Research @ Citi: AI’s Race for Real-World Data

Recorded: Aug. 21, 2026

Published: Aug. 27, 2026

Host: Elise Badoy, Head of EMEA Equity Research

Guest: Martin Wilkie, Co-Head of the Industrial Tech & Mobility Super-Sector

Transcript:

Opening Teaser (00:00)

Research @ Citi.

Elise Badoy (00:03)

Welcome to the Research @ Citi podcast. I'm your host, Elise Badoy. I'm Head of EMEA Equity Research at Citi.

In just a few days, Citi will host its landmark Global TMT Conference in New York. I think it's the largest ever we've hosted, bringing together leading executives, investors and industry experts from across the technology, media and telecommunications landscape. Ahead of the conference, we're discussing here on this podcast — and we're recording live from London — one of the most exciting themes shaping the future of technology and industry: physical AI.

Physical AI really is the convergence of artificial intelligence with the physical world, and it will enable machines, robots, and autonomous systems to perceive, reason, and act in real-world environments. As advances in AI increasingly move beyond the digital realm and into factories, warehouses, transportation networks, and indeed everyday life, investors are continuously asking us what this means for companies, for supply chains, and the broader economy.

To help us take listeners on that journey, I'm delighted to be joined by Martin Wilkie, who's the Global Co-Head of Industrial Research at Citi.

So, Martin, what is physical AI and the broad sectors that it affects?

Martin Wilkie (1:20)

Yes, thank you.

So physical AI is one subset of AI that we see really accelerating in the coming years because the technology is now allowing it to evolve. It is normally thought of as separate to agentic AI and generative AI, because physical AI is really all about the interaction with the physical world — using models to make things more efficient and, in some cases, autonomous as well.

Some of the examples of physical AI that probably come to mind immediately are things like autonomous cars, humanoid robots, drones. All of these are sort of obvious real-world examples of physical AI. But it's really anything where AI and physical systems are combined. This could be a self-learning warehouse to optimize logistics. And so all kinds of processes that link the physical world and use AI is sort of what we think about here.

Humanoids are a very important and very visible subset of physical AI, but in reality we're bringing AI to make autonomous physical processes that look nothing like humanoids in many cases. So this can really span across many, many industries. And one feature of it is that it can be quite industry-specific, which is unlike some of the large-language-model developments where you saw improvements in models that could be applied across industries very, very quickly. But of course, in the physical world, everything is very, very different.

When you think about autonomous cars, when you think about a factory, when you think about agricultural equipment, all of these examples are very, very different. And the solutions to them will evolve at different paces and they learn in very, very different ways because of the data that is gathered.

So physical AI is in many cases very industry-specific, even though the building blocks are similar, but how it's then applied to each industry will change over time because the availability of the data, the ability to integrate it is quite different by industry. But you mentioned which industries we will see this in: We'll see it in energy systems, we'll see it in transportation systems obviously, we'll see it in production, logistics. All these different areas, while having very specific form factors of the physical equipment, but they will all need to integrate with AI to really evolve.

When we think about what physical AI means, a lot of consultants and experts use this concept of vision, language and action. And what they mean by that is that these AI models have to sense the real world.

So the vision doesn't necessarily have to be a line of sight — it could be sensing in other areas, whether it's temperature or pressure or these kinds of things. But you have to get that real-world data and you feed that into a model. That's the language model. And then the action is what these systems do — so it's to move the robotic arm, it's to drive the autonomous car. And in some ways, that's really the difference of physical AI with a generative AI model, because you've got that real-world sensing up front. So you gather the data and you've got the real-world action at the end of it. You're physically doing something off the back of that.

And that really describes how physical AI would differ from a more typical description of generative AI. It does require a whole set of new interactions between AI models and the real world. And we'll come and talk about that in a moment. But ultimately, that's what we're trying to get to here — to allow AI to interact with the physical world.

Elise Badoy (4:41)

But so, Martin, obviously we've heard about the different challenges that we find with AI. But can you just explain to us what the different challenges are with specifically physical AI, perhaps compared to generative AI or simply large language models?

Martin Wilkie (4:52)

I think the biggest difference is how these models will be trained and how they understand real-world data.

When you look at how some of the early AI models were developed on text and images, oftentimes there were vast amounts of text available. And so these models — whether it's on the internet, on social media, or existing libraries of text — there was an unending amount of data that you could train a generative model on. Similarly, when it came to images, a lot of these were augmented in the training process by human verifications. So you had a model learning: Does this image contain a picture of a human, for example, in the case of an autonomous car, understanding its environment. But those pictures were verified by humans and the models would learn in that way.

When it comes to the physical world, particularly industrial processes, you don't have access to that data. Because in many cases these systems in the physical world weren't digitized historically. Many are today, so the Internet of Things has developed over really the last 10 or 15 years, so more and more processes are digitized. You can see this in your own home with your heating system, or things like that, where these are digital processes that might not have been, say, 10 years ago.

So there is more data available, but it's not nearly as much as you've got in other areas. And so getting access to that data is a constraint and it's something that is difficult to do. There's also the challenge of who the data is made available to, because in many cases a lot of this data is siloed. So for a factory owner, whether it's for security reasons or just because of the way the system was developed, that data is within that factory but it's not necessarily available to somebody else. And therefore, to train these models becomes very, very challenging.

And so one of the big differences is: How do you train these models? How can an AI system understand what an image of a certain thing is doing, or what they're reading from a temperature sensor or a pressure sensor or whatever else it might be. What does that actually mean? And so getting the volume of data is a huge challenge.

Another part — and this has become increasingly obvious over the last couple of years, and we've seen a lot of companies invest very heavily in this area, both organically and through acquisitions — is what we call data contextualization. And what this really means is: What is the context of that data? If I sense that something is moving at a certain speed, or that the pressure of something goes from X to Y, what does that actually mean in the real world? And if I have that measurement in one place, does it mean the same in another place?

We as humans would understand this naturally — you know, if I feel the temperature is hotter regardless of whether I'm in one country or the other, as a human being, the human brain understands how to contextualize what it's sensing. These models need to be taught how to do that. And so there are huge amounts of challenge around the availability of the data and then how to contextualize it. So I would say that this is probably the largest single challenge that physical AI has — how do you get the information on which to train these models?

Luckily, there is a parallel way of doing this which is more unique to physical AI than it is to generative language models, in that the real world has to abide by the laws of physics — whether that's thermodynamics or whether it's forces. And so your model can be augmented with laws of physics. You know that if a force interacts with another force, something happens — you don't need to have observed it because that can be codified. Similarly on things like heat and pressure — you know these things follow laws.

And so the good thing is that even though a lot of data is not necessarily immediately available, we know that there are lots of ways that the data has to follow certain patterns because of the laws of physics. So again, that's quite a big difference between how we think about some of the other areas of AI. And I do think that data availability and how you contextualize and use it is by far the biggest difference.

There's a second one, which arguably cuts across all types of AI, but I do think it's probably more important and more critical in physical AI. And that's the guardrails around it, given we're looking to automate and to bring intelligence to physical processes that clearly can have quite big impacts on the world around it. And the obvious example here would be an autonomous vehicle: How do you make sure it's safe? Who signs off? If it goes wrong, where's the liability? And so there are big questions around how that develops.

If we think about the early versions of physical AI, a lot of these were done in tandem with the industrial process. A lot of this was done on things like quality control — could I use AI to check the paint finish on a car, or some other measure of quality? And so it was a parallel process to what was actually being built or operated, but it wasn't inherently controlling the process. And we saw this in predictive maintenance for service, we saw it in quality control as I mentioned — areas like that.

But now, of course, we're moving into the possibility for physical AI to control the process, or in some cases even control the design. And therefore we do have to think about the guardrails around this, and in particular: Does there need to be a human in the loop? How do we think about governance? How do we think about safety? Where does the liability arise? If you have to get accreditation, whether it's for safety or for efficiency or whatever else it might be, who signs that off?

And particularly when it comes to autonomous processes or autonomous design, I think that's going to be a really, really important part of the story. A lot of that is regulatory as opposed to a technical development. And I think that in some ways that will drive, and in some ways may constrain, the growth here — because you are going to get requirements to make sure that these things are safe and that they can operate and coexist with humans.

And of course I don't just mean that in the sense of autonomous cars, but also just in terms of running a factory or running an energy system. Does the human operator ultimately have control, as opposed to it being fully autonomous? So some quite big differences there, and I would say those are the two big different challenges of physical AI compared to what we see in generative and other parts of artificial intelligence.

Elise Badoy (10:58)

I think we've just looked at the basics here, but what key trends should we be looking out for in the next few months, in the next few years?

Martin Wilkie (11:09)

I do think there are a couple of areas that are going to really drive where the winners emerge in physical AI. And a lot of this is around getting the data — so data is going to be increasingly valuable, understanding it, contextualizing it as I mentioned, and making sure that you can then use that to optimize your process.

And I think what you're going to see — we know that data is valuable already, but how you translate that and how you build your AI models alongside it — I think there's going to be increased recognition of the value of an installed base. I think when you look at a lot of generative models, they can be done on public sources, particularly for text and images and things like that. That's not going to be the same case in industry. As I mentioned, a lot of this is very industry-specific. So if you have a lot of learnings in the energy industry, those are not translatable at all really into transport or into manufacturing or whatever else it might be. And so the value in a particular end market in terms of the install base of gathering data is going to be very, very important. And that's going to ultimately drive the winners, I think.

I think what you are going to see is increased investment in effectively land grabs in install base — can I make sure that I have access to what is effectively proprietary data to train these models? And do I invest in the software and the systems to be able to contextualize that data? That's going to be really important. In parallel with that, simulation and physics-based models will fill the gap.

And so, as I mentioned, because there's a scarcity of data in many of these markets, but you can use physics-based models to plug that gap, then interacting between this really valuable install base of actual data, but then combining that with simulation and physics-based models is going to allow these physical AI systems to really develop. And so a lot of this is going to bring it all together and really separate the winners from the losers.

I think what you are going to see is the need to have this intelligence as an important layer on top of physical devices. So clearly, since the Industrial Revolution, companies could make phenomenal returns on producing physical equipment, whether it's in transport or manufacturing or whatever it might be. But more and more of the optimization around that is going to be in the data and understanding how to use that data.

And so I think you are going to see these two things come together — so that you can optimize the overall system, you're not just having the physical product but also the data. For me that's a huge driver — and really this convergence with physical assets and data — and I think that's something that's going to be a really important platform.

The second trend I think is going to be around edge data centers. Now, of course, another big theme that's been happening in the industrial sector is all the kit that is going into supplying data centers — all the investment that's needed for cooling and power and all these things. And so far, huge amounts of that have been on these ever-larger data centers. And so effectively the more complex models need more power to consume, and this has led to higher density of data centers, more power consumption.

There are challenges around that, particularly if you're running a physical system. Latency, which is simply how long it takes for the data to get from one place to the next, becomes a real challenge if you're running a physical system that is very fast-paced and potentially a very hazardous environment. And you can think of the obvious example: If I've got an autonomous car and it's trying to make a decision and it's transmitting information to a data center that's several hundred or several thousand kilometers away, and you have a challenge with connectivity, then you simply can't have that. And that's true in so many cases.

And even in places where you might not think that they're so hazardous, the amount of decision-making that is happening — even if you think about high-speed manufacturing processes — even just across a few seconds means that you can't really afford to have these latency issues.

When you combine that with the concerns over data sovereignty — so who owns the data — and also cyber security — so is it safe — means that this concept of edge data centers I think is going to become more and more important.

So effectively, you'd end up with this barbell of larger data centers that can be used for training and for calculations that are not latency-constrained, but also edge data centers in factories, in transportation hubs — it could be all sorts of places. And I think that's really what you're going to see, in that you need all of these systems to have immediate access to decision-making. Effectively, the brain needs to be where the physical devices are.

And so I think those are the two trends I look out for — on the data side, but also then on edge data centers.

Elise Badoy (15:59)

So we've looked a lot at the blue sky, Martin, but obviously as investors know in particular, they will be looking for challenges and there are suddenly some things along the way. Do you want to talk to us about the future challenges and various risks to the execution of the blue-sky scenario that we've discussed previously on this podcast?

Martin Wilkie (16:19)

So there are many, many challenges, and many of these we don't have the answer to yet. And some of them are political as opposed to technical, and so are probably more difficult to predict in terms of how and when they'll be resolved.

We talked about some of the technical challenges already when you touched on cyber risk. I do think a lot of the other ones that come into this really become more important when the human is no longer in the loop.

Now, I would argue that in the short run, that's actually not going to happen. In the very immediate term, many, many of these decisions remain under the control of the human. So you're effectively having agents doing individual tasks, but the human is ultimately still in the loop. You're not having a fully autonomous factory. You're not having a fully autonomous vehicle. You're not having fully autonomous design. And so that human in the loop remains important.

To move beyond that, there are big, big questions in terms of liability — you know, what happens if it goes wrong? — but also in terms of who signs it off. And in many, many industries, whether it's in something as obvious as aerospace, but even when it comes to construction, the number of safety certificates and the accreditation you've got to go through to sign off these products is absolutely immense. And so the argument that you can take the human out of the loop I think is quite challenging.

So that's going to be one key part to it — how do we think about that human in the loop part? And will that change over time?

The second one is really around the political acceptance around the workforce. And I think there's a lot of fear around this. There is a parallel argument, of course, that demographics are arguably very supportive, given aging populations and peak workforce already being reached, or in some cases being passed, in certain countries. And so there are areas, whether it's in Japan or China, where there's a recognition that even if it's not necessarily for society as a whole, there are certain areas that are understaffed and the demographics means that's not going to get any easier. And therefore physical AI, whether it's robots, in services, whether it's in factories, whether it's in transport, will simply be more needed.

So I do think on the political side there is this dual concern. One is really around: Do robots take your job? But the second one is: Well, actually, in some areas, we need to have it precisely because of the changing demographics. But as we know, these things can become quite challenging to debate and can raise a lot of concerns. So I think political acceptance is going to really matter.

And I think that's why it does matter for companies that supply into these areas — whether it's analyzing the data or providing the physical systems — that they're very cognizant of that sort of political risk around it. So definitely, that's going to be a big part of the debate.

Another one is cost — and I don't mean this necessarily in the dollar amounts; I mean more in terms of the structure of the cost. Because in many countries wages are relatively flexible: It's not true globally, and there are some areas where the wage structure is much more fixed than in other countries, where the wage structure can be a lot more flexible.

But ultimately, if you are putting in place physical AI, that's a lot of upfront capex. And so effectively you're moving from an opex-based P&L to a capex- and depreciation-based P&L, with a lot of upfront investment. And the challenge you have with that is: What happens if you get it wrong in terms of your assessment of the cycle, or just in terms of the growth of the market?

So if you're investing very, very heavily in physical AI systems — whether it's a transport network or whether it's a factory — and it turns out that demand is not as strong as you think, you've already made a lot of that investment. And so how can you get comfortable with the capital employed that goes alongside it?

Debates will come around as part of it. Can it be modular? So can you effectively scale this investment as the demand increases, as opposed to building it all upfront? Can it be leased? Or can you have physical AI as a service? And if that's the case, whose balance sheet does the investment sit on?

So I do think that, in the same way that it's true for lots of other investment in AI, this does require more capex. And ultimately this can be incredibly efficient, but it does mean that there's going to be more upfront investment. And the question around who finances it and whose balance sheet it sits on is going to be an important area. I don't see that as being a hot-topic issue today because we're really at the beginning of this process. But I think over time, that's going to become a much, much more pertinent question as well.

I think the final one — linking to that — is that the investment is likely to be pro-cyclical. We've talked about physical AI on these podcasts before, and one of the challenges we've had in the past is that when certain countries or sectors are in recession, they don't make these investments. Because effectively, where's the demand? What's the point of making the investment?

And so one of the challenges you have is it's not just whether the technology is available, but is the sector that you're selling into moving into an upcycle? Do they see the value of becoming more efficient, cutting costs, and all these kinds of things? And therefore, timing it is not just about whether the technology is available, but whether the end cycle in that sector is moving into a favorable dynamic as well. So that's going to be challenging as well in terms of seeing when the adoption picks up.

So yes, many, many challenges ahead, but I do think a lot of exciting developments as well. And even just in the past 12 months or so, I think there's been a lot of movement — not just in the quality of the models, but even some real-world examples where we're beginning to see increased penetration of physical AI. And I think this is going to be a hotter and hotter topic in the years to come.

Elise Badoy (21:45)

Martin, thank you very much for taking us on this fascinating journey. This episode of Research @ Citi was recorded on Friday the 21st of August 2026.

Disclaimer (21:57)