

Cumulative average annual return 15,8%. All returns after fees and commissions

| Year | Q1            | Q2         | Q3           | Q4          | Year   | Total  |
|------|---------------|------------|--------------|-------------|--------|--------|
| 2026 | <u>-16,3%</u> | 15,8%      |              |             | -3,7%  | 1355%  |
| 2025 | <u>-37,8%</u> | <u>22%</u> | <u>53,4%</u> | <u>0,5%</u> | 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 artificial intelligence.

## Recent developments

### ***We are ‘bitter lesson pilled’ when it comes to investing in real world AI.***

Tesla’s supervised Full Self-Driving (FSD) is the first real-world AI system operating at scale and used by paying customers. In Q2 2026 the company confirmed that demand for its vehicles is driven by the availability of FSD in markets where regulators have approved it. Unfortunately, key markets such as Germany and China continue to withhold approval, largely to protect their domestic auto industries. While this is regrettable, the regulators’ resistance is itself 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%, which is encouraging.

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 robot facilities 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–repeat 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 AI*

The most relevant advancements for our portfolio lie in robotics infrastructure, encompassing both hardware and software. For instance, Ashok Elluswamy recently gave a talk describing Tesla’s approach towards end to end real world AI. In this [clip](#) he highlights their approach towards using generative AI to simulate scenarios based on real world videos. This way the FSD can be calibrated to solve edge cases. Tesla’s approach is rooted in using end to end modeling coupled with large data pools. In short, Tesla 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 is getting better with each iteration.

We anticipate that AI software development will be driven by a combination of academic research, startups, and established companies. However, on the hardware and manufacturing front,, Tesla is poised to expand its lead. Ultimately, robots will be evaluated 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. Ultimately what counts is *lower cost, higher value and higher customer satisfaction*. As long as Tesla optimizes for these metrics, the sales and cash flows will follow.

Another critical driver of robotics infrastructure is the rapid advancement 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 technologies, namely **Isaac Gym** and the **Omniverse** platform.

Isaac Gym is a simulation engine tailored for reinforcement learning (RL). We see RL as the most promising approach, as it enables self-play mechanisms similar to AlphaGo Zero, fostering autonomous skill development.

Omniverse is to physical simulation what CUDA is to GPU computing: a foundational platform enabling broad, general-purpose applications. Over time, it will drive adoption, attract developers, and foster a global ecosystem. NVIDIA excels at two critical capabilities: first, designing highly complex technology; second, galvanizing the global developer community to adopt their 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 meaningfully better than Version 14, and the company prefers to scale with the stronger software. 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 is still 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 FUD (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 AI—from computer vision (AlexNet) to generative AI (ChatGPT), reasoning models like o1 and DeepSeek and most recently agentic work flows —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 through analogies. Instead, consider a framework where 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% boost in productivity would generate \$10 trillion in new value (which is 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 shady negotiations with OpenAI for vendor financing. Our concerns were warranted and CEO Jensen backtracked, which is a good sign. Longterm we see enormous potential for Nvidia as AI infrastructure and the newly emerging Space AI race will foster demand for sophisticated AI solutions solidifying Nvidia’s lead in this area.

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 areas like 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. However, the current AI stack remains highly inefficient, with CPU utilization for agent workloads often below 20%. Nvidia is positioned to address this challenge with its next-generation platform, Rubin.

Drivers of longterm wealth

Wealth creation is predicated on maximizing X/\$, whatever that X may be. Tesla is optimizing for miles/\$ and hours worked/dollar, while Nvidia’s key advantage lies in optimizing for tokens/dollar, which is the backbone of intelligence. You want to be able to produce large volumes of something that people desperately want at a lower cost than anybody else. That’s what wealth creation actually means.

In February 2026 Elon Musk announced the merger between X.ai and Space X. This is the first step towards a fully integrated intelligence power house. Tesla will play a pivotal role in this system, as Tesla’s robot business enables monetization of intelligence at scale. Somebody has to produce all the chips and robots to serve billions of humans with real world AI. That is going to be Tesla’s role. Whether that happens within a combined Space X and Tesla company or separate is secondary. As Tesla shareholders we prefer Tesla to deliver on the Robotaxi and unsupervised FSD launches in 2026 and 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 power house with Tesla, Space X and X.ai is beneficial as 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 |