

NASDAQ:INOD Q4 2025 Earnings Call Transcript

Generated on 6/10/2026

Background Audio | Music/Applause:

Thank you. Thank you. music music ¶¶ Thank you. Thank you. Thank you. Thank you. We'll be right back. Thank you. Thank you. Thank you. Thank you. © transcript Emily Beynon Thank you. Thank you. Thank you. © transcript Emily Beynon Thank you.

Conference Operator | Operator:

Good afternoon, ladies and gentlemen, and welcome to the InnoData 2 Report fourth quarter and fiscal year 2025 results conference call. At this time, all lines are in listen-only mode. Following the presentation, we will conduct a question and answer session. If at any time during this call you require immediate assistance, please press star zero for the operator. This call is being recorded on Thursday, February 26, 2026. I would now like to turn the conference over to Amy Agras, General Counsel. Please go ahead.

Amy Agras | General Counsel:

Thank you, operator. Good afternoon, everyone. Thank you for joining us today. Our speakers today are Jack Applehoff, Chairman and CEO of InnoData, and Maryse Espinelli, Interim CFO. Also on the call today is Anish Pentakar, Senior Vice President, Finance and Corporate Development. Rahul Singhal, President and Chief Revenue Officer, is unable to be here today, but looks forward to joining us on our next call. We'll hear from Jack first, who will provide perspective about the business, and then Maryse will provide a review of our results for the fourth quarter and fiscal year 2025. We'll then take questions from analysts. Before we get started, I'd like to remind everyone that during this call, we will be making forward-looking statements, which are predictions, projections, and other statements about future events. These statements are based on current expectations, assumptions, and estimates, and are subject to risks and uncertainties. Actual results could differ materially from those contemplated by these forward-looking statements. Factors that could cause these results to differ materially are set forth in today's earnings press release in the risk factor section of our Form 10-K, Form 10-Q, and other reports and filings with the Securities and Exchange Commission. We undertake no obligation to update forward booking information. In addition, during this call, we may discuss certain non-GAAP financial measures. In our earnings release filed with the SEC today, as well as in our other SEC filings, which are posted on our website, you will find additional disclosures regarding these non-GAAP financial measures, including reconciliations of these measures with comparable GAAP measures. Thank you. I will now turn the call over to Jack.

Jack Applehoff | Chairman and CEO:

Thank you, Amy, and good afternoon, everyone. Q4 was another strong quarter for InnoData. We generated 72.4 million in revenue, reflecting 22% year-over-year growth. This brought our full-year revenue to 251.7 million, representing 48% year-over-year growth for 2025. Our Q4 consolidated adjusted gross margin was 42%, exceeding our externally communicated target of 40%. Our adjusted EBITDA totaled 15.7 million or 22% of revenue, also exceeding analyst consensus by 1.2 million. In fact, our results exceed the analyst consensus across the range of key metrics, including revenue, adjusted EBITDA, net income, and EPS. We ended the year with 82.2 million in cash, up sequentially by approximately 8.4 million. We achieved these results while making meaningful growth-oriented investments in both COGS and SG&A. In COGS, we carried capacity ahead of revenue ramp, which consistently proved to be the right move. And in SG&A, we invested

in engineers, data scientists, and customer-facing account leadership, which investments also proved prudent, yielding innovation that has expanded our opportunities. We believe our business momentum to be at an all-time high. We are seeing robust demand across the entire generative AI lifecycle, spanning development, evaluation, and ongoing model optimization. And we believe we are gaining traction with a broad and diversified number of large customers. As a result of market demand and growing traction, we anticipate another year of potentially extraordinary growth in 2026. We currently estimate our 2026 year-over-year growth to potentially be approximately 35% or more. This estimate reflects active programs, recently awarded wins, late stage evaluations, and opportunities where we have clear line of sight. Because we are early in the year and because LLM initiatives spin up quickly, we believe there may potentially be significant upside to this range. However, we prefer to guide conservatively and adjust upward as visibility increases. At the same time, given the scale and complexity of the programs we support, timing variability and customer round schedules, budget approvals, or shifts in research priorities could influence the pace at which revenue materializes. Embedded in our outlook is the expectation that spend from our largest customer will increase somewhat in the year, and that the remaining customer base in the aggregate will grow at a faster rate. We expect this other customer growth to come from a mix of the MAG7, domestic AI innovation labs, sovereign AI initiatives, and leading enterprises. We believe this will meaningfully contribute to customer diversification. Our customers are moving fast driving shorter development cycles and responding faster to research breakthroughs in 2025 we succeeded in this environment in no small part because we followed the research anticipated customer needs and pivoted were required to illustrate. In the first quarter of this year, for our largest customer, we deprecated the meaningful number of post-training workflows, which represented in the aggregate approximately 20 million of annualized revenue run rate, but replaced them with a combination of new post-training workflows and scaled pre-training programs, an area of recent focus and investment. From a revenue run rate perspective, the net effects turned out positive. Indeed, we believe continuous innovation is critical to achieving our ambitious plans for 2026 and beyond. The truly exciting news is we believe we are entering a golden age of innovation at InnoData as a result of investments we have made and intend to make in the future. I'm now going to share some of our recent innovation initiatives. For competitive reasons, we'll be appropriately circumspect, but what we share will give you a meaningful window into how we're thinking, where we're investing, successes we're having, and how we intend to capitalize on the opportunity ahead. I'll briefly walk through our recent innovation in three areas, generative AI model training, agentic AI, and physical AI. Before I do, I want to underscore a unifying theme. Every innovation I am about to discuss is fundamentally a data innovation. Whether the goal is more capable LLMs, more reliable autonomous agents, or more intelligent physical AI systems, data quality, data composition, data validation, and data engineering at scale are at the heart of the matter. These are our core competencies. We'll start with generative AI training. Historically, customers told us the kind of training data they wanted. Increasingly, however, they're asking us to diagnose model performance, design the right training data sets, and demonstrate that those data sets will materially improve outcomes. Here's how that works. We begin by identifying performance gaps using our evaluation frameworks. We then engineer targeted data sets and validate their efficacy by fine tuning either the customer's model or a structurally similar proxy model. Only after we measure and demonstrate performance impact do we scale. This shifts the discussion from how much is the data to how effective is the data. We believe this shift is being driven by two forces, the accelerating pace of AI research and the cost and time incurred to train ever larger models. And conversations about data efficacy play directly to our strengths. We are also advancing methods for creating datasets that improve long context reasoning and AI models ability to observe and reason over very large amounts of information at once. This remains one of the industry's most important technical challenge. Solving it requires not just architectural improvements, but advances in the creation at scale of very specific types of structured training data. Creating training data that improves long context reasoning is a nontrivial problem, but we have made and are continuing to make meaningful progress on it. A secondary of innovation is around evaluating systems of autonomous agents and improving them through targeted data set creation. We believe that autonomous agents may represent the most significant business innovation opportunity since the advent of electricity. But companies quickly discover that many AI agents that performed impressively in controlled laboratory settings degrade in real world production. The real world is chaotic. It's shaped by edge cases, conflicting constraints, unpredictable user behavior, and adversarial conditions. Addressing this is

fundamentally a data challenge. Agents must be continuously trained and rigorously stress tested with datasets that are realistic, diverse, and complex. For this, we have developed a set of three highly complementary hybrid solutions. The first is an agent evaluation and observability platform. Data scientists can use our platform during development to visualize and annotate agent trace data, to build LLM as a judge evaluators, to create business aligned evaluation rubrics, to generate golden data sets for regression testing, and to generate test data at scale. Then, once the agent is deployed, Our platform can be used to continuously monitor its performance, perform root cause analysis and performance issues, and obtain mitigation data sets. We're pleased to share that we anticipate soon kicking off a managed services engagement with a hyper scalar in which we will use our platform to create test data and scale. perform automated evaluations and identify critical model vulnerabilities in order to improve performance of its customer facing intelligent virtual assistant. The second innovation is a managed agent optimization pipeline designed to systematically train for and therefore neutralize the chaos of real-world deployment at scale. The pipeline generates realistic test scenarios, automates evaluation, rigorously measures constraint satisfaction, and produces reinforcement learning datasets. Using this system, we have demonstrated improvements of up to 25 points and constraint satisfaction. Importantly, agents trained using conventional techniques tend to degrade significantly as task complexity increases. By contrast, agents trained through our pipeline sustain their performance under escalating real-world difficulty. In the most demanding scenarios, the performance gap between standard approaches and our system widens to more than 31 points. We currently have multiple AI innovation labs and enterprise customers actively exploring the system. The third solution we've designed to support enterprise agentic AI is an adversarial simulation system that generates high-quality, semantically diverse and scalable adversarial attacks to stress test agents. The system generates a full spectrum of attack types, direct jailbreaks, indirect prompt injection via RAG pipelines, multi-turn social engineering stenographic payloads, and compound attacks that combine injection techniques with domain-specific knowledge. Once vulnerabilities are identified, it generates highly targeted mitigation datasets to strengthen guardrails. We believe our system generates realistic adversarial attacks of scale in a meaningful way that exceeds existing alternatives. Many tools on the market produce simplistic or templated hostile content that lacks the nuance and sophistication of real-world threat actors, fails to scale across diverse scenarios, or relies on generic tactics that models quickly learn to anticipate and overfit to. But by contrast, our framework is designed to simulate adaptive, multi-step, and strategically coherent attack patterns, including highly sophisticated model extraction, cybersecurity, cybercrime, and Soberny threat scenarios that better reflect how advanced adversaries operate and allow our partners to stay ahead of emerging threats. The result is adversarial training data that is both scalable and durable, forcing models to generalize rather than memorize and enabling more robust real-world resilience. Our work is garnering interest from CISOs, and security leaders at some of the world's premier AI and cybersecurity companies, as well as relevant experts in government, and has led to early stage engagements with several of them. At a time when the cyber industry is experiencing significant disruption, these capabilities bolster our position in the emerging field of AI trust and safety, an area where we are meaningfully deepening work with several hyperscalers. We believe Enerdata is well positioned to emerge as a leader in prompt layer security, protecting AI systems at the point of interaction rather than relying solely on traditional perimeter or endpoint defenses. Taken together, we believe these solutions position us not just as a data supplier, but as a lifecycle partner in agent reliability. We believe 2026 will also mark the acceleration of physical AI, intelligent systems that perceive and interact with the physical world. While robotics provides the mechanical framework, physical AI provides the intelligence. The primary bottleneck in this domain is data set quality and scale. Manual annotation and static QA sampling simply do not scale to billion-sample corpora and continuously evolving environments. We have developed a large-scale data engineering system that incorporates structural validation, distribution monitoring, temporal consistency checks, and model-in-the-loop instrumentation. This enables us to identify and correct defects in datasets before they propagate into performance failures. We're already using components of this system in the high visibility engagements we recently announced with Palantir. We recently secured a significant engagement to create foundational datasets for next generation robotic datasets, including egocentric data. Egocentric data captures the world from the robot's point of view, what it sees and experiences in motion. We are also working with a leading robotics lab to create affordance data at scale. Affordance data teaches the system what actions are possible in a given setting, not just identifying objects, but understanding how they

can be used. Ego-centric data and affordance data taken together form the cognitive scaffolding that allows machines to act intelligently in dynamic environments. This work also positions us to support the development of so-called world models, internal simulations that allow AI systems to anticipate outcomes, reason about cause and effect, and plan several steps ahead. World models require richly structured data sets that capture interactions over time and the consequences of actions, precisely the type of data we are now engineering. We recently developed an AI model for drone and other small object detection that exceeds prior state-of-the-art benchmarks by 6.45%. In a field where progress is often measured in fractions of a percentage point, a 6.45% improvement is a material advance. The model improves detection fidelity under real-world conditions where small size feed, cluttered backgrounds, and environmental noise make reliable perception extraordinarily difficult. We believe this advancement has compelling dual-use implications that we are now actively exploring with potential customers. I'd like to underscore one of the important points I just made. For decades, InnoData has specialized in creating high-quality complex datasets. Today, these capabilities are central to unlocking the next generation of AI systems. Advanced LLM reasoning, agent reliability in chaotic environments, and robotics perception in the physical world all depend on engineered data ecosystems. And this is precisely where we operate. Our innovations in LLM training, agentic AI, and physical AI are not separate initiatives, rather, They are extensions of a single strategic advantage, our ability to engineer data that measurably improves model performance in real-world conditions. We believe our innovation pipeline will be margin-enhancing as well as revenue-enhancing. We expect early 2026 adjusted gross margins to be in the 35 to 40% range as we ramp up new programs, with normalization toward our target 40% or better adjusted gross margins as new programs ramp up and as innovation-driven workflows scale. Automation, synthetic systems, and evaluation platforms all structurally increase our operating leverage. I'll now turn the call over to Maryse, who will go through the numbers.

Maryse Espinelli | Interim CFO:

Thank you, Jack, and good afternoon, everyone. Revenue for Q4 2025 reached 72.4 million, up 22% year over year. Sequentially, revenue increased 15.7% from Q3, 62.6 million. Adjusted gross profit for Q4 2025 was 30.1 million, an increase of 6% year-over-year and 9% sequentially, with an adjusted gross margin of 42%. Adjusted EBITDA was 15.7 million, or 22% of revenue, and net income for the quarter was 8.8 million. To reiterate, this is net of significantly expanded data science and engineering efforts that are yielding the types of innovation Jack just spoke about. We ended the quarter with \$82.2 million in cash, up from \$73.9 million at the end of prior quarter, and \$46.9 million at the year-end 2024. And we did not throw down on our \$30 million Wells Fargo credit facility. As Jack mentioned, based on our current momentum, we presently forecast 35% or more year-over-year revenue growth in 2026. Thank you, everyone, for joining us today. Operator, please open the line for questions.

Conference Operator | Operator:

Thank you. Ladies and gentlemen, we will now begin the question and answer session. Should you have a question, please press the star key by the number one on your touch-tone phone. You will hear a prompt that your hand has been raised. If you wish to decline from the polling process, please press the star key followed by the number two. If you are using a speakerphone, please lift the handset before pressing any keys. One moment, please, while we assemble the queue. Your first question comes from Trevor Sutton of Craig Hallam. Please go ahead.

Trevor Sutton | Analyst, Craig Hallam:

Thank you, Jack. I feel like I just sat through an advanced AI data science class, so thanks for that. wanted to uh step back a little bit because i think people have the assumption that some of what's working for you is somewhat temporary and i think you've you've done an interesting job of kind of walking us through in past

quarters from post training as a start then pre-training and now there are dramatic other use cases including things like robotics and autonomous agents Can you just talk about the breadth of the things you're seeing and sort of where you see us in this continuum of data science opportunity for you?

Jack Applehoff | Chairman and CEO:

Sure. Thank you, George. Thank you for the question. So as we look out near term, 2026, we see ourselves as being incredibly well set up by the innovations that we invested in in 2025. And we see that innovation output as a flywheel, we're getting better, we're getting stronger, we're creating solutions that are solving problems that are the actual impediments that enterprises have when they're looking to integrate AI into their operations. So when you look across the spectrum of current capabilities in AI and future capabilities in things like agentic systems, you know, physical AI, robotics. All of this boils down to challenges in terms of data engineering. Of course, there are going to be continuous improvements in architectures. It'll be, you know, bigger models. There'll be narrower models for, you know, domain-specific, you know, challenges. But at the heart of it, in terms of making systems reliable, making them safe at an enterprise level, it's going to be about innovations such as the ones we're announcing today in data sets that are used for evaluation, data sets that are used, you know, for training and improving safety and reliability of models. So we think that we're at the very beginning and that our relevance is by no means diminishing, but only increasing. It's increasing not just at the level of foundation model builders, but it's clearly extending through the enterprise. We're super excited. about where we are right now and about the uptake that the innovations that we're creating are having and are going to be having over the next several years.

Trevor Sutton | Analyst, Craig Hallam:

That's great. And just one other question. Having lived through the last couple of years where you started the years with an expectation and you then ended up meaningfully exceeding those initial expectations, is Is anything set up differently going into 2026 relative to what you see in your sites relative to what you're committing to today?

Jack Applehoff | Chairman and CEO:

No, not at all. We're following exactly that same methodology. You know, we're really limiting our or we're taking a conservative approach to forecasting growth based on opportunities where we have a very clear line of sight. but where we can't predict a close rate, where we can't feel pretty confident in something happening, we're just not baking that into our guidance. Our aspiration is to surprise and to beat expectations. When I look at this year, I think it will likely be another year of doing exactly that. We're We're seeing enormous opportunity with a much larger set of customers. We think that that's going to result in growth. I think it's likely that we'll be increasing guidance as we move through the year. And I think it's going to be a year where we accomplish very meaningful customer diversification. On top of that, as we already discussed, I think it's going to be a year where, you know, we're starting to see increasingly hybrid human slash technologically driven solutions. That spells or presents the promise, I believe, for increased recurring revenue. I think it promises greater margins over time, greater stickiness, a whole lot of things that will over time be, I believe, consistently improving revenue quality as well on top of everything else. In terms of the work we do with foundation model builders, you know, we're seeing tons of traction, not just in our largest customer, but in others as well. We're very much aligned with what they're looking to accomplish in things like long-context reasoning improvements. We have innovations that are contributing to that. So we're tremendously excited about where we are right now.

Trevor Sutton | Analyst, Craig Hallam:

All right, good stuff. Thanks, Jack. Thank you.

Conference Operator | Operator:

Your next question comes from Hamid Khorband of BWS Financial. Please go ahead.

Hamid Khorband | Analyst, BWS Financial:

Hi. Just a first question. You were talking earlier about scaling your operations as revenue ramps. Do you have enough employees now? Do you see the need to add more employees? What's your timeline as far as expecting gross margin to move up from here? Thank you.

Jack Applehoff | Chairman and CEO:

Sure, thanks. So I think it really depends on what we're seeing. I think if we begin to project internally growth rates that are very significant, we're going to be making investments in order to ensure that we capture those growth rates. I do think that as a result of digesting some of those you know uh people investments that we're making in cogs um as a result of the innovations that we're discussing um you know different things like that I do think that we're going to uh you know I do think that we're going to be seeing movement you know back toward our target gross margins over time

Hamid Khorband | Analyst, BWS Financial:

Okay. And then is there a timing as far as this pipeline of deals that you're talking about with other customers other than your largest customer?

Jack Applehoff | Chairman and CEO:

So there are pipelines, but the deals that I'm referring to are largely deals that we're closing or have closed. So we're not depending on We're not speculating about what will be happening. These are things that are actively underway.

Hamid Khorband | Analyst, BWS Financial:

Okay. Thank you.

Conference Operator | Operator:

Your next question comes from Alan Clee of Maxim Group. Please go ahead.

Alan Clee | Analyst, Maxim Group:

Yes. Hi. For 2025, I think your adjusted EBITDA margin was around 23%. And I know it's important for you to reinvest back into the business for the health of the company. My question is, is there any reason to think that you would target a higher or lower adjusted EBITDA margin than what you did in 2025?

Jack Applehoff | Chairman and CEO:

We're very much focused on seizing opportunity right now. We believe that we can do that and stay profitable, but we also believe that it's more important to seize opportunity and to do some of the things that we are describing and prove out those innovations than it is to track adjusted gross margin percentages and try to maintain a certain percentage. So, you know, we're going to be actively reinvesting in the business. The more opportunities we see, to some extent, the more we'll be reinvesting. We do believe, though, that maintaining profitability is something that we can do while we drive very aggressive growth and while we become more progressively more critical to a larger and widening set of customers.

Alan Clee | Analyst, Maxim Group:

Okay. One of the bullet points you had on the innovation was the structural foundation for margin expansion through automation, synthetic data generation, and evaluation platforms. Can you explain a little what you mean of which margin extension are you referring to? Thank you.

Jack Applehoff | Chairman and CEO:

Yeah, so we're referring to overtime gross margin expansion. So a lot of the innovations that we're working on now and that we're bringing into the market are hybridizations of software and human teams. And I think that over time, we're going to be seeing the gross margins associated with those capabilities to be perhaps well in excess of the gross margins that we target today.

Alan Clee | Analyst, Maxim Group:

Got it. That makes a lot of sense. And the last question I had was just for first quarter 26, is there anything you'd want to point out in terms of that might stand out just in terms of I don't know, revenues or expense spend?

Jack Applehoff | Chairman and CEO:

Well, you know, I'm not going to say it's next quarter necessarily, but I think, you know, very soon we're going to be seeing quarters that, you know, from a revenue perspective are beating what our revenue was for an entire year three years ago. So that's pretty good news right there. As we move through the year, I think you're going to be seeing more proof points and more evidence and more engagement that we have with some very interesting companies around the innovations that we're describing. I think that we'll start to demonstrate that we're somewhat migrating from a vendor to a foundational layer within AI ecosystems, becoming someone that is able to unlock the promise of AI within enterprise engagements, a company that's able to help enterprises embrace complex agents that plan, call tools, execute complex workflows, and create a lot of value. You know, I think we'll be seeing that. I think we'll see evidence of that in first quarter. I think we'll continue to see evidence of that through the year.

Alan Clee | Analyst, Maxim Group:

Maybe one last quick one. When you were talking about your largest customer, I don't know if I fully understand, but you mentioned something about \$20 million that maybe is going to be replaced with more than that, or could you just explain what

Jack Applehoff | Chairman and CEO:

Yeah, I think the point that we were making there is how important innovation is to our company today and how it's becoming increasingly important. You know, there are things that we complete and we're starting new things, and by following the path, of innovation by, you know, what did Wayne Gretzky used to say, by skating to where the puck is going. We're able to deprecate things that the companies no longer require, but be there for them for the things that are the emerging requirements. Again, you know, we're seeing the emerging requirements to be more interesting from a business perspective. and a revenue quality perspective and a differentiation perspective than the things that came before. So the investments are proving out. They're enabling us to scale and increase the breadth of engagements. They're enabling us to win new engagements and new customers that some of which we think are going to be very substantial. They're going to really flower this year. That's going to address the diversification issue. So, you know, when we look at 2026, you know, we see a huge growth here. We believe that we're going to be increasing likely our guidance from what we're starting the year at. We think that the solutions and how we're embedded in workflows is going to be progressively more interesting and margin and revenue enhancing. It promises to be a tremendous year on all of those fronts.

Alan Clee | Analyst, Maxim Group:

That's great. Congratulations. Thank you. Thank you.

Conference Operator | Operator:

There are no further questions at this time. I will now turn the call back over to Jack Abuhoff. Please continue.

Jack Applehoff | Chairman and CEO:

Thank you, operator. So, yeah, to wrap up, 2025 was a great year, and 2026 holds the promise of being even better. In 2025, we delivered strong top line growth. We exceeded expectations across major financial metrics. We expanded margins. We strengthened our balance sheet. We invested successfully ahead of demand. And those investments proved wildly successful and set us up well for 2026. I believe that 2026, is likely to be an incredible year. We've got it to approximately 35% growth based on visibility today, but I believe there may be very considerable upside to that. We'll update you through the course of the year, much like we have done the last couple of years. I also want to underscore our belief that this year we will potentially diversify our revenue stream significantly. And we believe expertly engineered data ecosystems are going to be every bit as important as bigger models and new architectures will be in terms of advancing language models, media models, autonomous agents, robots, world models, and other kinds of AI that hasn't even been conceived of yet. So we're very excited about what lies ahead. We're very confident in our positioning. We're very committed to building one of the most important and we think most capable AI enablement companies in the industry. It's going to be an exciting year. So thank you all for being on the journey with us. Look forward to next time.

Conference Operator | Operator:

Ladies and gentlemen, that concludes today's conference call. Thank you for your participation. You may now disconnect.