Pony AI Inc.

Third Quarter 2026 Earnings Conference Call

Operator

Ladies and gentlemen, thank you for standing by, and welcome to Pony AI Inc.'s third quarter 2026 earnings conference call. At this time, all participants are in listen-only mode. After the management's prepared remarks, there will be a question-and-answer session. As a reminder, today's conference call is being recorded and a webcast replay will be available on the company's Investor Relations website at ir.pony.ai.

I will now turn the call over to your host, George Shao, Head of Capital Markets and Investor Relations at Pony.ai. Please go ahead, George.

George Shao

Thank you, operator, and hello everyone. We appreciate you joining us today for Pony.ai's third quarter 2026 earnings call. Earlier today, we issued a press release with our financial and operating results, which is available on our Investor Relations website. An earnings presentation, which we will refer to during this conference call, can also be accessed and downloaded on our Investor Relations website.

Joining me on today's call are, Dr. James Peng, Chairman of the Board and CEO; Dr. Tiancheng Lou, CTO; and Dr. Leo Wang, CFO of the company. They will provide prepared remarks followed by a Q&A session.

Before we begin, please refer to the Safe Harbor Statement in our earnings release, which applies to this call as we'll be making forward-looking statements. Please also note that we will discuss non-GAAP measures today, which are more thoroughly explained and reconciled to the most comparable measures reported under GAAP in our earnings release available on our Investor Relations website and filings with the SEC and the Hong Kong Stock Exchange. I will now hand it over to our Chairman and CEO, Dr. James Peng. Please go ahead.

Dr. James Peng

Thank you, George. Hello everyone, thank you for joining our earnings call. For Q3, we actually delivered another amazing quarter.

Let me walk you through four highlights for this quarter. First, total revenue grew 75% year-over-year, with Robotaxi revenue more than doubled and reaching 40% of total revenue. Second, our Robotaxi fleet now surpasses 2,700 vehicles, expanding across both domestic and international markets. This massive scale allowed us to interconnect our Guangzhou footprint into a single, contiguous coverage area spanning over 1,000 square kilometers. Third, we continue to expand globally, with Europe's first fully driverless operations in Zagreb, Croatia, and our entry into key UK markets. Fourth, we achieved strong growth beyond just the passenger vehicles. Robotruck revenue grew 76% year-over-year. We also initiated Gen-4 Robotruck commercial operations, and made solid progress on robo light-truck. This is our real edge — one core stack generalizes across vehicle platforms and operating environments, turning R&D into multiple revenue streams.

Let me first start with our progress in China. As regulation normalizes and entry standards rise, our first-mover advantage becomes more valuable. With scaled operation, our safety record keeps improving; and as a result, insurers are now actually pricing in our lower risk profile. Recently we renewed our vehicle insurance, and the insurance premium per Gen-7 Robotaxi came in at just 40% of those for average human-driven taxis, down from 50% last year. This is a clear indicator of our superior safety track record.

In terms of Robotaxi operation, we continue to deepen our leading position in high-value mobility markets.

In Guangzhou, building on the city-wide commercial permit we secured earlier this April, we have taken the lead in providing cross-district Robotaxi services linking Nansha district with the downtown central area. Our previously separate operating areas in the city's central districts and Nansha now form a single contiguous service area, covering more than 1,000 square kilometers. That's roughly nine times the size of San Francisco. Operating a fleet at this scale is a clear demonstration of our leading technology and solid dispatching capability. Through the PonyPilot App, users can now book rides across five districts including Nansha, Haizhu, Tianhe, Huangpu, and Panyu, directly lifting our trip distance and average price per ride.

In Shenzhen, our service at Bao'an International Airport is already open to users. We are now the only company operating a commercial, fully driverless Robotaxi service at an international airport in Chinese tier-one cities. Shenzhen is hosting APEC these days, and our Robotaxis are front and center in the city's showcase of innovation to the world. We're proud to be a living example of openness, innovation, cooperation and green mobility that APEC forum and Shenzhen champion to the world.

Meanwhile, in Shanghai, we have secured additional commercial permits for fully driverless operations in the core areas of Pudong District. Supported by these new licenses and an expanded fleet, we will very soon extend our coverage into key economic hubs, such as Zhangjiang Science City. This is home to numerous high-tech semiconductor and biopharmaceutical companies in Shanghai.

Beyond the tier-1 cities, we are also entering into new high-commercial value cities. We have recently initiated deployments in Hangzhou, Chongqing, and Zhuhai, laying the groundwork for further expansion of our operational footprint. We should also see Robotaxi launches in several other cities soon.

Now let's switch gears to overseas markets. We are accelerating penetration into high-value cities, especially those in Europe and the Middle East. Backed by a proven safety record and business model, we continue to capture shares in key markets. Under our joint deployment model, we secured mobility partners and started generating high-margin revenues.

In Zagreb, Croatia, we achieved Europe's first fully driverless passenger rides — going from launch to fully driverless in just five months. This is our joint deployment model in action. Our core technology and operating know-how provide the foundation, while our partners, such as Uber and Verne, facilitate local user acquisition, service operations, and regulatory adaptation. Our operations there have accumulated more than 200,000 kilometers and completed thousands of passenger trips. Our success in Croatia has become a replicable template for us to enter other global markets.

We continued to implement our previously announced pipeline of over 4,000 Robotaxis. In the UK, by collaborating with Uber we are advancing deployments in London and Manchester. The pace keeps on accelerating as road testing is already underway. We also prepare to launch Robotaxi in several key markets in Europe.

In Dubai, we are accelerating our commercialization progress. At last month's WETEX 2026 in the Dubai World Trade Centre, we were the only Robotaxi partner showcased alongside Dubai's Roads and Transport Authority, laying the groundwork for further expansion.

Some other international efforts include entering into Central Asia for the first time. Together with inDrive, we will launch Robotaxi operations in Kazakhstan.

Let me now turn to Robotruck, where we deploy the same underlying technology and safety framework that we developed for Robotaxi across more vehicles and use cases.

We unveiled a new Gen-4 heavy-duty Robotruck model in collaboration with GAC Commercial Vehicle Co., further broadening our Gen-4 product lineup. The new vehicle is built on the T9 battery-electric heavy-duty truck platform.

We are now entering operations at Jiangmen Port and Mawan Port — both in Guangdong Province. This is China's largest provincial economy and biggest export hub. The operation will generate ongoing revenue. With success in these two ports, we have built a solid deployment pipeline with various leading port operators globally. Our target is to reach hundreds of vehicles across ports next year.

In addition, we are tapping into urban logistics with our light-duty trucks, which recently entered production. We are also advancing commercial deployments across key logistics corridors in the tier-one cities in China.

At the beginning of this year, we have set three targets for the year. First, we will deliver 3.5x Robotaxi revenues over prior-year levels. Second, our fleet is scaling to 3,500 vehicles, and third, our global Robotaxi deployments will cover 20 cities worldwide. As we approach year-end, we are on track to reach all three ambitious targets that we planned earlier.

Looking ahead to 2027, we will scale up our Robotaxi and Robotruck businesses with full momentum. Robotaxi will expand across three fronts: deepening operations in existing cities, ramping up in new Chinese cities, and accelerating overseas in strategic markets. Our dual-engine domestic and global growth strategy will increase commercial vehicles in operation, directly driving top-line growth. Robotruck will deliver another powerful engine, as Gen-4 models and new initiatives begin unlocking new revenue streams. This is where our flywheel becomes truly powerful: greater scale yields superior technology, higher efficiency, and stronger economics.

With that, I will turn the call over to our CTO, Tiancheng, to elaborate on how our technology powers this flywheel. Tiancheng, please go ahead.

Dr. Tiancheng Lou

Thank you, James. Hello everyone, this is Tiancheng. I will elaborate more on technology.

Let me start with where we are. Back in early 2025, we were already running over a hundred fully driverless Robotaxis and continued to expand. Since then, we've scaled that operation more than tenfold.

The real story is how we scale. L4 Robotaxis and Robotrucks operate in the physical world, where safety is non-negotiable. Every second, they interact with vehicles and pedestrians in real time. When a decision goes wrong, the damage is done — there is no rollback, and no retry. That fundamentally separates us from digital AI applications, where models can fail and retry endlessly, and from L2 assisted driving, where a human driver remains behind the wheel as the ultimate safety fallback.

Like frontier Large Language Models, Physical AI also has two stages of training: pre-training and post-training.

Pre-training on massive multi-modal data establishes basic driving skills, but its performance ceiling remains bounded by human driving data — leaving long-tail hallucinations unresolved.

Breaking beyond the human safety ceiling requires post-training by reinforcement learning. In the latest AI engineering, over 90% to 95% of the core work happens here.

This is where our World Model comes in. *PonyWorld* is our closed-loop virtual driving school. Inside it, our AI driver executes billions of reinforcement learning and counterfactual scenarios to eliminate hallucinations and solve long-tail safety problems.

However, Physical AI faces a critical challenge that digital AI like coding do not: the sim-to-real gap. In coding, the environment is deterministic. For coding, the training environment and the real-world environment are the same one. In autonomous driving, our World Model must mirror the exact probability distribution of the physical world. Edge cases cannot appear too rarely, or the model won't learn to handle them; nor can they be artificially exaggerated, or the driving policy becomes overly conservative. The simulated distribution has to mirror reality.

This is our decisive edge. We built an AI-driven self-evolving closed-loop pipeline that continuously calibrates and upgrades *PonyWorld* using real-world fleet feedback. Every edge case encountered on public roads automatically sharpens the World Model's fidelity. As a result, *PonyWorld* not only evolves at a faster speed, but breaks through to a fundamentally higher performance ceiling — because it overcomes the limits of human comprehension and manual scenario design.

Today, *PonyWorld* generates 10 billion kilometers of training every week — scaling post-training in the virtual world before deploying onto public roads.

Because *PonyWorld* dominates post-training, it solves complex long-tail interactions with unmatched efficiency — directly powering our commercial expansion across four fronts:

First, breakthrough safety. Real-world risks come from rare, extreme interactions. By using Intention Enumeration inside *PonyWorld*, we systematically generate all possible intentions and trajectories of surrounding road users in a given scenario. This allows us to exhaustively stress-test every extreme interaction path in simulation and rigorously verify the completeness and safety boundary of our driving policy before deployment. This drives measurable safety improvements — reflected in our insurance premiums dropping to 40% of those for human-driven taxis, compared with 50% last year.

Second, rapid area expansion without retraining. Inside *PonyWorld*, reinforcement learning trains our AI driver to master all kinds of generalizable scenarios instead of a specific road or area. Wherever traffic rules and driving styles remain consistent, the model generalizes natively with zero retraining. Because our safety is proven at the scenario level, expanding our fully driverless coverage to over 1,000 square kilometers in Guangzhou — and scaling to tens of thousands of square kilometers next — is purely a function of deploying more vehicles, not repeating data collection or R&D cycles.

Third, scalable global replication. Entering new countries requires adapting to distinct traffic laws, left- or right-hand driving, and local driving patterns. *PonyWorld*'s High-Fidelity Interactive Smart Agents accurately replicate local road-user behavior, giving our AI driver localized interaction training prior to launch. This

directly powered Europe's first fully driverless operation in Zagreb in just a few months, and gives us a path to replicate internationally at scale.

Fourth, cross-platform generalization. Robotaxi and Robotruck share one unified core architecture. By Sensor-Agnostic Generalization, our model extracts a unified physical representation across diverse vehicle types and sensor setups, efficiently adapting by post-training to heavy-duty trucks and light trucks. One R&D engine unlocks multiple revenue streams with minimal incremental cost.

Finally, unlike other digital Large Language Models that expose public APIs and can be easily distilled, our models operate strictly inside our closed-loop system. Anyone who wants to replicate it must build this entire post-training infrastructure and real-world data flywheel from scratch. This gives us a much deeper moat than Large Language Models, making it far harder for latecomers to catch up.

Looking forward, this technical edge will keep compounding. As more real-world operating data flows back into *PonyWorld*, our simulation fidelity and training efficiency will continue to advance. That means when we expand into larger operating areas, enter more overseas cities, and extend into more vehicle platforms, the marginal R&D cost will keep coming down — and the speed of commercialization will keep going up.

With that, I'll hand the call over to our CFO, Dr. Leo Wang, to review our financial results. Leo, please go ahead.

Dr. Leo Wang

Thank you, Tiancheng. Hello, everyone. This is Leo.

I will now review our financial results for the third quarter of 2026. As a reminder, all comparisons are on a year-over-year basis unless otherwise stated. Please refer to our earnings release for detailed financial results and the reconciliation of U.S. GAAP and non-GAAP results.

Total revenues increased by 75% to US$44.6 million in the third quarter of 2026, driven by growth in our Robotaxi and Robotruck businesses. Together, these two business lines accounted for 80% of total revenues, compared with 66% a year ago.

As our business continued to scale, revenues grew much faster than operating expenses, reflecting improved operating leverage. Our operating margin improved from negative 274% to negative 154%.

Let me start with Robotaxi. Robotaxi services revenues maintained strong growth momentum, reaching US$17.8 million, up 165% year-over-year and 47% quarter-over-quarter. This made Robotaxi the largest contributor to our year-over-year revenue increase. Robotaxi revenues accounted for 40% of total revenues, up from 26% in the prior-year period.

Along with this revenue growth, our fleet continued to expand. As of September 30, 2026, our global Robotaxi fleet exceeded 2,700 vehicles.

Just as a quick recap, our Robotaxi business operates through two primary models: our self-owned fleet and the joint deployment model. In China, we primarily generate fare-charging revenues from our self-owned fleet. Overseas, we strategically partner with local fleet operators and mobility platforms under the joint deployment model, generating revenues from vehicle sales, as well as *Virtual Driver* licensing fee.

Drilling down to fare-charging revenues, we recorded approximately 400% year-over-year growth in the third quarter, supported by the continued expansion of operating areas and service offerings. In Guangzhou, operations in Haizhu, Tianhe, Huangpu and Panyu districts, which we highlighted in our previous earnings release, generated incremental revenue. In Shenzhen, our services to and from Bao'an International Airport and Shekou Cruise Port also contributed to third-quarter fare-charging revenues.

Under the joint deployment model, we will continue to advance deployments in markets with higher fares, while accelerating the rollout of our previously announced pipeline of more than 4,000 Robotaxis with Uber and other partners.

As commercialization progresses in China and overseas, we are on track to achieve full-year Robotaxi revenues of 3.5 times the 2025 level, alongside our year-end targets of 3,500 vehicles and a presence in 20 cities.

Now let’s shift to Robotruck services. Third-quarter revenues reached US$18.0 million, up 76% year-over-year and 35% quarter-over-quarter. The increase was primarily attributable to growth in freight transportation services, supported by our collaboration with Sinotrans.

In addition to the growth in our existing freight business, we are advancing Robotruck deployments in Jiangmen and at Mawan Port in Shenzhen, with commercial operations expected to generate ongoing transportation services revenues. As we expand into more ports, our project pipeline will continue to fuel top-line growth.

Lastly, Intelligent solutions revenues were US$8.9 million, broadly flat compared with the third quarter of 2025.

Let me now turn to gross profit and operating expenses. Gross profit increased by 69% to US$7.9 million as revenues grew. Gross margin was 17.8%, a slight improvement from 17.5% in the second quarter of 2026.

On the cost side, our safety record has helped us secure more favorable insurance terms. Annual insurance premiums for our Gen-7 Robotaxis were down to approximately 40% of those for human-driven taxis, compared with approximately 50% in 2025. While higher vehicle costs would normally push up insurance premiums because of higher repair costs, our safety record has more than offset this upward pressure. This translates into structural savings in fleet operating costs, strengthening our unit economics as we scale.

We also maintained expense discipline. GAAP operating expenses increased by only 3% year-over-year to US$76.6 million. Non-GAAP operating expenses, which excluded share-based compensation expenses, were down 2% to US$66.2 million.

Meanwhile, the increase in gross profit more than offset the modest increase in operating expenses, helping narrow our GAAP loss from operations to US$68.7 million, compared with US$69.7 million in the same period last year. The improvement was more pronounced on a non-GAAP basis, with non-GAAP loss from operations decreasing by 8% year-over-year to US$58.3 million, reflecting stronger financial discipline.

GAAP net loss was US$136.3 million, including a non-cash loss of US$77.0 million from changes in the fair value of trading securities. Excluding this fair value loss and share-based compensation expenses, non-GAAP net loss narrowed by 14% to US$48.9 million, from US$56.8 million a year ago.

Let me finish with our financial position and capital allocation. We maintain substantial financial resources to support technology development and commercial expansion. As of September 30, 2026, cash and cash

equivalents, restricted cash, short-term and long-term debt instruments for wealth management totaled approximately US$1.15 billion, compared with US$1.21 billion as of June 30, 2026, on the same basis.

Capital expenditures were US$17.7 million, a decrease of 22% year-over-year and 45% quarter-over-quarter. The decrease primarily reflected a disciplined pace of investment, while we continued to support the mass production and deployment of the Gen-7 Robotaxi fleet.

This quarter's revenue growth and improved operating performance reflect the progress of our commercialization efforts. Looking ahead to the fourth quarter and 2027, we will continue directing resources to core technology and key markets, translating our active pipeline and commercial progress into sustained revenue growth. At the same time, we will maintain expense and capital discipline to deliver greater operating leverage as our revenue base expands.

With that, I will now turn the call over to the operator for the Q&A session. Thank you.

Q&A

Tim Hsiao from Morgan Stanley

I have one question about overseas expansion. How does Pony stand out in overseas markets such as London, despite the tough competition with multiple players already in the area?

Dr. James Peng

Thanks for the question. This is James. Let me take this one. Yes, there are a lot of competitors in the European market, but I feel confident about our overseas market position, especially in Europe. Globally, when we look at it, outside China and the U.S., Europe is one of the most attractive Robotaxi markets, given its scale, pricing, and regulatory maturity. In that market, our competitive edge actually comes down to our proven technology, a strong safety system, and especially our very competitive unit economics. Because of our competitive edge, leading local mobility players want to work with us, which in turn can facilitate regulatory approval, allows us to scale faster, attract more partners and create more value across the whole ecosystem.

Let’s look at Croatia as a great example. We launched there in April this year, and by September we had already gone fully driverless. In just five months, we achieved Europe's first fully driverless passenger operations in more than 200,000 kilometers and thousands of passenger trips. Also in Zagreb, Croatia, that is a classic European city. Our Robotaxis actually operate in a mixture of scenarios with old-town roads, modern urban districts, and complex traffic conditions. This successful experience definitely can replicate to other European cities.

Now we are actually taking this playbook to the UK and other European countries. As Uber's partner, we are advancing deployments in London and Manchester, two of UK's most premium mobility markets. We actually started road testing in Manchester shortly after our announcement, which also shows that our execution speed is great.

With our proven safety record, large-scale operation experience, attractive unit-economics and strong partner network, we believe we are well positioned to be a leading player in Europe and especially in the UK. So I’m sure that we will capture a leading market share as the market matures.

With that, back to the operator.

Ming Hsun Lee from Bank of America

Hi, I also only have one question: Could you share some color on the operational and regulatory progress in China?

Dr. James Peng

This is James again. I'll take this question. I believe the current regulatory environment is by and large turning into a positive direction. Under the 15th Five-Year Plan, intelligent connected new energy vehicles are becoming part of the national policy agenda, targeting large-scale commercial deployment of autonomous driving for this Five-Year plan. This provides ongoing policy momentum at the central government to achieve that milestone, essentially have a regulatory base to shift Robotaxi from local implementation toward a more standardized national framework.

As this framework becomes more standardized, the safety bar certainly will be uniform and high for all cities across the whole nation. This is because all the other cities that want to operate Robotaxis, they need to follow the practices first established in tier-one markets. That will actually directly benefit operators like us who have already established proven technology, strong safety records, and real operating experience. We have already built that track record in China's most complex tier-one cities, and we are now scaling our active fleet, expanding coverage in the existing core cities, and entering new regions such as Hangzhou, Chongqing, and Zhuhai, with many other cities in the pipeline.

Therefore, my opinion is that as the industry enters a more standardized expansion phase, we believe Pony's first-mover advantage becomes even stronger. This track record and regulatory trust help streamline regulatory approvals, accelerate our engagement with both central and municipal authorities, and ultimately allow us to enter new markets faster and consolidate a leading position over time.

With this, back to the operator.

Yuqian Ding from HSBC

I got two questions. First, could you share with us on the update on Robotaxi unit economics? Second, how should we think about UE margins in domestic versus overseas markets as your fleet scales up in both places? Thank you.

Dr. Leo Wang

Thank you for your question. This is Leo. I'll take this one. I think UE is really, really important. Let me start with our China operation. We have already achieved positive UE margins in our Guangzhou and Shenzhen fleets. And with the network effects and also our continued efforts on cost optimization, the trend is actually keep improving. One direct example as we mentioned in this earnings call is about insurance: per-vehicle insurance premiums for our Gen-7 fleet have come down by another 19% this year.

So I want to emphasize that when we talk about UE, the revenue refers to the net value generated from ride-hailing services after deducting discounts and refunds. And also we cover seven cost items: first is the

depreciation of the vehicle, including the battery and also the autonomous driving hardware; second is the electricity charging; third is the salary we pay to the remote assistance and also ground supporting staff; fourth is the vehicle cleaning, service and the maintenance; fifth is the insurance cost; sixth is parking cost; and the final is the network communication cost. About half of the cost comes from the vehicle depreciation, and the other remaining half comes from the other six cost items. So when we talked about the UE, it's really a very comprehensive measure.

We're very proud that the Gen-7 platform gives us a very meaningful cost advantage, and because it is a fully automotive grade design, it is also helping us to lower our EV charging cost, and also helping us to lower the service and maintenance costs. In addition, our operational efficiency we believe remains to be industry-leading, backed by very competitive staff-to-vehicle ratios for both remote assistance and ground support staff. We will continue to improve our staff-to-vehicle ratio. This will keep labor cost per vehicle low even when the fleet keeps scaling.

And our overseas business carries a different structure, which is higher UE margin, driven by two factors. First, on the cost side — thanks to our stronger technology and optimized operations at scale — our total cost per mile remains very competitive in the whole industry. Specifically, our very high staff-to-vehicle ratio and also our superior energy efficiency give us a very good cost edge in some developed markets like Europe, where labor and electricity costs are very high. Second, the joint deployment model, which is asset-light for us: we don't bear vehicle capex. We offload depot infrastructure, EV charging, daily vehicle cleaning and maintenance, and other operational cost to our local partners, which would further improve our margin profile. We think the joint deployment model will become a high margin business and that can also scale.

With that, I'll get back to the operator.

Lawrence Ren from Citi

I only have one question on the AI side. As everyone in AI Industry is buying chips. With the recent news about H200 restrictions and the risk of more bans, where is your compute spend going? And how are you different from the rest of the industry to stay safe from these bans?

Dr. Tiancheng Lou

Thank you. This is Tiancheng. I can take this one.

In short, I will say we are not bottlenecked by AI compute. Let me break this down across inference and training part.

On inference, our Gen-7 Robotaxi has been running on automotive-grade chips since its launch last April. So these chips require much less computing power than general-purpose GPUs in data centers. We optimize the compute resource consumption of our on-vehicle models by optimizing model parameter size and improving the efficiency of our inference engine. So this delivers lower BOM cost, lower energy consumption, and increased driving mileage per full charge. That’s for inference.

On training, our compute dynamics are dramatically different from Large Language Models for two reasons:

First, our on-vehicle models are much smaller than Large Language Models, which means we don't need a huge cluster of cutting-edge GPUs to do pre-training or post-training. Most post-training compute going

toward running simulation inside our World Model — meaning training the driving model primarily relies on World Model inference, which requires far less compute.

Second, over the past few years we've kept making our world model iterate faster — we call it *PonyWorld* 2.0. It's self-evolving: it takes the road-test data coming back from our L4 fleet, works out where the sim-to-real gap is, and figures out which scenarios need better fidelity automatically. That cuts both the time and the compute we burn on training the world model by an order of magnitude.

In summary, our on-vehicle model architecture is more refined, and our engineering execution is stronger. That is what makes us the first company in the world to achieve L4 driverless driving on automotive-grade chips. Our per-vehicle BOM cost is highly competitive. The same strengths also make our data center compute more economical.

As I mentioned earlier, autonomous driving models fundamentally cannot be distilled. So competitors have to build the entire system from scratch. We also see that all the other major L4 players today rely on higher on-vehicle compute. That means larger model parameter size, and more compute consumed in training.

For an AI company, both model and engineering capability matter. I believe our technology lead will drive faster fleet expansion. It also gives us more flexibility when data center compute is in short supply. In addition, we have evaluated more than one China-made compute chips for model training, and we have successfully deployed our training pipelines. If needed, we can train our models on those chips.

With that, back to the operator.

Leo You from CLSA

My question is how do you think about diversifying the business? It seems that you are moving to more logistics fronts. Would you mind sharing more about that strategy? Thank you.

Dr. James Peng

This is James, and I'm glad to take this one. As most of you know, since our founding, our vision has always been autonomous mobility everywhere, and the logistics market sits at the center of our vision. Logistics itself is a massive market, and autonomous driving is a great solution for lowering down the costs and also filling labor shortages.

We leverage our L4 technology and operation across vehicle platforms and environments. We use exactly the same Robotaxi stack that we have developed over the years to power our Robotrucks and our Robo light trucks, across several scenarios, such as ports, long-haul corridors, and inter-city logistics.

We actually have already seen the results. Robotruck revenue grew 76% year-over-year this quarter. By collaborating with Sinotrans, we're deepening long-haul operations, such as the Beijing-Tianjin-Tangshan corridor. In ports, our Gen-4 trucks are in operation at Jiangmen and Mawan ports with an ongoing pipeline of hundreds of vehicles across leading global port operators. In the urban logistics, we provide our solution as light-duty Robotrucks. Compared with the more popular currently operating low-speed Robovan, our light truck offers 3 to 4 times the cargo capacity and 2 to 3 times the speed. So essentially, it is a great application for much higher-valued applications, such as supermarkets, cold chain logistics, and also the full urban supply chain. As a result, the potential for commercialization for long-haul trucks and light-duty trucks is actually enormous.

That's the power of our generalized platform — we can replicate, deploy, and monetize the same L4 capability across many different markets, which is exactly how we diversify. Of course, all of these efforts actually fit into our vision of autonomous mobility everywhere.

With this, I'll turn back to the operator.

Robin Leung from Bloomberg Intelligence

I have one quick question on the technology side. How do you think about companies plugging GPT-6 type models into embodied AI and eventually coming after the Robotaxi market? And if autonomous driving algorithms start to converge and look more like a commodity, how durable is Pony's moat and where does the advantage lie?

Dr. Tiancheng Lou

Thank you for your question. This is Tiancheng. I will take this one. In embodied AI, the industry SOTA still has not reached large-scale commercial deployment. That is why a general-purpose foundation model like GPT-6 can match it in selected scenarios — massive pre-training gives it broad world knowledge and enables impressive demos, such as robotics interactions and manipulations.

However, as I noted in my remarks, pre-training only establishes a baseline of "human-like" common sense — it does not eliminate long-tail hallucinations or cross the non-negotiable safety bar of the physical world. In L4 driving, over 90% to 95% of the core engineering happens in post-training. Operating on public roads means interacting with vehicles and pedestrians in real time, where there is no rollback and no retry. That's why a few-minute demo proves nothing about operating safely at scale.

Furthermore, training L4 autonomous driving directly on top of a large language model is fundamentally inefficient. Given the same compute budget, language is a slow and lossy intermediate representation for physical control. Even when humans drive, we rely on intuition and understanding rather than translating every perception and action into words.

This is also why L4 algorithms will not simply converge and commoditize. Reaching commercial scale requires mastering the post-training paradigm and closing the sim-to-real gap. *PonyWorld* provides a high-fidelity closed-loop reinforcement learning environment to scale post-training in simulation, while our fleet of driverless vehicles continuously feeds real-world data back to calibrate and sharpen *PonyWorld's* accuracy.

In short, foundation-model pre-training is only the starting line. Unlike digital Large Language Models with public APIs, our closed-loop post-training ecosystem cannot be distilled by other parties. Without a high-fidelity World Model for reinforcement learning and a large-scale driverless fleet feeding real-world edge cases back into that loop, no foundation model alone gets you to safe, commercial L4 autonomy. That post-training flywheel is what makes our advantage durable.

With that, back to the operator.