Research @ Citi Podcast, Episode 62: The Rise of Physical AI and What It Means for Us

Recorded: Jan. 9, 2026

Published: Jan. 15, 2026

Host: Elise Badoy, Head of Europe, UK, Middle East and Africa Equity Research, Citi

Guest: Martin Wilkie, Head of Capital Goods Research, Citi

Transcript:

Elise Badoy (0:00)

Welcome to the Research @ Citi podcast. I'm your host Elise Badoy, Head of Europe, UK, the Middle East and Africa Equity Research at Citi. I'm joined today by Martin Wilkie and we're recording in London. Martin is Head of Cap Goods Research here at Citi.

In this episode, we're really excited to dive into physical AI. So Martin, we've heard a lot about physical AI recently. What does it mean and how does it differ from the traditional fields of automation and robotics, which we have extensively written about as well.

Martin Wilkie (0:36)

Yeah, thank you. So the market for physical AI is getting a lot of press at the moment. It was a big subject at the Consumer Electronics Show in Las Vegas just even earlier this week. And the market refers to any physical process where you are learning from and applying artificial intelligence. So in the same way that AI is used in text, in large language models, in the physical world it's where AI meets real-world processes. And this is everything from AI being used in quality inspection for paintings on cars, all the way through to factory process, humanoid robots and autonomous vehicles.

So effectively, it's this adoption of AI into the real world. In the case of what I look at, it's the industrial world. It's a very broad set of applications where AI starts getting involved in areas, as you mentioned, like automation and robotics. So in terms of the industrial world, this is an evolution of adoption of AI that's already started. Of course, like in many industries, AI is used in cases of generative AI. So if we look in things like the service area, a lot of AI has been used already to help service engineers become more efficient.

And even things like agentic AI, where you are making decisions, that's used in things like purchasing and other decision-making in factories. But this is really one step further when you're actually getting the physical processes driven by AI. And that's all the way from design and then from autonomy into use as well. That’s where the big disruptions come from — and, of course, the big opportunities.

In terms of how we think about physical AI, I think we can really split it into two areas. One is around the design process: If you think about designing a car or an aircraft wing, you're really thinking about the very beginning lifecycle of an industrial product. And then, of course, all the way through to autonomy in use.

So in that first example of the design process, we've seen examples in the past where generative design has used AI to think about improving human decision-making on optimizing the design of something. But AI in a physical sense really accelerates that and starts coming up with designs that are optimal in ways that perhaps couldn't be done manually, or certainly not in any reasonable time scale. So it really turbocharges that ability to design better industrial processes, better products.

And then, arguably more famously, it's all about autonomy in use. This is when we start thinking about factories, about warehouses, about transportation, where you have autonomy in use. And you're using sensing, vision and other real-world data to come up with this ability to have autonomy in use. So it's really a massive accelerant for the automation and robotics that you've mentioned and a huge potential driver of productivity once we see this coming into play.

Elise Badoy (3:33)

So is this new? There have been investments in the past in digital twins and in the Internet of Things, IoT. Is physical AI something else? We've seen various iterations, so I'm just trying to make sure we understand the pivot here, especially for investors and our listeners today.

Martin Wilkie (3:51)

It is additional to those trends, which, as you mentioned, have been really a theme for the sector for several years now. But it's not just slightly incremental — it's massively incremental. I think of these digital twins and Internet of Things as the building blocks physical AI is then launching from. So effectively, we have this foundation. And one way we can think about this is, to have physical AI, you need the information, you need the data. And the industrial internet or the IoT, that was really putting connectivity onto everything. Many processes — whether it was your car, or whether it was a factory process several years ago — were effectively analog. There was a piece of machinery that was operating and producing an output, but there was no digital information as part of that, either for gathering data or for making decisions to actually use the product.

And of course the IoT effectively connected all these devices, and so you're suddenly getting data on everything from processes all the way through to transportation, to factories, warehouses and everything else. So in many ways, in the same way that a large language model needs huge amounts of text to base the decisions on, physical AI needs that same input, it needs data to drive its decision-making.

So the IoT really enables this because effectively what it's done is to give vast amounts of data into the physical AI process, and therefore it's one of the building blocks. And similarly for digital twins.

So, for our listeners, a digital twin is effectively a parallel version of a real-world process: Effectively, it's your 3D CAD image, if you like, as to what's happening inside a factory, inside a warehouse, whatever else it may be — a transportation system. It allows you to simulate what's happening in the real world, and effectively it allows you to tweak processes in a way that you can test it in the digital world before you do it in the real world. And this is not new, this has been happening for some time.

The difference here, of course, is that once you have that digital twin, then getting AI to effectively run these simulations to come up with better answers effectively turbocharges the use of those digital twins. And the examples here is that you get real-time learning by physical AI to come up with a better process.

If you think about running a warehouse, where do I store the product? What's the optimal route to gathering products from that warehouse before I deliver it to the customer? In the past that would have been designed by engineers, and in some ways would have been quite rigid. And of course in the physical AI world, it learns in real time. It learns that a certain product is more popular than another, it learns that a certain process will be faster than a different process, and therefore it self-optimizes.

But of course to do that, it needs to have all this technology. So we absolutely see this digitization trend that we've seen for a good 10+ years now really being the foundation, and physical AI is driving that to the next level.

It is worth adding, of course, that there are many, many things physical AI needs that goes beyond what we've had in large language models. And there are, of course, constraints in the real world that go beyond just looking at data. We think about the laws of physics, thermodynamics, forces, all these kinds of things. And so in many cases you need not just the software that we've talked about here but also simulation that understands what the limitations are of running a process, whether that's in transportation or in factories. And integrating these simulation packages with the digital twins, with the Internet of Things, allows you to build this foundation to really drive physical AI. So it's really building on these foundations and massively accelerating the adoption.

Elise Badoy (7:35)

So what does this mean for robotics? It seems it's not just humanoids.

Martin Wilkie (7:43)

No, that's right. And I think humanoids is a very headline-grabbing example, partly because they seem to be so disruptive. And of course, in a sci-fi sense of the word, the humanoid robot is often seen as the natural progression of autonomy and robotics. But in reality, humanoids is just one subset of robotics, and even robotics is one subset of autonomy for physical AI.

Remember, we think about autonomy in many things, whether it's transportation, entire systems, not just in robots. But even within robotics, the humanoid form factor is one example. One thing we've highlighted in the note that we published just at the end of last year is that a lot of what we call task-specific automation is really going to gain investment from here as well. Because we've looked at and talked to companies looking at how factories are operated, and in more than 90% of cases, the work done in factories could effectively be done by a single arm.

So if you think about a human in a factory today that could be sorting or picking or moving or something like that, even though that person obviously has two arms and two legs, what is needed for that specific task is not necessarily the entire functionality of the human body, and therefore why do we need the entire functionality of a humanoid?

So we really see parallel developments here where the humanoid will develop in many applications where you genuinely need to replicate the full range of abilities of a human, but in many other cases, where you just need really a single arm — which, as I mentioned, is very much the case in the vast majority of factories — we see this task-specific automation really being the most dominant use of physical AI.

So we see intelligence getting built onto that. Today, of course, robotic arms are used in many, many areas, whether it's welding or painting or lifting or picking and all these kinds of things. But in many cases, these are relatively dumb processes. And of course adding sensing, adding vision, other sources of data gathering, allows that intelligence to be put onto that robotic arm. And so we do see that you get intelligence being brought into many, many other areas of factories as well.

It is worth highlighting that in many cases, the humanoids can be bettered for specific applications, particularly if it's load-bearing capabilities. There are other ways to lift heavy loads beyond a humanoid. In the case of picking and sortation, there are some applications that are clearly much faster than the human hand. And this is why we really highlight that you get this adoption of physical AI across many, many different form factors, not just in humanoids.

One example we've given is that in the same way that the horse and cart wasn't replaced by a robotic horse, we show that specialized robotic and automation applications will develop to really optimize those specific tasks. And so you're having form factors that are not necessarily humanoid. Of course, in many applications outside of industry — particularly when we think about service and consumer markets and things like that — the humanoid opportunity is substantial. But we do highlight that inside factories, this task-specific automation is arguably just as important.

Elise Badoy (10:55)

And when people talk about cloud vs. edge, what does this mean in the case here for physical AI? Will one be likely to dominate?

Martin Wilkie (11:05)

So this is a big debate that we have with investors. And the reason this has come, I think, to be such a focal point for the debate is around the investments that have happened so far. Because arguably what's happened over the last several years is that everything has moved to the cloud, particularly when we think about the IoT — the Internet of Things — that has moved data and design decisions to the cloud. And there were huge reasons behind that. And of course we see this even outside of industrial applications, that moving to the cloud has huge advantages in scalability, in ease of deployment, in security, in cost. All these reasons have driven the cloud adoption, not just in enterprise software but also in industrial vertical software and in particular on the IoT.

But there are limitations with the cloud. Some of these limitations can be latency. This is literally just time — the data doesn't travel instantly from one place to the other. Of course cyber concerns: Do you want your highly sensitive factory processes to be adopted or taken outside of your premises? And in many cases linked to that data, sovereignty requirements — in some industries, legally you have to keep the data inside a certain physical area, and so you can't necessarily have it on the cloud that could be based somewhere else.

And so it does matter that we think about this edge vs. cloud debate. We do see both really mattering. In industrial terms, edge devices are those devices that are at the edge of the network, which is a similar terminology that you have in consumer and other markets. But when we think about industry, edge devices are the industrial-hardware components that can gather real-world data to be digitized. So this is sensors and switches and things like that. But it's also the items that are then running autonomously, whether it's your car, whether it's a robot.

So these edge devices are both gathering data and they're actually in use. And of course, in reality, you need to get both sides of the equation right. You need to get the latency and data-sovereignty questions right as well. And so we absolutely see both of these happening.

I think a lot of the design process where arguably latency is less of an issue is going to remain on the cloud. And so we continue to see a lot of the software movements toward generative design areas like that remaining very much cloud-based. Arguably, that also spurs adoption. Many of these business models moving toward software as a service, for example, really benefit from the cloud because it gives you massive scale benefits. It reduces the cost, which of course increases the adoption rate.

And so on that side, huge amounts of it still will happen on the cloud, but equally at the edge there are significant advantages, particularly with latency. One example often given here is with an autonomous car. If the autonomous car has a latency issue when it's running, in the real world driving down the motorway you absolutely can't afford to have any latency of a decision between a real-world action and the car deciding to stop or start or whatever else it might have to do. And it's the same in factories: Do you really want to have the risk of having any slight slowdown in the data? And for that reason, we see both of these things getting adopted.

So we don't think that physical AI is going to suddenly move all to the edge and that's going to mean that data-center investment and things like that are going to somehow slow down or not be needed. We absolutely see both of these happening. And it does mean that you have these two parallel sets of investments on the cloud, which is going to be obviously dominated by large data centers, and then at the edge, where you very much get infrastructure and intelligence built inside industrial facilities, whether that's in warehouses or in factories.

Elise Badoy (14:52)

When you look at the success factors, what will be critical? How will competitive markets be built?

Martin Wilkie (15:00)

So we see three key pillars of success for industrial companies here. The ability to create digital twins really matters because to optimize an industrial process, you have to effectively simulate that before you can think about applying it in the real world. And so really having the understanding of how the real world operates — whether that's the layout to the factory, whether it's your transportation network, whether it's the road network or whatever else it might be — you need to have information in all of that.

So in the same way that large language models need massive volumes of text to learn from, industrial AI needs industrial data. And this really falls into two camps. You've got real-world data directly from intelligent devices or from asset-management software. So this is effectively that you have sensors, you have all sorts of cameras, whatever else it might be inside factories and warehouses that's gathering information. And that's giving you real-world information and visibility about what's happening that you can digitize.

And equally, when we think of asset-management software, these are effectively the software packages that are running an oil refinery or running another industrial process. And having access to all that data, regardless of how good your model is inside AI, if you don't have that data then effectively your model is worthless.

And so really having access to that data gives you this enormous advantage. And of course, the install bases of many industrial companies is that they have vast amounts of on-the-ground equipment that is gathering information already. So we absolutely see one of the big critical success factors as being access to this data and having the domain knowledge of understanding what it means. But of course the data in one industry can be very different to another industry, and understanding how a utility works, how an oil company works, how a car company works is very, very different. And therefore having that data and the domain know-how really, really matters.

In terms of how that's then used and developed from a business-model perspective to create these moats, we do see opportunities both in the cloud, as we mentioned — a lot of that is in design, in training and simulation — and at the edge, which is really around autonomy, quality and prediction. And while AI is both potentially a disruptor and an accelerant to the adoption, we do think that, again, this data is arguably the biggest part of the moat.

And what you are seeing is this is driving acquisitions. You're seeing companies that are perhaps wanting to get much more onto the data side of things making acquisitions in software, or they're making acquisitions in companies that are gathering the information, whether that's in sensing or elsewhere. And so really it's building this mode to have the ability to get the data and to process the data.

So we've often said that even for software, the industrial install base and the domain expertise becomes the biggest part of the mode. And we very much think that this continuation of ensuring that you understand the data and you understand the application of that data is really going to give companies the ability to create what I think can be very, very disruptive processes in physical AI.

Elise Badoy (18:17)

And then finally, Martin, many regions of the world have been in industrial recession. When do we see this investment actually happening?

Martin Wilkie (18:25)

So we are, I think, at the very beginning of this journey, and that's partly driven by the technology. We're just at the beginning of the technology becoming available. But then, as you say, we've been in an industrial recession. So part of this is also that the industrial cycle has arguably been sort of against us for the last two or three years.

We have seen some of this being adopted already, so if we look back about five or so years, you were beginning to see already some very small cases of AI being used in manufacturing, and that was typically predictive maintenance. This was recognizing certain patterns — that if there was a certain vibration or noise or whatever else it might be, the AI could recognize that this was predictive that an engine was about to break down, or something like that. So we've begun to see small applications like that already. We're now beginning to see some of those areas moving across into things like quality control. And so many car companies now are using AI already to detect whether the paint finish on the car is at an acceptable level, and typically the success rate, if you like, of AI doing that vs. the human eye is already better.

So we've seen this beginning to get adopted already, but we really, really are at the very beginning of this cycle. And what we highlight is that industrial-technology investments are usually pro-cyclical. What that means is that if I'm a capital allocator inside an industrial company, I typically don't make the investment in a downturn, even if the ROI is favorable on paper, because arguably the ROI is even better if I wait until an upturn.

Many of the drivers of physical AI are there all the time — labor costs, labor availability, supply-chain challenges, all these kinds of things — but, arguably, they become even more pertinent when you move into an upcycle because labor costs get even higher, labor availability becomes more challenging. Oftentimes, the inflation of raw materials and so forth go up. And of course, your desire then to become even more efficient becomes much greater in an upcycle.

So moving into that industrial upcycle I think is going to really matter. And that's not been where we've been for the last two or three years. I mean, post-COVID we've had an industrial recession for different reasons in different parts of the world, and industrial production has, on balance, been in a recession from 2023 up until about the middle of 2025. So there is a cyclical part to this that we think is beginning to inflect.

But just as importantly, we think that the technology is now becoming available. And we begin to see this happening already. I think we can already look at some of the use cases as to how this can show itself into the numbers.

We've looked at things like the industrial-robot market. Most of the third-party forecasts are that this market can be growing at a mid- to high-single-digit percentage rate, which is already somewhat attractive, but it's not exponential or anything like that. But one of the things we show in the note that we published back at the end of last year was if you begin to get industrial robots that can take more and more tasks away from humans — and a lot of that is because there is no available labor and effectively you're backfilling for the lack of availability of humans — these robots of course are much more efficient, they can run multiple shifts per day, they can often be more productive. And the adoption of this can really, I think, massively accelerate that adoption and potentially change the growth rate of that industry from being mid- to high-single-digit to potentially being 20%+ CAGR over the next several years.

And so we are really at the beginning of that process. And in many cases, we think this can be additive to capex budgets for companies over the next few years. In many ways, simply put, you're effectively changing your wage bill into a depreciation bill. You're effectively moving away from running processes by people to running processes through autonomous software and design or robots and so forth in factories. This can really be additive to capex growth, and we've shown that we think this can add a mid-single-digit percentage to annual capex growth initially over the next five to 10 years.

And so this is really giving us, I think, what can be a very, very appealing upcycle in the adoption of physical AI. And of course, it cuts across many, many different areas, whether it's in robotics, where it's in design simulation. But this is where we see this growth picking up. As I mentioned, it's linked to the capex cycle and the industrial-production cycle, and therefore we still think we are very much at the beginning of this journey.

Elise Badoy (23:13)

Martin, thanks so much for this fascinating overview. This episode of Research @ Citi was recorded on Friday, Jan. 9, 2026, and I'm your host, Elise Badoy.

Disclaimer (0:00)