EPISODE 1133

[INTRODUCTION]

[00:00:00] JM: Factories require quality assurance work. That QA work can be accomplished by a robot with a camera, together with computer vision. This allows for sophisticated inspection techniques that do not require as much manual effort on the part of a human. Arye Barnehama is a founder of Elementary Robotics, a company that makes these kinds of robots. Arye joins the show to talk through the engineering of Elementary Robotics and his vision for the future of the factory floor.

We're looking for writers and podcasters for Software Engineering Daily. If you're interested in either of those roles, email erica@softwareengineeringdaily.com. We'd love to get your applications.

[INTERVIEW]

[00:00:45] JM: Arye, you welcome to the show.

[00:00:47] AB: Awesome, Jeff. Thanks for having me.

[00:00:50] JM: You work on Elementary Robotics. Explain what you do.

[00:00:54] AB: Yes. So I'm CEO here at Elementary, where we focus on automating quality assurance and quality control with visual AI and robotics.

[00:01:07] JM: And what does that mean in more detail?

[00:01:09] AB: Yeah. In detail, as physical things get made in the world, they're going through production lines, and we are really helping those brands and manufacturers drive new insights around which parts pass inspection and can be put together into an overall product to make sure we're limiting failures or cosmetic issues that reach the final customer. So we're helping things get put together appropriately and making sure there's no defects on those parts as they get put together. But then, also, we're really trying to help drive insights for our customers so that they better understand kind of the root cause as well of where these issues might be coming from. So, overall, trying to reduce scrap and waste in our industry and really superpower manufacturers to make quality goods.

[00:02:03] JM: Quality assurance for physical items has historically dependent on manual inspection. How have improvements in robotics changed quality assurance?

[00:02:12] AB: Yeah. So there's this joke in the industry regarding manual inspection that if you ask 10 inspectors what a defect looks like, expect 11 answers. And all of these jokes obviously have a bit of truth to them. And that's because on the manual side, you can think about sitting there, looking at an item all day, every day. You've got visual fatigue. You've got mental fatigue. And so there's – Generally, it's also somewhat subjective. Something that I think is a defect. You might be looking at slightly differently.

And so all of this kind of adds up to just general inconsistency in inspection. And so bringing tools that are kind of data-centric tools like machine learning and then automation allows you to take some of the subjective nature out of that to make things more repeatable and data-oriented and then allow those humans to kind of put their knowledge into the system and then operate quality and operate root cause, which we just touched on, which is a bit more of the kind of high-level goals as supposed to just the actual visual eyeballs on the product itself.

[00:03:23] JM: Telling me more about some of the recent developments that have made it possible to use AI for quality assurance.

[00:03:30] AB: Yeah. So there are some exciting developments in AI that have continued to accelerate. There's, I think, first of all, just the overall maturity of AI and any use at scale. And so, just at the very beginning before even kind of any new architectures, there's just the maturity of AI as a technology that can be deployed into production. The tooling that needs to be used for that, the data handling. And so, all of kind of the pipes to make AI possible have accelerated to allow this to be more scalable in an enterprise manufacturing environment.

And then there's also new research that a lot of us are looking at or making use of in the industry that's kind of cutting edge around using lower data sample sizes. So, not every manufacturer, when you go in, is going to have a data sample of thousands and thousands of objects already built for you. And so you've got to be able to get useful and start providing value with lower and lower data sample sizes. And there's a lot of research in that space that's also having an impact on our ability to apply it in manufacturing.

[00:04:38] AB: What are the areas of QA where computer vision is mostly utilized?

[00:04:43] AB: So, by that, do you kind of – I guess, for our audience, the difference in terms of what I think you're asking on the computer vision side versus machine learning side, is there are traditional computer vision use cases in quality assurance and manufacturing that are a bit more rigid. So, things like edge detection, blog detection, presence detection, some computer vision algorithms for dimensioning. And so, whereas machine learning is more sample-based and kind of labeled data and training-based.

And so there are areas where these computer vision tools can be quite useful given a very rigid set of rules that you know you're going to be inspecting for. And so, generally, I think you can think about the two, computer vision versus machine learning in that way of how good could a rules-based system be to find these defects versus are you looking for more generalized things and you don't want as much of the rigid rules-based system and you want a more flexible system. And so, it's our job to find the balance between those two.

[00:05:51] JM: What's the goal of Elementary Robotics?

[00:05:54] AB: Yeah. We are focused on allowing visual AI to be as easy to use and easy to deploy in the manufacturing world as possible in order to make better, more scalable products.

[00:06:09] JM: And how did Elementary get started?

[00:06:12] AB: Yeah. So, taking a step back, just a bit of background on myself as it leads into that. I previously cofounded a company called Melon, which allowed me to do everything from live in Shenzhen, China for six months and bring up manufacturing production lines, shipped thousands of units of hardware. We were the startup in residence at IDEO, the design firm where I really fell in love with design. So, we were the startup in residence at IDEO, the design firm.

And then eventually Melon was acquired by a private equity firm, and we got rolled up into a company called DAQRI, which had raised about $500 million working on industrial augmented reality. And there I really fell in love with these use cases in the industrial world. And I specifically saw a lot of maturity and scale that could be had on the computer vision, machine learning side that was happening in AR, but wanted to bring that back and apply that to robotics and use robots for how we could bring this amazing kind of software on the CVML side to life in the manufacturing world. And so, that was some of the initial inspiration behind Elementary, was using robots to bring to life some of the great software I had seen in computer vision and machine learning.

[00:07:33] JM: Elementary started with investments from some companies as customers. How is the development process been with those companies that you started out with? What were the requirements?

[00:07:44] AB: Yeah, it's really awesome having these customers that are also partners early on in the lifecycle, because it really enables us to do diligent customer discovery and have really deep discussions around pain points, and problems, and integrations. And so, I think we've learned a tremendous amount from that. And as a kind of design and product-focused founder like myself, it's always great to be as involved. You want to be as involved with customers as you can be as early of a stage as you can be. Get as much feedback as possible. And so, it's been super beneficial to us as a company, and allowed us to make sure that we’re solving real problems for real customers.

[00:08:27] JM: What does a 3D inspection entail?

[00:08:29] AB: Yeah. So one thing that's really unique about Elementary systems as you’re kind of bringing up here, is our ability to do inspections of 3D objects from multiple angles given the integration between our vision systems and our automation robotics systems. So, we have the elementary gantry, which allows us to do 3D motion. And then pairing that with vision systems that are then able to take these multi-angled inspections.

And so, you can imagine an automotive. These are big, very different geometric shaped parts. And so you can't always just say, "Hey, I'm going to take a top-down static image of this. You really want that flexibility and you want to look at things from different angles. You can imagine even as a human, when you're looking at your iPhone, I always use this example, but this isn't exactly it. But when you're looking at your iPhone and there's a scratch on it, you can't always see it. You have to move it. Look at it in different angles. Hold it under different light to actually be able to see that. And so, similarly, these multi-angled inspections can really drive value to finding new defects in kind of different ways that you might not have been able to previously.

[00:09:43] JM: What kinds of inspections are done on products for Q&A? Could you maybe give an example?

[00:09:50] AB: Yeah. So, it varies across industries. But things that kind of stay consistent are there are assembly inspections where all of the right pieces put on this product. Was this product assembled properly? So, there're the assembly verifications. There is also kind of final inspections. So, things like was the label put on properly? How does it look cosmetically? Are there any scratches that a customer wouldn't otherwise want to see? Are there any defects on the paint that a customer wouldn't otherwise want to see?

So, as you think about how a product gets put together, there are different inspections that run throughout those stages of that process. So, assembly and all the way through to kind of final cosmetic verification.

[00:10:39] JM: Elementary has human in the loop at the core of the design. Can you explain what human in the loop involves?

[00:10:47] AB: Yeah, absolutely. Great question. So, on the machine learning side, as we've all seen, there is still a requirement in a lot of cases where someone needs to be doing something like setting a threshold or labeling data. And so, you brought up earlier in the conversation that many inspections have previously been done manually. And to that note, you want to take that knowledge. You don't want to just rely on the data that is coming in, but you also want that human knowledge that's kind of been learned and you want to still keep that in the loop. Have that manage what's going on. Have that label data and set the proper thresholds for that production line, because each production line is going to behave differently.

And so, you still really want to keep the human in that system to oversee it. I think Google has a great study about doctors plus machine learning make the most accurate predictions on X-rays as supposed to just machine learning or just doctors. And so, kind of similarly, we want to keep that human in the loop and have them be a part of the analysis as well as some of the kind longer-term root cause insights that are being driven.

[00:12:01] JM: Tell me more about how remote inspection works.

[00:12:04] AB: Yeah. Especially given where the world is at right now, remote is super important, because it's hard to get boots on the ground in factories at the moment. So, you definitely want to be able to use IoT networks to be able to support your customers remotely. And, also, one thing we focus on is enabling our customers to be able to support themselves remotely, because even some of them are having issues getting in their facilities. And so, just having that flexibility where they can access our devices remotely is super interesting and definitely providing value back to a lot of our customers.

[00:12:43] JM: What does the hardware component of Elementary consist of?

[00:12:47] AB: Yeah. Definitely, for those listening, check out our website, elementaryrobotics.com. We show some of our hardware solutions there. We show the elementary gantry, which is a 5° of freedom system that allows kind of this multi-angle inspection that I was talking about earlier. So, it’s basically path planning and moving a camera around an object for these 3D inspections.

[00:13:13] JM: And what were the functional requirements when designing that hardware?

[00:13:17] AB: Yeah. That's where to your question about polling in customers early. Just really digging into diligent customer discovery of based on their requirements. So, what sizes of their parts? What speeds of their production lines? And then how do we fit all that together and make these machine learning inferences? What compute does that mean? What resolution on the camera? This multifaceted problem space where we had to work all those together, kind of like it in any of these hardware-software products, there's always these tradeoffs between what do you solve in hardware? What do you solve in software? How flexible do you make it and scalable, versus focused and vertical? And so we had to work through a lot of that problem space and are continuing to as we continue to grow our product offerings.

[00:14:05] JM: So, Elementary Robotics can make 3D images of a product to detect defects. How does the initial data collection work?

[00:14:14] AB: Yeah. As you set it up, it's running. It's really easy for the human to program, first of all. And so, you can go in and you typically have a sample object and you can program it around that sample object. And then on the data collection side, once you have that set up and easy to run, you can run it in a data collection where you're not trying to make any inferences in the beginning. You're just collecting the data needed to then go deploy your inspections. And so that's where that kind of human in the loop piece comes in again.

[00:14:47] JM: How much data do you require for a specific product?

[00:14:50] AB: Yeah. I wish there was a hard and fast answer to that. It's less and less every day for us as a company. But it also depends on the type of product we're inspecting and the inspection requirements. But, generally, we are not talking thousands of images like some machine learning classifiers. We are talking tens or hundreds of images. So, significantly less given our machine learning approach. But, still, you need a robust sample size to be able to test on it and have confidence in.

[00:15:23] JM: What types of machine learning models do you utilize? Are the inspection images compared to like a baseline just like you would do and in GANs?

[00:15:34] AB: Yeah. There're a couple different approaches here, but we generally focus on anomaly detection for our customers. So, learning what good looks like and then being able to, from there, find anomalies and defects on their parts.

[00:15:49] JM: And are there separate models for separate kinds of tasks?

[00:15:55] AB: So, the models and the inspections that we offer to customers and let them set up are, yeah, focused on the type of inspection there trying to run. So, if they want to be running anomaly detection, that's one piece. If they have data and want to run a classifier on a specific type of defect, that's another inspection they can set up. And so, it is driven by the type of inspection they're setting up.

[00:16:18] JM: Tell me more about how a model gets deployed. Is it getting deployed on-premises? Or is getting deployed at the edge? Or is it getting deployed to cloud resources?

[00:16:29] AB: Yeah. Our models run on the edge largely due to just speed requirements and inference requirements on our systems in the manufacturing environment.

[00:16:41] JM: Okay. And tell me more about the edge deployment. What cloud provider do use? Or what kind of edge infrastructure do you use?

[00:16:49] AB: Yeah. So, we have the cloud piece, and that's where we do the retraining. So, it doesn't have to be. Some customers want their server running on-prem. And so, you can do that. If somewhere you're doing the retraining as you've collected the data, and then we kind of over our IoT like are pushing that back to edge compute where we're running the models and making these kind of real-time inferences. That edge compute – I think, right now, our models do run on Nvidia GPU's. But then kind of flexible in terms of what we're using on the CPU side.

[00:17:27] JM: What is the architecture look like in more detail?

[00:17:30] AB: Yeah. So, Elementary is pretty focused on automation scalability. So we've tried to take that into our software side as well, and we largely work around kind of a microservices approach to make things flexible and adjustable in terms of how we're deploying different modules? Where? Just as we just discussed, kind of edge versus cloud. And so we try to make that possible given it is a world where different customers that we're deploying into where different customers have very different requirements around where compute can live and where data can live. So, with that approach, just trying to enable all of those different types of deployments.

[00:18:11] JM: So what happens when you're doing a QA and a defect is found?

[00:18:17] AB: Yeah. So, on our side, we flag that, obviously, and we store that data and we save it as a fail. It then depends how we are integrated into the customer line, whether we then send them a signal maybe to their PLC system that controls their production line. Maybe to something elsewhere that's kicked out and then sent to rework. Or that kicked out and put in a scrap pile. Hopefully we're catching it early enough where it is kind of sent to rework on the earlier side and it's not a big change for them. And so, we're generally flagging that for them. And then depending on their production line, there are different processes that they might go through with that part.

[00:18:58] JM: Elementary can also help find and define new defects. Explain how this works.

[00:19:05] AB: Yeah. So, that largely comes back to your question around CV, computer vision versus machine learning. And so that the rules-based approach versus the anomaly detection machine learning-based approach, and the ability to, on the machine learning side, take a more data-driven approach to finding defects where it's finding these anomalies that you might not have programmed the rules defined. But now because they're there different from what the system has learned good should look like, they're getting flagged to you and you can start seeing, "Hey, maybe we didn't really have a problem in our process, and we are finding defects or anomalies, and we just didn't know to be looking for those."

[00:19:44] JM: Tell me more about your stack for the backend.

[00:19:47] AB: Yeah. So, at a high-level, happy to kind of share that I think to your question earlier, which really hit home on what's super important right now during COVID. On the backend side, we have enabled this kind of IoT link, which is super important these days not only for our customers as I kind of mentioned previously. But one other thing to add in there is just how valuable that's been for us as a company.

So, during these more remote times, we can't always be co-located with robots. Sometimes we've shipped them home with people. But, you've also got to be able to remote-in. And so we can do remote testing, remote development and deployment over that kind of IoT link. And that's been super valuable to our development team. And then, obviously, from there, there's all of the remote data access and storage in the cloud. And so, just really enabling that to be a more IoT-centric device has been super valuable to us as a company.

[00:20:46] JM: And can you tell me more about what cloud services have been useful to you?

[00:20:51] AB: Yeah. Now, we're very customer-driven. And so every customer kind of has different thoughts around which cloud we use and kind of maybe they're already using a cloud. And so, we have worked to this point to enable that from multiple cloud providers just given really wanting to enable our customers. But on our side, obviously, you can imagine we're working through with these different cloud providers that all have solutions around that machine learning piece, the data handling, retraining and data storage side. As you think along the lines of massively scaled machine learning image-based backend and cloud pieces, that's where we focus our efforts currently.

[00:21:39] JM: When you're looking at your stack, like are there are any – I don't know. TensorFlow, or more SageMaker, or some database that has been particularly useful?

[00:21:50] AB: Sure. Yeah. I mean, on the ML side, our team is definitely centralized around PyTorch and the services over there.

[00:21:58] JM: So, you've been in stealth mode for a while. You recently became publicly commercially available. How was the reaction been?

[00:22:05] AB: Yeah, it's been awesome. I am really excited to come out of stealth mode. Makes talking about what we do easier. And so it's been really great and great to just be more open with customers and see their interests come in. So, that's been really awesome, and I'm very glad we finally came out of stealth and are able to share the message more broadly.

[00:22:29] JM: What has been your biggest technical challenge so far?

[00:22:32] AB: Yeah. Robots are really a full stack problem. So, that has been a really awesome challenge to solve. And I think part of what gets our engineering team so excited to work on this every day is just you're integrating all the way down to embedded devices running motors and running all of the kind of controls code all the way up through the system, through the embedded ML, all the way up to the cloud and interfaces and web app.

And so, it's just a very broad team skillset and kind of diverse backgrounds that we pull from to enable a full stack solution like that to work together. And then go through your own testing and your own QA and be able to find if you're crushing bugs. Are they on the hardware side? The software side? And so, really enabling great full stack development has been a challenge and something that we love working on and kind of gets us excited every day.

[00:23:34] JM: Where you personally spend the most of your time?

[00:23:36] AB: Yeah. As most kind of founders and CEOs, it's always shifted over time with Elementary as we've continued to grow. I do stay really close to customers. And so, kind of especially now as well too, there's just a huge focus for me over there. So, rolling up my sleeves and working with customers and trying to always be listening to the voice of the customer and making sure we're solving a big problem for them. But, also, as I mentioned earlier, I kind of fell in love with design in the early days of my first company when I worked at – When we're the startup in residence at IDEO. So, still really love the design, human-centered design side to this. The interfaces that we're building, the product side and kind of collaborating with the product and engineering teams. So, I love that and I'm super passionate about side. But we continue to scale, definitely, and kind of laser-focused on the customer and customer growth side as well.

[00:24:35] JM: What does that design process look like in more detail?

[00:24:38] AB: Yeah. So on the design side, it's scaled as we've scaled as a company. And so, kind of matured the process to have specs and requirements and interviewing customers to understand their problems. Trying to get all of that information upfront so that as we go into the visual design process, it's based on data that we've been collected and then going through the user experience into the then user interface design and making sure that all fits within the overall architecture of the system. And then working with engineering to make those tradeoffs always of what the dream design does and looks like and behaves like versus how do we implement it and the timelines for those things. And so, there's a lot of data we're pulling in externally, and then a lot of internal discussions and tradeoffs that we make through the process. And kind of you end up with a really multidisciplinary result to push the design forward.

[00:25:41] JM: Can you tell me more about what programming languages you use?

[00:25:45] AB: Yeah. It is, again, back to the full stack piece. There's a bunch of different teams contributing. And so, there are different places that are more or less specialized. I said we used PyTorch on the machine learning side. Obviously, Python is very widely used. So that team is kind of deeply integrated on the Python side, and that carries throughout the stack to some degree. But then as you get closer to the robot and the embedded side, you start programming more on the microcontrollers, embedded C, C++ and kind of dig in over there to get things performant.

And so, it actually does scale. And we enable most of that through this microservices approach that I was kind of mentioning. We have this internal piece called Atom, which is our kind of internal core microservice OS.

[00:26:39] JM: What does that mean? Internal core microservice OS?

[00:26:43] AB: Yeah. It's basically our internal OS that our systems run off of that is a microservice architected approach to allowing these different modules in our system. You've got the ML module. You've got the camera module. You've got the embedded compute module to talk to each other and to do that through different programming languages.

[00:27:06] JM: Tell me more about that communication layer, that unifying communication layer.

[00:27:11] AB: Yeah. Well, our CTO would probably be able to dig into that a lot more. But what I would say is you can visit our GitHub and actually see some of that. That's probably the best way for people to dig into it.

[00:27:23] JM: Can you give me your perspective on the future of the factory floor? What it looks like in a decade?

[00:27:30] AB: Yeah. That's an awesome question. So, I am really excited about like more and more data-rich factory floor, and that's what I'm seeing, and that's kind of where the trends are going, and that's really where Elementary is also trying to help our customers. So, a data-rich factory floor where we can close the loop on more systems, that data is getting fed back into decision-making, fed back into designs and just really kind of closing the loop on a lot of what we do based on what's being collected and outputted and where anomaly detection is running and things you're finding. And then making more real-time changes based on what you're finding. And the final product being shipped out is still being tracked and fed back into manufacturing. Overall, just creating a stronger feedback loop around that more real-time data is a vision of the factory that I find extremely compelling and I'm working to help build with Elementary.

[00:28:26] JM: Do you see yourself expanding into products that would fulfill other areas of the factory?

[00:28:32] AB: No. We're really vision AI-driven right now, and there's a kind of huge market in front of us to grow there and a lot of problems to solve. And so, at the moment, we're just kind of laser-focused on that.

[00:28:45] JM: What are the other adjacencies in the computer visionary that you could expand into?

[00:28:50] AB: Yeah. Cameras are used a lot with robots right now. So, not areas that we are expanding into, but there's a lot of robot guidance, robot ML+ cameras for picking and closing the loop on different robotic systems. So, there're a lot of other areas robots are kind of closing the loop on robot systems. But, currently, not areas that we are looking at or planning to go into. We are just really focused on kind of area where we can be applying elementary systems to provide more insights around product quality, product inspection and quality assurance throughout the factory floor.

[00:29:33] JM: You've mentioned cameras several times. Do you build those cameras yourself? Or is there an off-the-shelf provider you can go to for cameras?

[00:29:40] AB: No. We don't build the cameras. Cameras are super specialty. And so, we do use industrial cameras and have kind of partners on that side of the world that we work with.

[00:29:52] JM: And what is the market for off-the-shelf cameras like? Is it competitive? Is the cost dropping? Is the cost prohibitive? Is it already cheap? Tell me more about that.

[00:30:03] AB: Yeah. I think the market for off-the-shelf cameras is really interesting. We are using industrial cameras. And so they're not super off-the-shelf. There are some websites you can go to in order them, but you do build relationships with these companies as well. And they're not cheap in the sense of that it's the same cost as the camera in yourself phone. But overall, the costs are going down just as camera use cases are just scaling. As AI scales, as vision scales, camera use cases are scaling.

So, I think there's a really good ROI in our industry for these types of use cases. Although, you've also got to get the compute to the right price point, and you got to take that full stack approach that we are mentioning to get the overall system there. But there are some really great camera providers out there on the market that are constantly pushing that forward.

[00:30:59] JM: Do those cameras all have similar APIs? Or what is the typical API between the camera and the software you build around it?

[00:31:07] AB: Yeah, they don't, but they usually expose a lot of similar things. But they all have slightly different APIs or languages that they support. And so, there is some customization around different cameras.

[00:31:23] JM: And are there any bottlenecks in camera technology? Like – I don't know. Frames per second? Recording rate or something that's important for improving the fidelity of downstream engineering applications?

[00:31:35] AB: I would say, currently, less on the camera side. More on actually the compute an inference speed side. So, more on the latencies around making the decisions and less on the core camera piece.

[00:31:50] JM: Yeah. Can you elaborate on that? So, like we're talking about PyTorch. Python is not the fastest language. But I don't know if the interface between Python and Torch. I don't know what Torch is actually written in. Maybe it's in C or something. Maybe the underlying infra that's executing is faster. But maybe you could tell me more about the inference side of things and how that –

[00:32:11] AB: Yeah. And, well, all of that is in Python, you can still – Kind of once you've got the trained model, you can definitely kind of take that down and shrink it and kind of get it running embedded and use different techniques to get it running faster embedded. And so, it doesn't have to be based on PyTorch and kind of those architectures. You can definitely push to get it faster given you're trying to make edge inferences. But still, there's just some limits even as fast as you get that around running ML on the edge. Passing data back-and-forth between CPU, GPU, shared memory and kind of how all of that adds up to speed of inferences that you're able to make. But I will say, a lot of people are doing a lot of great work in this space outside of Elementary as well. Nvidia is doing a lot of great work here. And so, that is definitely in an active area that people are pushing forward.

[00:33:05] JM: Tell me more about your perspective on the machine learning ecosystem. What are the shortcomings? What is machine learning getting really good at? What are you intrigued by?

[00:33:15] AB: Yeah. Based on deployments and kind of just being customer-focused right now, I think something it's always on my mind is just data handling and pipelines. And I think that's still a super important piece of making all of this work in scale. Now, I obviously love reading all of the latest papers and looking at self-supervised learning and different ML approaches. But at the same time, I think on the really practical side, I find just how do you pass data around? How do you label it? How do you deal with datasets and manage datasets is still just super important and really fascinating to watch where that tooling goes.

And, generally, as an ecosystem, the machine learning side, I just am super excited about where everything is going. Where the community is going? Where research is going? I'm really dedicated to seeing that continue to scale and help enterprise businesses.

[00:34:13] JM: If I am a QA tester at a factory and my factory uses Elementary Robotics, what does my job look like?

[00:34:22] AB: Yeah. So you asked earlier about human in the loop, and I think that's kind of the key. We see two different things. One is quality managers and quality technicians moving towards operating multiple systems. Overseeing Elementary systems and helping re-label data. Look for root causes and really kind of manage quality now through something that helps them scale up in terms of throughput. And then, also, focus on other areas and focus on other problem areas or areas that need their support throughout their production line.

[00:34:59] JM: And do you see the role of robotics as augmentation to the QA tester? Or does it obsolesce the QA tester?

[00:35:11] AB: Yeah. I'm very focused on and convinced it's augmenting the QA processes. So, everything from – As I mentioned, like you're shifting those people to still managing these devices. They're managing the thresholds of the inspection. Digging into the data for root causes, as well as there are certain times people running QA also have an assembly task. And so, now you're shifting them back to having more of their focus on the assembly task and giving them more time than having to kind of split attention and multitask between two things. So, I think it's kind of either enabling them to do more of other work or enabling them to manage higher throughput quality through the system.

[00:35:57] JM: You have one of the most impressive webpages I've seen together with a loading bar, because it's so graphically intensive. Could you explain a little bit about what people would see if they went to the Elementary Robotics homepage?

[00:36:11] AB: Yes. I would definitely check it out, because as you mentioned, it is very visual and there's a lot of motion on there. And so, I don't know if my words can even do it justice. Definitely, it's elementaryrobotics.com. And what we really tried to do with that was just showcase this motion plus vision concept. So, that's why it's graphically intensive and has different animations is to give people as you scroll through our website a visual insight and kind of look into running inspections of 3D parts and kind of multi-angled inspections. And so we thought it'd be to really help tell the story we wanted to build that into the website.

[00:36:53] JM: Can you tell me about building the proof of concept? The first version of the Elementary Robotics camera system?

[00:36:59] AB: Yeah. So, we've built a lot of different first versions throughout our time at Elementary as we've iterated through and kind of gone through diligent customer discovery to find that right that right product market fit. And so, in the early days as I mentioned, we were putting in place the infrastructure and architecture to enable all of these different systems to work together. And so, I think that was kind of the early engineering Lyft, was enabling that architecture and infrastructure that would say, "Hey, you can use motors and motor control with cameras and robots in different gears and joints." So, really, kind of taking that full stack approach to enable all of that and then balancing the, "Hey, we also know these gears are probably going to change and some of the motors might change." And so, kind of balancing that, "Hey, we're building the first version with wheels, and we want that to work so that we can get customer feedback. But we also know things are possible and likely to change." And so how do you allow for that flexibility in the future as well? And so, those are a lot of different tradeoffs and things we're thinking through in the early stages.

[00:38:11] JM: Tell me more about some of the hard canonical engineering problems that you run into it Elementary.

[00:38:17] AB: Well, I think there's always a question, and this is just a general enterprise startup question of integrations. And I think everyone's running different stacks, whether it's different ERP systems, or it's different PLC systems controlling their production line. There is always different integrations, and as CEO, it's always my job to look at the roadmap and sit there and say, "When do different integrations make sense? We know they're possible. But given the engineering resources we have balancing integrations with new feature development, just trying to make those right tradeoffs and what unlocks the next tranche of customers. The next value prop to our customers?" And so, how do we balance that alongside? Maybe the next customer really runs a different PLC system than we've ever integrated with. But, also, we've got to prioritize some other core software engineering work.

And so, I think that's something I think all enterprise companies deal with, is making those integration tradeoffs given that customers – There's usually one – There's the 80-20 rule. There's usually something that the majority of people are using, but then you always run into those different cases as you scale to more customers.

[00:39:34] JM: Are there any other canonical engineering problems? Things that you're seeing today that you think you will be seeing again and again and again in the next five years?

[00:39:42] AB: I think even as machine learning gets better, there's just always going to be improvements there. There's always going to be how do we use less data? How do we augment our data? And so, I think there're just open-ended questions there that we'll continue to revisit. And I think those are exciting challenges ahead for us.

[00:40:03] JM: Do you spend much time reading papers or talking to researchers? Or do you feel like most of your consumption of machine learning stuff is further down the pipeline? It's like well-developed productized machine learning?

[00:40:16] AB: Yeah. So we've got an amazing advisor that we work with on the academic side. And so, we chat with him and pick his brain. See what's coming down the pipeline. And that's been an amazing resource for us to have. And then, yeah, I'm known for paying out over Slack channels just at all different hours the things I'm finding online as I'm just looking through papers and publications. Maybe there's a new conference and kind of who won best paper? So, I find all of that fascinating. And think you have – That doesn't mean we're implementing it, but I love being aware of it, reading it and kind of discussing it all with the team.

[00:40:57] JM: Let's scale back and return to the beginning of our conversation. Could you just run through in more detail what happens in the production pipeline where Elementary Robotics is involved?

[00:41:10] AB: Yes. So, scaling out to little bit more of 10,000 foot view of a production line, you've got parts coming from suppliers that are going to get put into whatever you're producing. So, you've got these kind of supplied components, and you want to verify that those supplied components are good before you put them into your product. So, there's an opportunity for QA there. Then, as you're assembling your product out of those parts, there are verification steps throughout that process that you can use QA to make sure, "Hey, did we put this together properly along the way?"

And then at the end, you've got kind of the final piece, "Okay, overall, it's been assembled. We've verified along the way. And now we want to do final assembly verification," and overall that's usually where cosmetics can come in more, and how does the overall final product look before we ship it out the door? And so, there're just multiple areas throughout that stack that you can be looking and kind of running quality.

[00:42:15] JM: And let's again return to the future. So, let's say five years down the line, what is in the future of Elementary? What are the specific fields that you're excited about?

[00:42:27] AB: Yeah. So, great question. As we mentioned, on the what's hard part, I also think that's a huge opportunity for the future, which is that ML is continuing to grow and evolve. Whether it's the techniques, whether it's synthetic data, or being able to run more kind of simulated in terms of training and how that augments the data that we need to be collecting. And so, I think there's a lot of opportunity for growth there too.

Just continue to invest in the core principles of what we do at Elementary around ease of use, ease of set up, time to ROI. And so, things that I think we continue to invest in and grow just to make our customer experience better, easier, faster and continue to unlock more value for our end-users.

[00:43:16] JM: Can you tell me a bit about the sales and integration process? So let say I'm running a factory or I'm operating some multi-factory process engineering facility set. What's the process for getting a sale closed and then getting integrated?

[00:43:34] AB: Yeah. So we were pretty hand in hand with our customers right now especially being early-stage enterprise company. We really want that voice of the customer. So early on, we're talking to them. Understanding their pain point. Understanding where they're struggling with inspection. Where they'd like to augment their current capabilities? Whether it's improving throughput or improving yields and finding defects they weren't finding before. Really, just understanding the problem space at the beginning.

And then from there, being able to work with them to scope that out and scope out what does the line look like? How our parts being fed? And what are we integrating with overall?

[00:44:20] JM: Has transfer learning been useful at all to you in getting one model to be essentially reused by another system?

[00:44:29] AB: Not at the moment given our current architecture. I don't believe so. But, yeah, I don't think at the moment we are digging too deeply into that just given kind of the initial architecture of what we're doing.

[00:44:43] JM: You mentioned edge computing a little bit earlier. Do you see any potential to use things like Lambda at the edge, or like Cloudflare workers, or Cloudflare keyvalue stores? Any other edge computing products that might be useful to you?

[00:45:02] AB: I definitely think we can continue to explore edge products. So, yeah, I think there's definitely a bunch of those that could be interesting. And I know on the engineering, our CTO and our team are continuing to look into kind of performance at the edge and different solutions there. Also, Nvidia is putting a bunch of work into enabling obviously great performant edge solutions and scalable solutions. So, I think it's an area that a lot of companies are investing into. And we are totally happy to work with and partner with others on that and make use of what the community is putting out in that space as well.

[00:45:43] JM: Do you work with any human in the loop infrastructure providers, like Scale.ai?

[00:45:50] AB: We don't at the moment.

[00:45:52] JM: Could you see that happening in the future or being useful?

[00:45:56] AB: Possibly. Definitely, always keeping an eye on it. I know scale is on a really exciting trajectory. So, definitely, watching it and thinking through where it could potentially be useful for us either now or in the future.

[00:46:09] JM: Cool. Well, Arye, it's been a real pleasure talking. Do you have any other subjects you want to discuss? Things in Elementary's trajectory that I haven't really thought of to mention?

[00:46:21] AB: No. This is awesome. I really appreciate it. I think I would just mention that if this idea of full stack solutions and full stack problems resonates with you, the idea of deploying machine learning to improve quality and kind of improved yields throughout manufacturing and to bring visual intelligence to manufacturing. If any of those either challenges or workspace really resonate with you and to the listeners out there, please check us out. We'd love to talk to you either over LinkedIn or look at our job board and see what opportunities are out there. We are hiring. We did close a series A recently led by a great VC firm threshold. And so, we are definitely growing and it's an exciting time here. And so just anyone out there, we'd love the opportunity to see if there's a fit on the Elementary side as well.

[00:47:18] JM: Okay, Arye. Well, it's been a real pleasure talking. Thanks for coming on the show.

[00:47:21] AB: Cool. Thank you so much. And I really appreciate it. It's a lot of fun.

[END]