

Rachid Finge (00:01):

Hi there, and I'm so happy you made it back. Welcome to season two of the Made by Google podcast. I'm your host, Rachid Finge, and I'm here to introduce you to the people who work on Google's products. We had amazing conversations in season one and we have a fantastic lineup of guests for this season as well. Now, in case you stumbled upon our podcast and aren't subscribed yet, this is a great time to do it. So hit subscribe, follow, or whatever it's called on your podcasting platform so you won't miss any episode of this season. Now, if there's anything clear to me about 2023, it's that 2023 is shaping up to be a year in which AI is a big talking point for us at Google. Of course, AI really is nothing new. We've been an AI first company for years and we've been working on AI for decades. And if there's one area where AI is having an incredible impact on millions and millions of people already, then it's in photography, it's often called computational photography, and it's basically the art of AI helping you to take better pictures or making your pictures better after taking them. And that's what we're talking about today. Two features made possible by AI to help you fix up stuff after taking a photo sometimes even many decades later. So from the Google Photos team, please welcome Zachary Senzer.

Rachid Finge (01:36):

Zach, welcome to the Made by Google Podcast. Happy to have you today. Please tell us more about your role at Google and how you ended up here.

Zachary Senzer (01:43):

Yeah, thank you so much for having me. So I'm a senior product manager at Google Photos, leading our editing and computational photography experiences, and I started at Google as an associate product manager. We call that APMs. Mm-Hmm. <affirmative>, where I had the opportunity to learn what product management was, and I truly love the experience. I then joined Google's full-time APM program starting off on Google Assistant before ultimately rotating onto Google Photos where I've had the chance to work on many different parts of the product over the years, including editing, which is where I am today.

Rachid Finge (02:14):

Now, people who listen to the Made by Google Podcast frequently know that I like to go into internal directory where we have personal mission statements. Every Googler has one. Yours is "Help you bring out the best in your photos and videos", which I guess is pretty self-explanatory, but what does that mean to you?

Zachary Senzer (02:31):

Yeah. So we all take photos and videos for a reason such as sharing a really important life memory with family, but many times myself, and I'm sure you find that there's this gap between the photo and video that you have. in the photo and video that you actually want. And it could be that the photo of a moment isn't quite how you remembered this moment. Maybe it's very blurry or perhaps you wanna do something a little more creative before actually sharing this photo or video. And so my mission is to help you bring out the full potential in your photos and videos such that you can achieve what it is that you're trying to actually achieve with them.

Rachid Finge (03:04):

Right. So if I took a nice photo with my pixel camera and I made it even better with Google Photos, I have you to thank you for.

Zachary Senzer (03:11):

Exactly. You have me and a lot of other wonderful people on my team.

Rachid Finge (03:15):

Today's guest works on one of Google's most beloved products Google Photos. Zachary Senzer is a senior product manager and has his focus on the editing features. It's not only the basic editing stuff that Zach works on, like fixing exposure and colors, but as you'll hear shortly, he works on much more advanced editing tricks too. Zach is happy to admit that he himself is not a great photographer, which gives him an edge in deciding which editing features most people will appreciate. Zach also has a surprise or two for us at the end of the episode. I hope you'll enjoy our conversation. Now, today we're talking about two really neat photo features on Pixel Photo Unblur and Magic Eraser. But before we get into the nitty gritty on how they were, could you perhaps set the scene for how computers are changing photography and what actually computational photography is?

Zachary Senzer (04:09):

So a lot of these early camera phones were all about these old school type optics with bigger and better lenses, but computational photography recognizes that camera should be more than these physical optical machines. Right. An AI camera relies on this special software to give people the ability to snap images that previously only high-end cameras were able to capture. And for example, technology that can snap a dozen or so pictures in rapid succession and then use this AI software to align and combine them into a single image with great lighting and free of any blur from camera shake. And this technology is actually what powers pixel's night site feature in the camera where photos shot in potentially very low light ends up looking really bright, crisp, and detailed, which is a result that previously would've required this professional camera a tripod in quite a bit of photography know-how.

Rachid Finge (04:57):

Yeah. I remember seeing the first photo like taken with night site, which was like crazy, like black magic almost. Now what about those imperfect shots for people that perhaps don't have a pixel or, you know, when Pixel wasn't around?

Zachary Senzer (05:09):

That's a great question and it's actually a key focus for us. On the Google Photos team, people stored their most precious moments within Google Photos, and for many of these moments, people just simply didn't have the chance to capture them on a high-end smartphone or a very high-end camera or maybe they didn't have a chance to actually configure their camera just right, even if they had one of these high-end devices. And computational photography and AI more broadly allow to improve these images after they were captured, whether it be a photo that you literally just took an hour ago, or a photo that you took decades ago. And this is actually the theme of our most recent fixed on Pixel Super Bowl ad that just aired, which shows how features like Photo Unblur and Magic Eraser use AI to fix these imperfect shots.

Rachid Finge (05:49):

Yeah, it's a great ad. Amy Schumer had a lot of access to remove in in that commercial, what I remember from it. So we'll talk about magic Eraser in a minute, but let's get to Photo UnBlur. First, people who aren't familiar with the feature, could you quickly describe what it is and how you would use it on Pixel?

Zachary Senzer (06:06):

Sure. Photo Unblur uses machine learning to improve your blurry pictures. As the name suggests, it works on photos that were again, recently captured, or much older photos that you have in your library. Mm-Hmm. <affirmative> helping you relive the moment as clearly as you remember it. And if you come across a blurry photo in Google Photos and then tap edit on it, the app then suggests Unblur in the editor. And once you tap Unblur, Google Photos will start analyzing the photo and then shows an edited

version that ends up having less blur in visual noise. And there's also a slider so that you can go in and adjust the amount of Unblurring before you ultimately save and then share that photo.

Rachid Finge (06:42):

Now I think we could probably talk for hours about how this actually works behind the scenes, but could you give us sort of like a 101 version on what happens when you tap the LER button?

Zachary Senzer (06:53):

Photo Unblur uses the series of machine learning models that run on a user's pixel device, and it uses it to detect and reduce the blur in visual noise. And I think a big question is how do we actually know what that Unblurred image should look like?

Rachid Finge (07:06):

Exactly.

Zachary Senzer (07:07):

And these machine learning models we're trained using these data sets that contain images with these unwanted elements of an image. So for example, images that have blur and have visual noise in it. And we compare that to another dataset, which has a set of images that don't have these unwanted elements in the image. And the source of blur in an image can vary from a lot of different reasons. Maybe it's an imperfect lens, imperfect focus. Maybe the object you're trying to capture is moving around or the camera itself is shaking, or it could be a combination of these different things. And the photo Unblur feature can address more challenging scenarios where there's significant motion blur, noise compression artifacts, and mild out of focus blur. But we find that the best results that users get with this feature is when the blur is due to moderate camera shake and slightly blurry images that also include faces. And in the event that maybe your image isn't blurry, but there's some visual noise, or potentially it has some JPEG compression artifacts, Unblur will then enhance it and remove most noise in compression artifacts without actually altering the underlying characteristics of the photo. And overall to basically summarize the photo Unblur feature helps you improve the quality of the whole photo and also make specific improvements to any faces that are present.

Rachid Finge (08:17):

I guess as a product manager, you're probably one of the first people to experience a feature like this, and maybe you also get to decide, well, is it compelling enough? What do you remember from, you know, developing it and perhaps seeing it for the first time doing something great. Do you remember, was that like a personal picture or?

Zachary Senzer (08:32):

Yeah, some of the first things that I did when I was getting my hands on this technology was growing through all the older photos that I have within Google Photos. photos of me when I was a baby or a toddler back when the cameras were not the high-end pixel devices that we have today. A lot of them suffered from that blur and visual noise. And I was amazed how I was able to take these images and transform them into images that looked like they were captured on very high-end devices. This sort of transformation of these really old precious memories was a pretty magical experience to go through for the first time.

Rachid Finge (09:02):

Definitely. And I do have the same experience with pictures that it's just incredible how this is even possible. And I guess the same goes for a Magic eraser. People might be slightly more familiar with that because that feature launched with Pixel six instead of seven, so it's been there slightly longer. Again, could you quickly explain what is the Magic Eraser?

Zachary Senzer (09:21):

Magic Eraser can basically detect distractions in your photos, whether it be small people in the background or things like power lines. And as was the case with Photo Unblur, you can get a suggestion to use Magic Eraser when you open a photo in the Google Photos editor that the app detected might have things in it that you wanna remove. Or if we don't detect that, you can then go into the tool section of the editor and use it on any image that you have. And when you enter Magic Eraser, just like Photo Unblur, it'll analyze the image and in the case of Magic Eraser, it'll figure out things that you want to suggest to remove. And then you could tap, erase all to remove those suggestions. Or let's say we didn't suggest exactly what you're trying to remove. You can also circle or brush what you want to remove. And the best part of my opinion is you don't need to actually be precise and Magic Eraser will actually figure out what it is that you're trying to remove. And in the event that you don't want to remove a distraction entirely, let's say it's a really key part of that image, but you want it to blend in a little more and be less distracting, you can then use the camouflage mode within Magic Eraser to change the color of these distracting objects in the photo. And in just a few taps, the objects, colors and shading will blend in really harmoniously and naturally with the rest of the photo.

Rachid Finge (10:31):

Now if you think back to what you just told us about Photo Unblur, you basically teach a computer or a phone, you know, this is what a great picture looks like and this is the same great picture, but now it's blurry and perhaps, you know, at some point the AI kind of learns how to fix it. Right. But with Magic Eraser it seems much more complicated in a way because now the eraser needs to sort of guess well was like behind a person in the picture. How would you teach a computer that trick?

Zachary Senzer (10:58):

Yeah, so I'd say there's two main parts of Magic Eraser when it comes to AI. First is figuring out what are the distractions? And then once we know that there are distractions, what would the image actually look like with those distractions removed? And Magic Eraser uses more machine learning models that run on device that detect these distractions in the images. Again, let's say it's small people in the background and it'll suggest them to remove. And these models were trained to basically understand what the likely subjects of a photo are and then who or what might be in the background that you might want to remove. And again, in the event that you wanna remove something that's maybe not suggested to you, you can do these very coarse gestures on the image like a circle. Right. And we could then predict what you meant, refine the selection to only what you're trying to remove, and then erase it from the image and to your key question about what should actually replace the distraction that's there. Magic Eraser will then analyze what's nearby in the image, and based on learning from the training data, it'll fill in the background and predict what the pixels would look like if the distraction wasn't actually there.

Rachid Finge (11:58):

Again, you were probably one of the first people to experience this. Do you have, because I know there are a lot of people who use the feature, it's been around for a little bit longer than Photo Unblur, any particular examples you've seen maybe in testing or maybe even in the, in the real world where you were like, 'Wow, this is incredible'?

Zachary Senzer (12:14):

Yeah. The focus of the feature has been in, in our marketing a lot of removing background distractions, but we've definitely seen people use Magic Eraser in very creative ways. So the Super Bowl ad that you talked about earlier shows Amy Schumer removing X's from her various images. And then on social media, I've also seen a lot of creative ways that folks have used the feature, whether it be to do things

like clean their rooms or reduce the number of people that appear in long lines. It makes you wish sometimes that those tests are actually that easy in real life.

Rachid Finge (12:41):

Oh, definitely. Yes. I would like a real life magic eraser, so if you could work on that, that'd be wonderful. Outside of Google Photos now, I think this is one of the best examples, you know, Photo Unblur magic Eraser, but you know, the whole computational photography, I think it's such a compelling example of where AI touches people's lives in such a great way. Where do you think we're going next when it comes to AI and photography?

Zachary Senzer (13:07):

Overall, I'm really excited about continuing to leverage AI to democratize these traditionally challenging tasks in photography to help people enhance their content and be overall more creative. And AI has the power to simplify these tasks with a very intuitive interface and make it so that you don't need to be this expert to make your photos and videos look amazing. And a lot of what we've talked about so far today has been on the photo side, but I'm personally really looking forward to doing more with videos. We find that people are capturing and creating way more with videos than ever before. And getting the perfect video is arguably even more challenging than it is for photos. You have many different tracks going at once, many different aspects and dynamic nature of the scene, and there's so much potential in the space to truly revolutionize what people can do with their videos.

Rachid Finge (13:53):

Yeah, I guess you will be busy trying to figure that out because that seems like, you know, most videos have 30 pictures a second, basically, so it must be at least 30 times harder to do that.

Zachary Senzer (14:04):

That. Definitely harder for lots of exciting stuff on the way.

Rachid Finge (14:06):

Oh, that's wonderful. Hey, you mentioned, you know, democratizing stuff. I was just wondering what do you mean by that? Because I, I hear that a lot in tech circles, but I'm not sure if most people know, you know, what us, techies, mean when we say

Zachary Senzer (14:19):

That. Definitely. So a lot of times these really complex edits or this complex functionality and the operations that we talked about today, whether it be removing something from an image or taking a blurry image and making it look sharp was previously only accessible to people who had either very high end devices, high end software, or just general ability and expertise to know how to do these different operations. And personally, I'm not this professional editor or someone who has all their expertise and software required to make these really complex edits. And so for everyday people like me, I'm trying to figure out how I can get those really, really awesome photos and videos that I see out there in the public. So when I talk about democratization of these different sorts of edits and features, to me means making it so that myself and any sort of everyday person can go in and get those really awesome photos and videos in just a tap without that expertise required.

Rachid Finge (15:13):

It sounds, you know, as a product manager it's probably a superpower to actually not be a pro photographer in this case, but you know, basically be like the rest

Zachary Senzer (15:22):

Of us. It definitely helps me empathize a lot more with different people who are ultimately going to be using this. And it also again, breathes in that sense of magic when I get to try these features for the first time, given that I don't necessarily have that familiarity of using this really fancy software or have that domain expertise in terms of how to actually perform these different editing operations.

Rachid Finge (15:40):

Now, Zack, you know, regular listeners often made by Google Podcast know that we close each episode with a top tip for our listeners. What is the top tip you would share with them?

Zachary Senzer (15:54):

To date, most of our newest computational photography features have only been available to those who have Pixel six, six A and seven phones, which have this Google Tensor chip inside that helps us run all these machine learning models on the user's device like Magic Eraser and Photo Unblur. And for the launch of Magic Eraser on Pixel six and later the six A in Pixel seven, we worked really closely with our AI researchers at Google to create machine learning models that run well on Tensor. And since then we've been working really hard to optimize those models on other devices. And we've recently announced that Magic Eraser and other exciting benefits are now available to people with older pixel devices and to Google one members on Android and iOS devices. Right. And I'm really excited that more people will be able to use this magical technology that powers Magic Eraser to improve their photos. And so would really love for everyone to give it a try. And we can't wait to see all the awesome ways that you'll use the feature.

Rachid Finge (16:48):

That's a great tip. So for people listening, do look for Google One and get some amazing benefits from that in Google Photos. Zach, thank you so much for talking to us. Thanks for sharing the tips and talk to you soon.

Zachary Senzer (17:01):

Again, thank you so much for having me.

Rachid Finge (17:04):

So there you have it, Magic Eraser and Photo Unblur now available to more people than ever through Google one. So make sure to check out Google one and don't forget it comes with many more benefits from Extra Storage and Google Drive to a VPN to protect your internet traffic. Thanks again to Zach for coming on the Made by Google podcast and I'm already looking forward to the next episode, which is about spatial audio. That's when sound from movie seems to be coming from every direction, including behind you. And all you need is a pair of pixel buds. So more about that next time. Thanks again for listening. Take care and talk to you soon.