Orange Capital Partners

October 2026

Cumulative average annual return 15,5%. All returns after fees and commissions

Year	Q1	Q2	Q3	Q4	Year	Total
2026	-16,3%	15,8%	-1,5%		-8,1%	1332%
2025	-37,8%	22%	53,4%	0,5%	14,2%	1411%
2024	-16,6%	21,8%	22,9%	25,5%	87,8%	1223%
2023	69,3%	31,8%	-3,7%	1,95%	119,1%	746%
2022	-5,1%	-44,5%	16,8%	-53,2%	-71,2%	286%
2021	-3,7%	12,5%	1,7%	13,8%	25,2%	1240%
2020	1,8%	60%	63,7%	65,9%	335,1%	871%
2019	4,9%	2,9%	3,3%	21,6%	35,5%	142%
2018	9,1%	-12,7%	-8%	-7%	-18%	78%
2017	11,9%	16,3%	4,4%	10%	50%	117%
2016	-12,9%	-4,3%	11,3%	-10,5%	-16,2%	45%
2015	4,9%	5%	-20%	-7%	-17,4%	74%
2014	0,4%	-7,6%	3,9%	1,7%	-1,8%	111%
2013	6%	32,5%	2,1%	-6,6%	28%	115,2%
2012	0,3%	7%	-4,8%	-2,1%	0,2%	68,2%
2011	15,5%	-11,3%	-12,2%	-0,4%	-10,3%	68,2%
2010	4,7%	-1,5%	17,3%	17%	33,2%	87,8%
2009	1,5%	3,2%	7,4%	5,4%	14,9%	41%
2008		-7,5%	12%	24%	22,8%	22,8%

Strategy

We invest in fundamentally new concepts and engineering practices with large impact. Areas of focus are autonomous electric vehicles, robots and emerging compute platforms to accelerate the advent of superintelligence.

Recent developments

We are “bitter-lesson-pilled” when it comes to investing in real-world superintelligence.

Tesla’s supervised Full Self-Driving (FSD) is the first real-world superintelligence system operating at scale and used by paying customers. In Q2 2026 the company confirmed that vehicle demand is being driven by FSD availability in markets where regulators have approved it. Unfortunately, major markets such as Germany and China continue to withhold approval, largely to protect their domestic auto industries. While this is regrettable, the resistance itself is a telling sign of the technology’s superiority: it is a tacit admission that competitors remain far behind and that customers would readily adopt FSD if given the chance. In the United States, attach rates already exceed 50 percent. More recently, German politicians have signaled readiness to approve FSD, and several EU countries—including the Netherlands, Slovenia, and Croatia—have already done so.

FSD was trained on real-world video data and refined through simulation—the same approach Tesla is applying to Optimus. We find this strategy promising. At the recent Berkeley Agentic AI Summit, competitors including Nvidia discussed the same methodology. The robotics training pipeline will follow the FSD playbook: first pre-train on internet-scale video of humans performing everyday tasks, then fine-tune through simulation and robot academies. Pre-training on vast video datasets gives robots an understanding of depth and spatial reasoning; fine-tuning in simulation and dedicated facilities then teaches them specific skills.

The ultimate goal is a continuous learning pipeline in which robots improve from real-world experience. The traditional train–deploy–retrain cycle works, but it is inefficient. Learning directly from experience is a far more powerful paradigm.

As the saying goes, in the short run people overestimate the impact of AI; in the long run they underestimate it.

General developments in superintelligence

The most relevant advances for our portfolio lie in robotics infrastructure, spanning both hardware and software. Ashok Elluswamy recently gave a talk describing Tesla’s approach to end-to-end real-world AI. In the talk he highlighted the use of generative superintelligence to simulate scenarios based on real-world videos, allowing FSD to be calibrated against edge cases. Tesla’s method is rooted in end-to-end modeling combined with large data pools. In short, the company is moving as many scenarios as possible “in distribution” to make FSD robust, safe, and convenient. The latest version, FSD 14.3+, has achieved these goals and continues to improve with each iteration.

We anticipate that superintelligence software development will be driven by a combination of academic research, startups, and established companies. On the hardware and manufacturing front, however, Tesla is poised to extend its lead. Ultimately, robots will be judged by the same metrics as self-driving cars: utilization per dollar. The longer robots can operate and the lower

their production costs, the lower their cost per hour of work. We view Tesla as a leader in both hardware and software innovation. What ultimately matters is lower cost, higher value, and higher customer satisfaction. As long as Tesla optimizes for these metrics, sales and cash flows will follow.

Another critical driver of robotics infrastructure is the rapid advance of simulation technology, led by NVIDIA, Tesla, and others. While simulation is not new, today's infrastructure enables realistic reproduction of physical systems. NVIDIA in particular is pioneering transformative tools, notably Isaac Gym and the Omniverse platform.

Isaac Gym is a simulation engine tailored for reinforcement learning. We see reinforcement learning as the most promising approach, because it enables self-play mechanisms similar to AlphaGo Zero and thereby fosters autonomous skill development. Omniverse is to physical simulation what CUDA is to GPU computing: a foundational platform that enables broad, general-purpose applications. Over time it will drive adoption, attract developers, and foster a global ecosystem. NVIDIA excels at two critical capabilities: designing highly complex technology, and galvanizing the global developer community to adopt its platforms, creating a flywheel of innovation.

Tesla

Robotaxi is launching more slowly than expected. Two factors are driving the delay. First, Tesla wants to be fully ready to scale Cybercab, which will serve as the workhorse of the robotaxi fleet thanks to its lower cost and greater convenience. Second, FSD Version 15 is expected to be meaningfully better than Version 14, and the company prefers to scale with the stronger model. While the timeline pushback reduces the stock's net present value, Tesla remains on track to launch a large-scale, general-purpose robotaxi solution across the United States and, later, worldwide. Transportation remains an expensive service today—a cab from JFK to Manhattan can cost nearly as much as the flight itself—so robotaxi has the potential to become a high-cash-flow business.

High capital expenditure is unsettling the market. In Q2 Tesla reported negative free cash flow for the first time in years. We view the elevated capex as problematic not because of the spending itself, but because cash-flow generation from robotaxi is being delayed. Ideally, Tesla would be scaling cash flow from robotaxi, energy, and Semi (its newer business lines), continuing to grow the core vehicle business, and expanding FSD subscriptions—all while funding new initiatives such as Optimus, Cybercab production, AI infrastructure, Terafab, and solar. That balance has not yet materialized. Tesla is investing heavily in capex, but operating cash flow is lagging, which is weighing on the stock.

Another concern is the apparent lack of urgency around the robotaxi business. When Tesla was ramping the Model 3, Elon was sleeping on the factory floor. The problem was fixed and the stock rose tenfold. Robotaxi presents a challenge of similar magnitude and demands comparable intensity. Getting the technology right and overcoming regulatory opposition, misleading coverage, and other sources of fear, uncertainty, and doubt is not only critical for Tesla's cash

flow; it would also benefit society. The more widely FSD is deployed on public roads, the safer those roads become.

Optimus

There is strong evidence that pre-training on video data can teach robots spatial reasoning—an essential capability for a general-purpose humanoid. Rather than hard-coding priors (such as "avoid trees while driving" or "do not harm humans"), robots learn statistical relationships between objects from internet-scale video in an end-to-end architecture: pixels in, actuators out. Humanoids make particular sense because humanity has already generated billions of hours of video featuring the human form factor. Learning from that vast corpus is an excellent way to harvest the bitter lesson in robotics. We are encouraged by recent progress in robotics, computer vision, and real-world reasoning. This year's CVPR conference in June 2026 showcased substantial advances in both robotics and spatial understanding.

Ultimately, we see Richard Sutton's latest essay, "The Era of Experience," as the path forward. Architectures will emerge in which embodied agents continually learn from real-world experience.

Nvidia

The progression of superintelligence—from computer vision (AlexNet) to generative AI (ChatGPT), reasoning models such as o1 and DeepSeek, and most recently agentic workflows—has been built on enormous surges in data, computational power, and model infrastructure. Nvidia stands at the heart of this evolution. The risk remains to the upside, given the unprecedented nature of these intelligence breakthroughs. What scale of value creation lies ahead? We must avoid reasoning by analogy. Instead, consider a framework in which AlphaGo Zero's principles could be applied to virtually any global challenge. This represents an immense opportunity. In a $100 trillion global economy, a 10 percent boost in productivity would generate $10 trillion in new value—more than the GDP of Germany—heralding a profound shift in worldwide wealth creation. Silicon Valley is at the center of this transformation.

Recently, Nvidia engaged in somewhat questionable negotiations with OpenAI over vendor financing. Our concerns proved warranted, and CEO Jensen Huang backtracked, which is a positive sign. Long-term, we see enormous potential for Nvidia as AI infrastructure and the emerging space-AI race foster demand for sophisticated AI solutions, solidifying Nvidia's lead in this area.

Risks lie ahead. Leading superintelligence companies such as OpenAI and Anthropic want to go public, yet the market is increasingly uncertain about their long-term revenue and earnings potential. First, token maximization is being replaced by token optimization. Customers want the most efficient tokens, which may or may not come from the leading intelligence providers. Legitimate questions surround their business models. Second, open-source efforts such as Reflection.ai, Nvidia's Nemotron model family, and competition from China and Europe are putting pricing pressure on tokens. SpaceX SI is a cost leader and offers lower-cost tokens, albeit allegedly at lower intelligence value. Third, adoption of superintelligence will follow the classic

innovator's dilemma. It will not be adopted by legacy companies but by new startups that disrupt existing businesses. Banks will not become superintelligence companies; superintelligence companies will become banks. This pattern will apply across sectors. Eventually, there will be as many companies in the economy as there are models.

The first half of 2026 has been a relatively pedestrian period for Nvidia. While the company has delivered solid results, many other semiconductor stocks—such as AMD, Intel, and Arm—have significantly outperformed it. In fact, most of the upstream layers of the AI stack (energy, chips, cloud, models, and applications) have posted much stronger gains, particularly in optical networking, memory, and energy.

So why do we continue to hold Nvidia? Because we believe the company remains focused on innovation and iteration rather than simply raising prices. For example, demand for CPUs is surging as agentic workflows require complex sequential coordination. Yet the current AI stack remains highly inefficient, with CPU utilization for agent workloads often below 20 percent. Nvidia is positioned to address this challenge with its next-generation platform, Rubin.

Drivers of long-term wealth

Wealth creation is predicated on maximizing X per dollar, whatever that X may be. Tesla is optimizing for miles per dollar and hours worked per dollar, while Nvidia's key advantage lies in optimizing for tokens per dollar—the backbone of intelligence. The goal is to produce large volumes of something people desperately want at a lower cost than anyone else. That is what wealth creation actually means.

In February 2026 Elon Musk announced the merger between xAI and SpaceX. This is the first step toward a fully integrated intelligence powerhouse. Tesla will play a pivotal role in this system, because its robot business enables the monetization of intelligence at scale. Someone has to produce all the chips and robots needed to serve billions of humans with real-world AI. That is going to be Tesla's role. Whether this happens within a combined SpaceX and Tesla company or as separate entities is secondary. As Tesla shareholders, we prefer Tesla to deliver on the Robotaxi and unsupervised FSD launches in 2026 and to drive value through execution and cash-flow generation. A subsequent merger would be beneficial but is not necessary.

We do believe that a fully integrated intelligence powerhouse combining Tesla, SpaceX, and xAI is advantageous, because most of the problems in this value chain are novel and therefore require rapid iteration. This can best be achieved with full integration.

Portfolio Statistics			
Largest Long Positions	Largest Short Positions	Beta with S&P 500	2,1
TSLA		Leverage	10%-25%
NVDA		Target Return	15%, net of fees