hmm. <Affirmative> from hardware tooling to software algorithms to device capacity. Right. How many processes can run at the same time to render that end result? But this is one of those big missions that's going to be a little bit slower burn than the initial real tone launch on Pixel.

Rachid Finge (23:33):

But I'm sure the, the end result every year of as you update real tone, probably bring it to new places. It will definitely be worth all the work and seeing the results.

Florian Koenigsberger (23:42):

I certainly hope so. Yeah.

Rachid Finge (23:43):

Absolutely. Now Florian, we close every episode with a top tip for our listeners. How can our listeners get the most out of real tone? Or how can they be more inclusive in their own photography?

Florian Koenigsberger (24:00):

That's a beautiful question. So a couple things. First, I'm going to harken back to frequent faces, I have found if folks are open to that experience, that it really does help in the long-term, allowing the camera to learn a little bit about who is being photographed, and again, not which specific people, but which faces are appearing more frequently, leads to more accurate results over time. I can say that as somebody who photographs my family often, where we have a range of different skin tones that can really be helpful. My father is a white German man. My mother is a black Jamaican woman. My brother is a little bit darker than me. So we've really got that tonal range in there. Right. On the question of how to be more inclusive in your own photography. I think if this is your first time hearing about the notion that folks with darker skin tones or folks of color might have a different experience with photography, I would give that a Google. There is a lot of written material, some of which has been published by our expert partners about ways that you can improve your photography and take into account tonal range, especially in images where you might have somebody with really light skin and somebody with really dark skin. There's a beautiful explainer video by Vox that goes into the history of the Shirley card and some of the racism built into camera technology in the early days. I think that's foundational reading,

and if I could give people one recommendation for a text that really shifted my lens pardon the pun, on this work is the vision and justice issue of Aperture Magazine guest edited by Dr. Sarah Lewis. I think that that should be part of every curriculum in America and potentially in the world, but it is a deep examination of the historical African-American contributions to the mediums of photography and film, and it really helps people understand like, oh, there's a whole other side to this medium and its development that the technology industry, I think is behind on. Right. And hasn't historically paid as much attention to that could paint a really beautiful future for so many people around the world. So those, those would be a couple of my tips.

Rachid Finge (26:03):

So in other words, read up. I think that's a great way to close out. Thank you so much, Florian for talking to us and looking forward to see what you do next with real- tone.

Florian Koenigsberger (26:11):

Thank you so much. It's been an absolute pleasure, Rachid. Thank you.

Rachid Finge (26:15):

Well, I've heard about Florian's work before, but it just hits differently when you hear him talk about it. It's hard to understate how much work there's left to be done in cameras when it comes to fair representation. So I'm glad we have Florian and the team working on it. Just Googling real tone or in fact, Shirley Cards will teach you a lot about the topic, and there's a lot to explore. Thanks for listening to this episode of The Made By Google Podcast. The next one is out next week, so if you're not subscribed yet, just hit subscribe or follow or whatever you call it in your podcasting app. And make sure you never miss an episode of The Made by Google podcast. Thanks for listening. Take care. Talk to you next week.