

Sustainability Now Podcast

Can AI Really Extract Emissions Data?

Transcript: 21 August 2026

Bentley Kaplan

Hello and welcome to the weekly edition of Sustainability Now, the show that explores how the environment, our society and corporate governance affects and are affected by our economy. I'm Bentley Kaplan and very glad to be your host for this episode. On today's show, we're going to answer a burning question that climate investors have been asking themselves, their research providers and their new best friend, Claude, which is: can AI collect emissions data for me? Because it's fast, very fast, pretty cheap compared with a human, and mostly accurate. So in today's show, we're going to corner one of MSCI's foremost experts on carbon emissions to ask what can go right when you unleash an AI agent on emissions data, and what can go wrong. Thanks for sticking around. Let's do this.

Large language models and the arrival of generative AI has been a seismic event. Conversations between philosophers, thought leaders, politicians and tech CEOs about how this new technology is going to disrupt, upend, improve, destroy and revolutionize multiple industries and civilization itself are definitely worth a listen to. But between you and me, they can be a little bit hard to process. A disbelieving head shake might be the most rational response as you get back to work.

But as you would have noticed, and as your friends, parents, partners, children and pets would have told you, AI agents are already bleeding into the smaller corners of our lives, taking over or helping with the tasks that many of us do daily. In some cases it's repetitive work, the grind, simple tasks that need to be repeated thousands or tens of thousands of times, and leaning on an AI agent saves us that drudge. It also can speed the process up almost inconceivably and redirect that time, effort or salary to other endeavors.

For investors specifically, many would be looking around at the tasks that they do or pay others to do, trying to figure out how much of it can be offloaded onto an AI agent. It's an option that can be really enticing at a time where budgets are under pressure, fees are being compressed on the revenue side, and costs for market data aren't falling at the same time. An ideal candidate would be something like data collection, and for scenarios where, firstly, the data collected are pretty simple but needed at scale and, secondly, where that data isn't necessarily driving returns but it's often needed for regulatory purposes or to meet client mandates. Which is what leads us to carbon emissions: scope one, two and three. The basic building blocks of so many climate-related investment strategies and reporting requirements. Companies are putting these numbers out there on a more regular basis. And sure, scope three reporting isn't great yet, but AI is great at plugging gaps.

Now, somebody who might have more mixed feelings about turning AI loose onto the collection of emissions data is Tanguy Séné, my colleague based in MSCI's Paris office. Tanguy spends a lot of time

and brainpower thinking about scope one, two and three emissions. But don't feel bad for him. He loves it, trust me. And getting this task away from him and into the arms of an AI agent was never going to be easy, but I thought it was worth a shot anyway.

Hey Tanguy. So welcome to what I think is your first appearance on the podcast. Um, and I know this is something that you've been thinking about. It's definitely something that investors have been thinking about as well, you know, about replacing human researchers with buffed-up AI agents, faster, cheaper emissions data, you know? So given your position and the way you've worked with emissions data over the past several years, is there something you would push back on as a viable strategy?

Tanguy Séné

Well, I think it's completely accurate. AI can easily get, and fast, scope one, scope two, scope three emissions data when these are disclosed by companies, of course. It's so fast and it's so good for incorporating data collection learning into processes. We use AI for precisely that. We've been applying machine learning and natural language processing to data collection for years. And by that I mean way before 2023 and the LLM boom. As a side note, I should say machine learning and NLP are a form of AI, just not the generative kind. And they are fit for purpose for data collection.

That being said, a couple of examples. In extraction we use NLP and machine learning to identify the entity, which issuer, which security, which subsidiary, and then to detect and extract facts from reported documents, tables, charts, numbers across various sources and languages. We use machine learning in quality assurance to monitor the data continuously and flag anomalies. So for instance, implausible for the company, inconsistent with peer group, inconsistent emissions data with related data points. We use AI in geospatial data collection, so asset-level satellite data interpretation. We use machine learning for contextualization, identifying mislabeled data, separating a company's claim from the reported figures. So AI, the way we use it at MSCI, is at every step in data processes. And in addition, what it automates is learning that we've accumulated over years in this process. So all the experience we build in data collection, we can integrate it into AI-built processes towards more and more automation. But we need a human in-the-loop to sign off on certain automated recommendations.

Bentley Kaplan

Okay, so that's maybe more of a cordial start than I thought we would get to. You know, it's not, you know, as you're saying, it's not that AI doesn't have a place in our team's workflow. In fact, it's proving very useful, um, in the right kind of applications. But, but I do think those applications you're mentioning, you know, they feel like roles that are kind of maybe nipping around the edges, you know, quality control, error checks, maybe roles that don't challenge the central position of human researchers in the process, you know? So if you're a cash-pressured investor, you're looking for spots where you can get more expensive humans out of a process, where you can bring down costs, you know, with negligible, that's open to argument, negligible impact on your outputs.



So let's say that I am a cash-strapped asset manager, for example, and I still need my emissions data. So what's to stop me from bringing in an AI agent, pointing it at a list of one hundred and fifty companies, say, in my portfolio, and asking it to pull out scope one, two and three emissions for, like, the last fiscal year based on what these companies have put in their reports? You know, maybe I'm going to have to run that prompt a couple of times and spend some time checking a few numbers. And if there's some obvious outliers, doing some, some manual review. But in a few hours, you know, there it is, a great little Excel file with numbers for all of my companies. At face value at least, that seems pretty good to me.

Tanguy Séné

Allow me to take a step back. Is it really what you want? Scope one, scope two, scope three in general? The problem might be more complex than it appears. The whole point of a category such as emissions scope one is comparison, right? So categories exist so that you don't have to compare apples with oranges, or scope one with scope two. But maybe you need a more granular, complex categorization. So to give you an example, to bake an apple pie I need apples rather than oranges, definitely. But to bake a French apple pie, I want golden delicious apples specifically, because lady apples are too small to fit the purpose of how a French apple pie is meant to look.

So that's why I'm asking you about the level of complexity that is best for your objective. In your scope two emissions that you want to collect, are you certain you want to mix market-based, carbon-neutral-adjusted, in the same column as location-based? Combine these numbers and they mean completely different things within scope two from an accounting standpoint, and especially if big tech companies are in your portfolio.

Bentley Kaplan

Okay. So so you're saying, um, you know, fine, unleash your AI agent, go get emissions data. But what you're also saying is that your AI agent doesn't necessarily have the context that needs to know, you know, about what emissions data to collect, or how something like scope two emissions can be reported differently, even by the same company.

Tanguy Séné

Exactly, expert context is everything. That's more than 200 interpretation decisions in our data methodology documentation. And more to come with new standard climate disclosures. Think carbon credits for instance.

Bentley Kaplan

Can we spend a little bit more time here? So let's, let's go back to that original example, you know, of me sitting in front of my AI-collected data sheet with emissions for one hundred and fifty companies. You

know, what are the, what are the first things I should think about before I go and use these numbers in my processes or reporting?

Tanguy Séné

So here you face two potential situations. One is the figures are correct and comparable. The other situation is the figures are correct but not comparable. Taking two examples: approximately thirty six percent of annual scope one and scope two disclosures do not specify whether that figure covers the entire company, scope one and scope two, or only selected activities. So that's a huge proportion of your companies. Second example, in roughly eight point five percent of emissions disclosures we looked at, the combined scope one and two figures, the sum does not equal the addition of the separate scope one and scope two. Why? Because companies may report the total scope one plus two on a different accounting basis than the other parts. So one company, one of your two companies, might be reporting its full global footprint, and the other one reporting its scope one plus two only for its European or US operation. And so then the gap is not a difference in emission performance. It's a difference in the accounting method used. You cannot compare these two numbers apple to apple.

Bentley Kaplan

Mhm. Okay. So I mean I take your point here. If you set your AI loose on the world and it does actually do a super accurate job and pulls in emissions numbers verbatim that companies report across all three scopes, you're saying that the reporting from those companies may not actually tell you if those numbers cover all their operations. For example, if not all operations, then which parts? And second, you know, the type of calculation method companies are using, or what assumptions that they make.

And I think, you know, we've discussed this before, um, a great example of, of this challenge is in, you know, big tech, quote unquote, companies with their scope two emissions. And, you know, for listeners maybe who aren't engaged on this, what is quite an intense debate about something called market-based emissions, um, in short, it's a question about whether a company's emissions should reflect the electricity it's actually pulling off its local grid, so it's location-based emissions, or if it's allowed to adjust that figure based on clean power that it's paid for on paper, you know, its market-based emissions. Because those two measurements of scope two can be a long way apart. A data center running around the clock in a region where most power is either gas or coal based is physically responsible for a lot of emissions. But if the company running that data center then goes and buys enough renewable energy certificates, its market-based emissions can come out close to zero, you know, for argument's sake. So both of those numbers are legitimate depending on your perspective or the framework that you're working from. They're just answering different questions. And with AI pushing data center demand up sharply, the gap between those two numbers is widening, which is exactly why the argument has become so heated.

And I think, you know, to bring it back, you know, that's the risk that you're highlighting here, right? At face value the numbers are accurate. But there is context that's missing. And that context is, you know, is really meaningful.

Tanguy Séné

This is an important, important point. The missing information here is not hidden somewhere in the report. So coverage of operations, or different accounting bases. The missing information you need for comparability was never written down. There is nothing to extract. No human nor AI can retrieve a piece of information that does not exist.

Bentley Kaplan

Okay, so let me, let me play devil's advocate on you here, because I think many people would make the case that, yes, this challenge is here. AI would have difficulty dealing with, you know, these gaps or, you know, these assumptions. But that's not something that we don't already face as human beings collecting or processing this emissions data. We're looking at the same reports as this AI agent, with the same lack of information about calculation methods or operational coverage, or even just the lack of reported emissions in some cases. You know, we have standardized frameworks like the GHG Protocol which have helped in getting more comparable data reported, but it's still a bit of a wild world out there in terms of, you know, what and how companies report, especially if we look across different regions and different company sizes. So why couldn't you just give an AI agent a list of, say, metadata it should collect, including things like, you know, operational coverage not indicated, alongside the actual emissions number? Then you go ahead and build the same emissions data set that a human team can put together, but, you know, much, much faster.

Tanguy Séné

So yes, you can ask AI to collect all kinds of metadata you want, assuming you know exactly what you want and unintended consequences about the data. Ask AI for operational coverage, for scope two method. And if the model is hallucination free, it will probably return an accurate answer very fast. But a lot of the time that answer is not stated. So there's always a judgment to be made about missing data. So what do you do then? So now you have to decide, and who is going to make the decision. That's the important point. There is no answer to that in the company report. So what interpretation decision are you going to make in the face of data gaps? And it has to be the same decision, whatever it is, apply the same way for every issuer every year. And this decision has to make sense for your portfolio, investment processes, clients, auditors. So you need to own these judgment calls. But the difference with the human experts or data provider is this judgment is made potentially silently. You do not own, or you might not see, the judgment that is made by classifying unlabeled scope two as the same thing or comparable thing to market-based, location-based. So this is the hard problem that AI cannot solve on its own. Extraction is getting cheaper, but decisions in the face of missing data still have to be owned by someone.

Bentley Kaplan

That's, that's, I suppose, what's really at the center of this question. It's not just about getting numbers from company reports and making a few minor tweaks. It's about years of expertise, understanding what the different numbers might mean, how they're derived, what they measure and, critically, how they're being used, and being able to transparently sort of show and explain the process for getting to those numbers. And if you try and fold all of that into an AI-based process, it sounds promising, right? But you risk losing transparency and control over decisions that experts know are meaningful but generalists or AI models might not. So very much a user warning on that one. You know, maybe to sort of wrap this up, with AI helping in some of your processes there, what opportunities has it given you through having that extra bandwidth that's been freed up from the more sort of humdrum verification, quality assurance and so on?

Tanguy Séné

From what I see among my data collection colleagues, there is an exciting focus away from data extraction to data interpretability and data structure. And that's because of AI. Instead of working on the tedious process of how to extract the data at scale, eleven thousand plus companies, etc., more and more the focus is about determining what a figure means, determining whether to believe it, against which context, and designing a consistent, practical methodology about data gaps. Part of the job that is still relevant for humans asking better questions. So AI handles the volume. Experts handle new ambiguous cases, and the judgment behind them. And an AI answer is only as good as the human question behind it. And formulating the right questions requires knowing the domain well. So experience, knowing how modeling works, all the frustration about the inconsistency, all the dialogue with corporates. And that's not something you can prompt your way into.

Bentley Kaplan

So let's try and tie all of this up in a nice bow, especially if you find yourself still thinking about French apple pie, which is delicious, but not the point of this episode. So, for free, here is a mental Post-it note with three bullet points.

Number one: for AI, getting emissions numbers was never the hard part. Extraction has been solved. But what AI can't always do is tell you whether two companies' numbers mean the same thing, because the information that would tell you that isn't always reported. To quote Tanguy, there's nothing to extract.

Two: when the data is missing, somebody has to make a call. And that's true whether that somebody is a person or an AI agent. The difference is whether you can see the call being made, and whether the same call gets made the same way next year, and for all companies.

And three: when the inputs don't line up, a model doesn't hand you an obviously bad answer. It hands you a confident one. And that's the risk. Not AI failing loudly, but AI failing quietly and precisely.

And that is it for the week. A massive thanks to Tanguy for his take on the news with a sustainability twist. If you're keen to find out more about the role of AI in climate data collection, Tanguy has co-authored a blog with our colleague Marion de Marillac. It is up on MSCI's website, free to access and

read in its entirety. And the piece is called The Promise and Limitations of AI in Climate Data. It's a great blog. Gives some very practical examples of where AI can be used and where those potential shortcomings might be. It really won't take longer than your next cup of coffee, and you will be well equipped to wade into any discussions about collecting climate data with AI, with your colleagues, parents, pets and children.

As always, I do want to say thank you very much for tuning in. If you like what we're doing, then let us know. Drop us a review, rate the show on your platform of choice, and tell a friend or a colleague about this episode, especially one who won't shut up about market-based scope two emissions. Thanks again, and until next time, take care of yourself and those around you.

The Sustainability Now podcast is provided by MSCI Solutions LLC, a subsidiary of MSCI Inc. Except with respect to any applicable products or services from MSCI Solutions, neither MSCI nor any of its product or services recommends, endorses, approves, or otherwise expresses any opinion regarding any issuer, securities, financial products or instruments, or trading strategy and products or services are not intended to constitute investment advice or recommendation to make or refrain from making any kind of investment decision, and may not be relied on as such. The analysis discussed should not be taken as an indication or guarantee of any future performance, forecast or prediction. The information contained in this recording is not for reproduction in whole or in part, without prior written permission from MSCI Solutions. Issuers mentioned or included in any MSCI Solutions material may include clients of MSCI or suppliers to MSCI, and may also purchase research or other products or services from MSCI Solutions. MSCI Solutions materials, including materials utilized in any MSCI Sustainability & Climate indexes or other products, have not been submitted to nor received approval from the United States Securities and Exchange Commission or any other regulatory body. The information provided here is as is, and the user of the information assumes the entire risk of any use it may make or permit to be made of the information. Thank you.