2024 REPORT

Building AI-Powered Products

The Enterprise Guide

BROUGHT TO YOU BY

V7

EXPLORE V7 →

Index

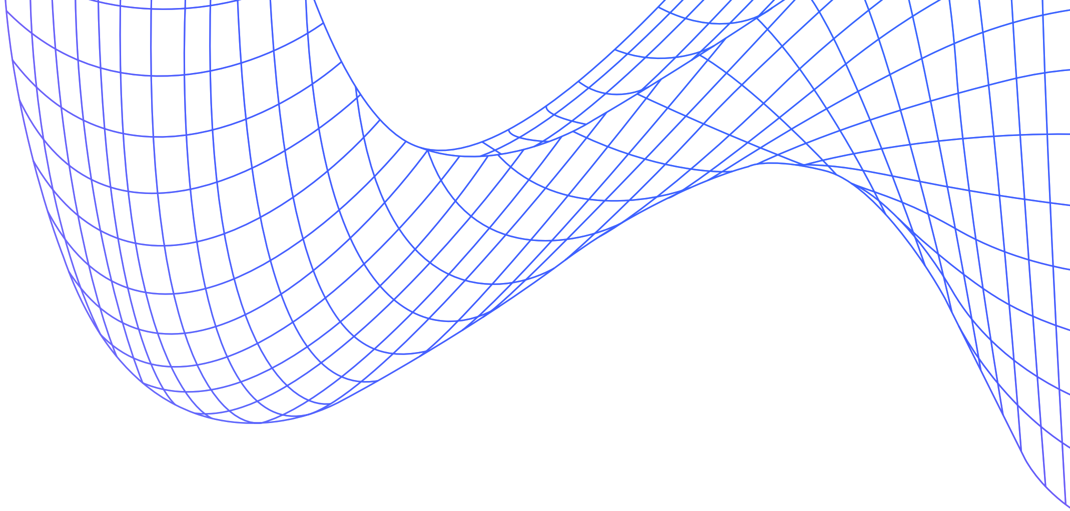

WHY BUILD?

The Case for AI-Powered Products

With the rise of ChatGPT, AI-generated images, and incumbent co-pilots, AI has well and truly gone mainstream. And yet, unlike other tech trends from the start of the decade, AI drives far deeper than its hype, delivering measurable business impact and value to users.

With the power to revolutionize industries and reshape human-computer interaction, the impacts of AI are already being felt globally. Development times are being halved, proofs of concept are realized at breakneck speeds, and companies are racing to compete in who builds an AI-first product first. So much so, that the AI industry is expected to grow 37% yearly, for the next seven years.

Put simply, AI has lit a fire beneath enterprise businesses—and the time to move is now.

Despite this sense of urgency, there are pitfalls along the way, many of which define the difference between a successful product, and an expensive failure.

In this guide, we break apart the process of building an AI product and provide you with a clear outline of what you'll need to address.

From the quality of your training data to establishing product-market fit, consider this guide as a handbook for your AI-powered product development process.

What Can You Expect From This Guide?

Getting Started With AI

While AI promises a host of rewards, the road to reaping them can be complex. In this section, we tackle the impact of AI on enterprises, emerging industries, and key challenges that lie in wait.

Building AI Into Your Product

What are the essentials of a well-built AI-powered product? What's the process? And what is the development strategy of AI-product success stories? We outline key components of building an AI-powered solution.

Your Strategy Recap

A companion checklist to take with you throughout your AI product journey, ensuring you stick to best practices at every stage.

GETTING STARTED

Setting the Scene of Enterprise AI

Dubbed the "AI Gold Rush", we're in perhaps the biggest AI hype cycle we've seen so far.

As a result, a battle has begun to realize the possibilities of AI, with incumbents and startups alike clambering to stake their claim. Much like the tech boom of the 90s, there will be winners, and there will be losers.

In this section, we tackle the impact of AI on enterprises so far, outline some of the challenges to AI adoption, and discuss the emerging use cases preparing to take the spotlight.

The Impact of AI on Enterprise

Revenue generated from AI continues to grow year on year, with a global market value of $ 100 billion dollars that's expected to balloon. While that figure alone is impressive, let's explore some emerging market trends within enterprise AI.

AI: THE TECH OF CHOICE FOR THE C-SUITE

Accenture states that 73% of companies in 2023 are prioritizing AI over all other digital investments, with 90% of business leaders applying AI to tackle aspects of operational resilience.

Enterprise businesses are the most likely to embrace AI, with 68% of large businesses adopting at least one AI technology in 2022. Of those large companies, 20% used at least four AI technologies to bolster their day-to-day operations.

AI: PRODUCTIVITY AND PROFITABILITY FUEL

While AI is set to automate a series of professions, the World Economic Forum states that an estimated 97 million new jobs will be created by 2025 thanks to AI. Businesses will benefit from a labor force empowered to focus on more meaningful work, directly impacting productivity and profitability rates. In fact, the overwhelming use case of labor-focused AI thus far has not been to replace human labor but to maximize it and convert it into meaningful outputs.

In this same vein, AI's impact on labor productivity could raise global GDP by 7%, according to research conducted by Goldman Sachs. 40% of these economic gains will come from AI-powered product enhancements, according to research from PWC.

AI: CONTROVERSY AND RISK

Amidst the successes of the AI industry, controversy has reared its head. Concerns have arisen around the legal liability of AI, lawsuits have surfaced from the artistic community - lambasting creations from Midjourney and Stability AI - while some leaders have stated that AI poses a "risk of extinction" for human society.

While calls have been made to halt AI production, others have highlighted the rapid developments happening as a result of AI for good. In May of 2023, two AI breakthroughs hit the headlines - the first being the discovery of a new antibiotic to combat drug-resistant bacteria, and the second being the use of AI to enable a paralyzed man to walk again.

Against the backdrop of anti-AI rhetoric, a more nuanced take has emerged, highlighting both the regulatory needs and revolutionary capabilities of AI. In response, and with the acknowledgment that AI is here to stay, regulators worldwide have scrambled to keep up. Similarly, governments have rapidly sought to capitalize on this growing market, making moves to empower AI innovation and clear the route to its success.

"It's natural to wonder if there will be a jobless future or not. What we've concluded, based on much research, is that there will be jobs lost, but also gained, and changed. The number of jobs gained and changed is going to be a much larger number, so if you ask me if I worry about a jobless future, I actually don't. That's the least of my worries."

James Manyika
Chairman and Director of the McKinsey Global Institute
"Salesforce Sessions: The Future of Work in the AI Age"

Which Industries Most Commonly Use AI?

While the impacts of AI have been felt across every conceivable industry, there are a number of sectors predicted to leverage AI intensely in the near future. According to IDC's 2023 Infobrief, banking, retail, professional services, and manufacturing will account for over half of global IT spending on AI in 2026.

"Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years."

Andrew NG
Co-Founder of Google Brain
"Stanford Future Forum Sessions"

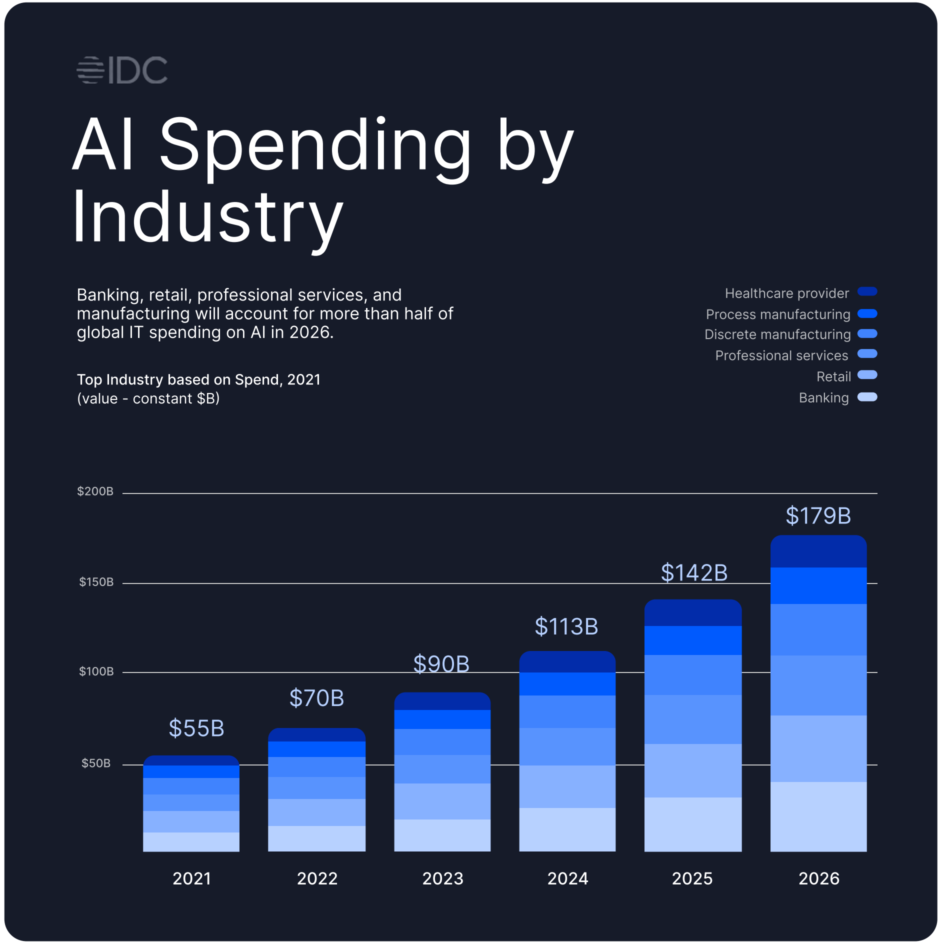

Let's take a closer look at the industries expected to flourish with the support of AI →

EXPLORE V7 USE CASES →

AI INDUSTRIES

Healthcare

The AI healthcare market is forecast to be worth almost 188 billion U.S. dollars by 2030, with an annual growth rate of 37%. Thanks to AI's ability to resolve human error, bolster human labor, and commit to repetitive tasks, healthcare has had a series of breakthroughs with the capacity to better the lives of millions. From AI-enabled prosthetics to physician copilots, the AI healthcare market is booming, driven by high-quality training data and state-of-the-art annotation platforms.

We asked Ankush Khona, a senior healthcare technology leader with over two decades spent advancing healthcare, by leveraging data, AI, and analytics, the greatest challenges he's faced in AI product development - and how he overcame them.

"Healthcare data often exhibits intricate patterns, correlations, and interdependencies that are difficult to fully understand. Access to quality data and ensuring that the data accurately represents the diversity and complexity is a significant challenge.

To overcome this, we collaborated closely with healthcare professionals, domain experts, and clinicians to ensure the test data captures the complexity, variability, and nuances of real-world data. Involving experts in the design and validation of data can help enhance its quality and realism. My advice to other teams embarking on AI development? Develop a comprehensive data and AI strategy that's aligned with your organization's goals, and identify specific use cases, prioritizing them based on their potential impact and feasibility

Next, address ethical considerations related to AI, such as bias, fairness, transparency, and accountability. Make sure you develop policies and guidelines for responsible AI use.

Finally, establish mechanisms for continuous improvement and evaluation of your AI initiatives. Monitor and analyze the impact of implemented AI solutions on key performance metrics, collect feedback from users, and iterate on the models and processes to drive better results."

Ankush Khona

EXPLORE V7 USE CASES →

For those that invest in this rigorous approach to AI development, a number of rewards await...

"Organ non-use due to size mismatch is a significant issue in transplantation. With V7, we were able to achieve 97.0% model accuracy. V7 makes it possible to calculate the exact size of the organ and do a side-by-side comparison of the donor and recipient organs. It's a real paradigm shift".

Dr. Rowland Petit
Chief Science Officer @ InformAI

InformAI
Advancing Healthcare Through Informatics

MODEL ACCURACY WHEN USING V7
97%

COMPANY INDUSTRY
Healthcare

TYPE OF DATA
DICOM

FAVOURITE FEATURE
S3 Integration

AI INDUSTRIES

Agriculture

With the rise of agri-food solutions and advancements in smart farming, the agriculture industry is reaping the benefits of AI. The industry has boasted great successes across vertical farming, crop protection solutions, and autonomous harvesting robotics - and shows no signs of stopping. It's perhaps no surprise then, that the Agriculture market is expected to reach 7.43 billion by 2030.

Securing solutions to food scarcity, disease, and climate crisis concerns, this industry promises to solve the immediate and future risks facing the globe at large.

> "We needed a tool that could do annotating and data versioning because we distribute our tools to farms, and we need to make sure that they have the same version of data for the same models. V7 met our needs."
>
> Raymond Tunstill
> PhD Researcher, CTO of FruitCast

UNIVERSITY OF LINCOLN

MODEL ACCURACY 95%	COMPANY INDUSTRY Agriculture
YIELD PREDICTION TIME 5 weeks	FAVORITE FEATURE Data Versioning

EXPLORE V7 USE CASES →

AI INDUSTRIES

Energy

The energy market has gone from strength to strength in recent years, with environmental concerns driving many of the much-needed innovations within this space. As businesses race to reduce fossil fuels, deliver renewable energy sources, and equip developing regions with energy, AI has quickly become a crucial component of reaching a net-zero world.

Projected to reach a market valuation of $19.8 billion in 2031, this space is already measuring greenhouse gas emissions with AI-enabled drones, delivering energy-efficient smart grids, and maximizing renewable energy outputs.

> "Annotation systems are playing catch up trying to project where research and cutting-edge development would be and how to make it configurable. Luckily, solutions like V7 came up, which are maturing in the industry, helping companies like ours to scale up and stay ahead with our R&D."
>
> Suchet Bargoti
> CTO, Abyss Solutions

ABYSS

DATA LABELED 2TB+	COMPANY INDUSTRY Energy
TYPE OF DATA 2D/3D, thermal data	ENGINEERING TEAM SIZE 100+

AI INDUSTRIES

Finance

Fintech has had an extraordinary few years, thanks to the rise of neobanks, cybersecurity advancements, and the digitization of global banking infrastructure. With the rise of AI, the rate of innovation within the finance industry is expected to keep accelerating.

Estimates indicate that the global AI in Fintech market size will reach $61.30 billion by 2031, growing at an annual growth rate of 22.5%. Advancements are expected across fintech-focused chatbots, algorithmic trading, and anti-bank fraud solutions.

> “Our Large Language Models will enable us to tackle many new types of applications, while it delivers much higher performance out-of-the-box than custom models for each application, at a faster time-to-market.”
>
> Shawn Edwards
>
> Chief Technology Officer for Bloomberg
> ”Introducing Bloomberg GPT”

AI INDUSTRIES

Retail

Finally, AI poses an incredible opportunity for the retail industry, from personalized shopping experiences to inventory management and demand forecasting. Having already undergone a digital transformation, fueled further by the COVID-19 pandemic, retail is now uniquely poised to go one step further. Estimates indicate the market value of retail AI to reach $85 billion in 2033.

EXPLORE V7 USE CASES →

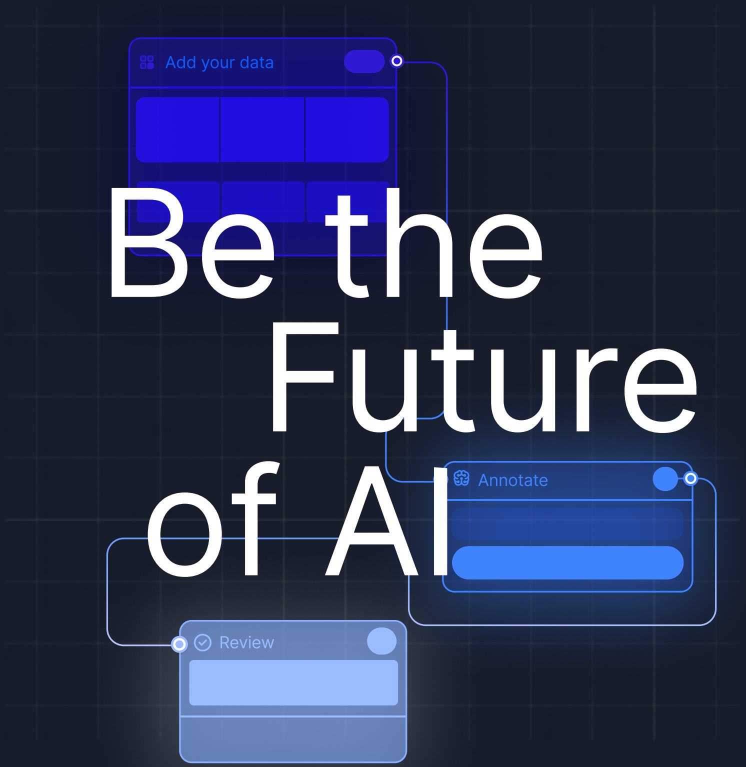

The Challenges of Enterprise AI

At the enterprise level, AI-powered product development has a series of complex hurdles. While adoption rates of AI within enterprises far exceed that of startups, the fact remains that AI development poses a series of hurdles.

We've tackled a few make-or-break factors for any AI-powered product, and outlined some of the common pitfalls that lie in wait →

According to IBM's Global AI Adoption Index, the biggest barriers to AI adoption are:

While the IBM AI in the Enterprise report reveals that while leaders believe AI will digitally transform their business, they also believe:

COST-EFFECTIVE AI ADOPTION

AI development costs can vary wildly, quickly spiraling out of control when a development pipeline is marred by inefficient tools, lacking infrastructure, and needlessly ineffective processes. Costs will spread across hardware, software, labor, training, and maintenance, with the unpredictability of the process often turning development into a financial black hole.

While AI can deliver phenomenal profits and cost-saving measures, for those unable to clearly define the ROI of their efforts, stakeholder buy-in will continue to be a challenging task.

Despite these challenges, solutions are coming to the fore, with the cost of AI adoption decreasing year on year. In fact, according to the ARK Invest Big Ideas 2023 report, the cost to train an LLM similar to GPT-3 has reduced from $4.6 million in 2020 to $450,000 in 2022 - a decline of 70% per year. Similar decreases are occurring within hardware, thanks to innovations from IBM and NVIDIA - making AI adoption a more feasible option for those once barred by eye-watering costs.

Finally, platforms for training data and model management have evolved to allow enterprises to clearly track, manage, and measure their development pipeline - making it much easier to control project costs.

With the right infrastructure in place, enterprises can now review the performance of annotators, assign the most crucial tasks to experts, and automate parts of the training process with highly trained models that won't compromise on quality. Similarly, these platforms have evolved to maximize the value of onsite machine learning engineers, while limiting the degree of training needed to create high-performing models. With these systems in place, reducing the cost of AI adoption is immediate, while showcasing ROI becomes an intuitive part of the process.

For those that go it alone, restricted funding for AI development, over-scrutiny, and ill-fitting solutions will continue to grind innovations to a halt.

Our advice for enterprise? Carefully explore the solutions on the market, and secure the best one for your business. An inability to find them will fail your project before it even starts.

Later in this guide, we'll explore the pros and cons of building your own AI infrastructure, versus using an out-of-the-box option. Before we dive into the next looming challenge, let's recap with a number of ways enterprises can promote cost-effective AI adoption.

How to Reduce the Cost of AI Adoption	V7
Leverage streamlined infrastructure that will minimize costs throughout the AI development lifecycle.	✓
Optimize the labeling process with annotation features that prioritize speed and accuracy.	✓
Maximize AI engineer productivity by ensuring your people are focusing on tasks of value, rather than manual or infrastructure maintenance work.	✓
Avoid high product failure rates with reliable QA processes that won't exhaust time and resources.	✓
Avoid a slow training data process by leveraging platforms that can handle datasets with complex/dense annotations, domain-specific file formats, and can scale as your project grows.	✓
Prioritise platforms that allow for experimentation, and rapid creation of MVPs and POCs, to secure meaningful funding for AI initiatives.	✓
Train and utilize models that can automate elements of annotation and QA.	✓
Leverage out-of-the-box solutions that ensure engineers are only focusing on crucial tasks.	✓
Empower managers to assess the performance of annotators, reserve the most crucial tasks for expert labelers, and balance precision with cost-effectiveness.	✓
Ensure your AI tooling vendors sets your team up to succeed, with robust customer support, valuable documentation, and the guidance needed to squeeze every last ounce of value from your training data platform.	✓

PACE OF INNOVATION

Perhaps one of the more unique aspects of AI-powered product development is it's ability to move at breakneck speeds. However, for enterprise customers, therein lies a great challenge. For those that fail to rapidly innovate, short product lifespans, high customer churn rates, and obsolete products become a looming possibility.

"If you look at it at a micro level, you used to do credit card fraud without AI, now you do it with AI. You used to do customer service without AI, now you do it with AI. So within each silo it will have incremental benefits. But I think the biggest benefit will come when you rethink the entire business model."

Dr Kai-Fu Lee
Former President of Google China
"Boston Consulting Group: How AI Will Reshape Companies, Industries, and Nations"

Below are just some of the common issues enterprises face when trying to innovate at pace:

- Shortage of AI talent.
- Inaccessible or poor-quality AI training data.
- Lacking quality controls that produce inaccurate models.
- Ineffective tools or platforms that slow down the process.
- Data is too complex to manage.
- Costs are prohibitive.

For businesses seeking to realize the benefits of AI, it's become clear that it's essential to build an AI product pipeline that considers cost, training data, accuracy, ease of use, and speed. Failure to do so seems to guarantee late launches at best, and product flops at worst.

In fact, according to Gartner, approximately 70-80% of AI projects fail, thanks to poor quality data, an inability to manage it, and lacking infrastructure adequately equipped to extract value from this data.

How to Improve the Pace of Innovation	V7
Speed up model iteration cadence with tooling that can balance speed and precision.	✓
Increase your team's output velocity by ensuring they focus on crucial tasks, automate components of the process, and reach accuracy with rapid tooling.	✓
Increase collaboration across stakeholders with intuitive dashboards and workflows.	✓
Manage the entirety of your company's training data corpus in one spot, with UI that makes it easy to extract data, organise datasets, and scale.	✓
Leverage existing frameworks and technologies to get your team up and running faster.	✓
Prioritize infrastructure that makes integration simple and secure, allowing you to rapidly develop building blocks for your product development process.	✓

DE-RISKING AI DEVELOPMENT

AI development can be vulnerable to risk, from cyber security concerns to bias and ethical challenges. In the context of enterprise, a sprawling scope of risk will halt production before it even begins. Scrutinized across legal, compliance, infosec, and sustainability measures, de-risking AI is a crucial step for enterprises. Failure to do so can result in reputational damage, legal liability, eye-watering fines, and faulty - at best - products.

With EU legislation like the "Artificial Intelligence Act" to contend with (carrying a maximum fine of 30 million euro or 6% of annual worldwide turnover) the answer is clear: de-risk, or face the consequences.

"Companies who use AI responsibly will have a competitive advantage, gaining the trust and loyalty of their customers, their shareholders and all stakeholders. Those that don't will hurt their brand reputation and, the health of their businesses."

Maria Luciana Axente
Globally Recognised AI Ethics Expert
"The Innovator: How Companies can Ensure They are Implementing AI Responsibly"

Fortunately, enterprise-ready architecture exists for secure AI development, with built-in data explainability, security, and governance structures. While the temptation may be to curb costs with cheaper, or free tools, the fact remains that these tools are often prone to security risks. More mature platforms like V7, provide rigorous infrastructure that de-risks development by design.

In the words of V7's Co-Founder, **Alberto Rizzoli**:

"Time and regulatory barriers are the two biggest risks in AI development. Both are mitigated by a well-run build vs buy process, offloading regulatory compliance to mature vendors and speeding up the development process by focusing internal resources on what the company will have an unfair advantage on".

As you prepare to dive deep into development, you'll need to strategize against risk, embed trustworthy controls, and house your product in architecture that shields against common enterprise risks. Without this, you'll struggle to justify the development of your products from the get-go.

How to De-Risk AI Development	V7
Leverage infrastructure that prioritizes data explainability and good governance structures.	✓
Track and visualize your data labeling and model development progress, with a UI that makes it simple to identify bottlenecks.	✓
Integrate your project within a wider AI ecosystem, sharing resources and policies, with safety in mind.	✓
Improve team synchronization with a standardized approach to AI development, in addition to collaborative workflow processes.	✓
Ensure data and processes are interoperable.	✓
Prioritize compliance and risk management, with infrastructure that considers GDPR, FDA, HIPAA, ISO27001, SOC 2, etc.	✓

Best-in-Market Products

The market is flooded with AI-powered products, each vying for the attention and loyalty of the consumer base at large. Your products face an uphill battle to secure market dominance, reliable profit margins, and customer satisfaction - compounded by the rate of change within AI development.

So, who comes up top?

Overwhelmingly, best-in-market products leverage best-in-market tooling, cutting-edge frameworks, highly responsive UI, and processes that can reliably deliver better products.

Platforms like V7 give enterprises a market advantage by accelerating their model development with automated workflows that find consensus, debug, retrain, and ship into production with full visibility and accuracy. Brand-damaging recalls and sky-high margins are cut, in favor of bulletproof innovation that shuts out the competition.

How to Build Best-in-Market Products	V7
Prioritize infrastructure that will rapidly inject new knowledge into AI, and secure first-to-market benefits.	✓
Prioritise fit-for-purpose features, like AI and human workflows, powerful annotation tools, and effective dataset management.	✓
Develop with certainty and confidence in your data quality and accuracy, with enterprise-ready infrastructure.	✓
Ensure the test and validation process for your trained model is measurable.	✓
Accelerate the improvement and reliability of your data with model integrations, performance visualization, and error reduction features.	✓

GETTING STARTED

What Do Enterprise Customers Want?

The success stories of the AI Gold Rush have overwhelmingly come from the incumbents of the enterprise world: Adobe, IBM, Apple, Amazon - the list goes on. However, these companies aren't exclusively rewriting the rules of product development, but rather using AI to improve the functionalities that customers are facing in incumbent products.

For some companies, developing user-facing AI solutions comes naturally as it solves well-known problems of their customers. For example, selecting people, clothes, hair, or other irregular objects in Photoshop used to be a painstaking and tedious activity involving multiple steps. Now, AI-based object detection allows users to create segmentation masks automatically. Graphic designers can use interfaces and selection tools leveraging ML instance segmentation technologies.

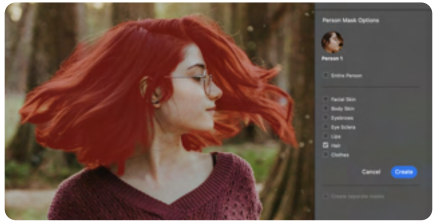

With many products, however, AI models will solve tasks that happen behind the scenes, rather than being directly accessible to end users.

Increasingly, the role of AI is becoming clear. Rather than reinventing the wheel, AI is instead becoming an undeniable cog in the development process of future products.

With this in mind, it's crucial to ask the question, "What do my customers actually want?"

During a webinar in May 2023, CTO and Executive Vice President of Microsoft, Kevin Scott, perhaps most succinctly addressed this conundrum when he said:

"One thing in particular that you have to bear in mind, is that the model is not your product. The model is just infrastructure that is enabling your product."

With this in mind, avoid one of the key failure points enterprise businesses make:

Build for the problem, not for the tech.

Questions to Ask When Exploring Enterprise Customer Needs:

How can I use AI to improve the user or customer experience of my consumer base?

How can AI redefine the way my product or service is used?

How can AI be a cost-saving measure - either within my business or for my customers?

How can AI improve the speed of my delivery?

How can AI enable my team to strategically pivot at the right time?

Why do AI Products Fail?

An estimated 70-80% of AI projects fail, with Gartner indicating that just 15% of use cases leveraging AI would be successful in 2022. The good news is, the failure points of AI projects have now become reasonably predictable. Below, we tackle the top reasons AI products fail.

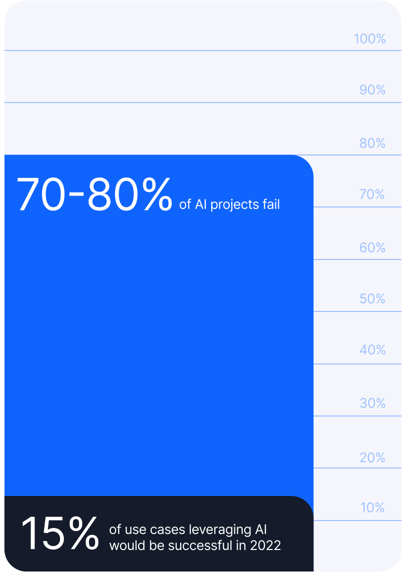

INSUFFICIENT TRAINING DATA

Overwhelmingly, AI projects fail as a result of lacking data. In a Gartner CIO survey, 53% of organizations rated their ability to mine and exploit data as “limited”. The knock-on effect? Poorly trained models, an agonizingly slow development process, and a final product that blends into the crowd.

With this in mind, AI success stories share one thing in common: a rigorous data collection process that is unique to their use case, continuously updated, and sufficient in scope to produce accurate models. Later on in this guide, we’ll discuss in more detail how to compile quality datasets.

"AI is the process of compressing training data. This data distribution represents the "knowledge" of an AI model. If a task or user interaction isn't represented in the distribution, the AI will fail. Labeling the right amount of data to contribute to AI knowledge is an exhaustive process, yet the primary correlator of model accuracy. Without the right training data, there is no engineering process that can save you".

Alberto Rizzoli
Co-Founder of V7

Insufficient Training Data: Takeaways

- Quality data is crucial to the success of your product.
- Unique data will be your competitive edge as an enterprise.
- Develop a data mining and exploitation strategy from the outset.
- Be strategic about how you extract value from your data.

SOLVING THE WRONG PROBLEM

Perhaps one of the most common reasons for AI product failure is poor product market fit. While AI promises a host of benefits, it can be easy to get lost in the weeds. Rather than seeing your model as the final product, it’s crucial that you take a holistic approach to AI-product development, ensuring it directly benefits company-wide goals.

It’s crucial to ask: What problem are you actually trying to solve? And, your follow-up question should be: Why is AI the best solution for that problem? Not only will this help you get stakeholder buy-in, but it will also help you to define what success looks like for your AI-powered product.

It can be exciting to ideate the myriads of problems AI can solve within a business, however, few problems solvable by AI today are true “painkillers”. It’s important to focus on whether a problem at hand brings measurable value to the business, or is a “vitamin” (a nice-to-have).

Countless AI products have failed on launch, as a result of decision-making that looked past a data-driven approach and instead favored the alluring promises of AI hype.

AI can, and will, fuel your ability to deliver seismic shifts to your industry. However, to do this, you’ll need to truly interrogate the what, why, and how of your project. Countless companies get lost in the technological bubble, and fail to meet the interests of the market in the process - ensure you’re not one of them.

“It’s crucial to think about the need and the use case.

Let’s take context. The context in which your model is being consumed is the most important factor in determining what kind of model you’re training in the first place.

If it’s being hosted on a compact edge device, you don’t have all the memory in the world to play with. If it's a consumer product, latency matters. If it's hosted on a supercomputer in a research center somewhere, you can take your time to focus on performance.

Until you’ve defined this and communicated it between model development and model consumption, you’ve set yourself up for failure. Start with the need and the use case.”

Kevin Chang
Training Data Operations Consultant

Solving the Wrong Problem: Takeaways

- Secure product-market-fit.
- Use platforms that allow you to rapidly create POCs/MVPs to test your hypothesis.
- Define the ROI for the business.
- Holistically treat the project like any other software project.

TIME AND COSTS

Countless projects fail thanks to hidden costs, missed deadlines, and an unwieldy product development process. In many cases, businesses underestimate the time and expense involved in creating an AI-powered solution and find themselves hemorrhaging cash due to an ineffective AI assembly line.

Worse yet, these businesses fail to earn market share, surpassed by competitors that beat them to the market with a solution that made better use of AI and was more cost-effective.

When building an AI-powered product pipeline, it is crucial that you address points at which time and resources can be lost. Fortunately, out-of-the-box solutions exist, with an express ability to monitor, control, and save time - saving costs in the process.

"AI re-invents itself every 18 months with major breakthroughs. As such, time is your biggest expense. Building on a "previous generation" stack leads to worse products, and a loss of talent that wishes to push the bleeding edge. The two biggest time sinks in machine learning are data labeling and internal tool development, both of which can keep progress hostage for months. Outsource these whenever you can."

Alberto Rizzoli
Co-Founder of V7

Time and Costs: Takeaways

- Bypass the time and financial investment of building an AI solution, and opt for an out-of-the-box platform that prioritizes easy integration.
- Automate components of the annotation and QA process with reliable models.
- Ensure expert labelers are only assigned the most crucial of tasks.

LACK OF TALENT

Another big cause for failure is an inability to secure or retain top AI talent. An enterprise AI project will invariably require exceptional AI, machine learning, and computer vision talent. Unsurprisingly, this particular candidate pool is hotly sought after, which means your ability to retain and reward your people will be paramount.

In the context of product development, few things are so demoralizing as a failing development pipeline. Innovation grinds to a halt, teams waste weeks on ineffective tasks, and fail to realize the impact of their hard work through market realization.

Talent Takeaways

- Equip your people with the best tools for the job, ensuring they can collaborate, focus on quality work, and reap the rewards of their efforts.
- Set up safeguards to mitigate against model failure rates, product recalls, and plummeting morale.

"Hiring quality ML talent is almost impossible, but when world-class developers have been identified, it is every leader's responsibility to equip them with best-in-class tooling that ultimately keeps them motivated and successful, and drives up staff retention rates."

Matt Brown
Head of Sales at V7

63% of businesses claim the largest skill shortages are in AI and machine learning.

The average income for an AI engineer in the US is between $78,000 to $150,000, but graduates from leading AI labs have starting salaries of $200,000 and above.

Fewer than 10,000 people have the skills necessary to become AI experts, which often come from building AI.

POOR ADOPTION AND LACK OF COMMITMENT

AI adoption can be an uphill battle, intensified by disparate approaches to development, sprawling tools, and silos of development. Worse yet, developing an AI product pipeline is expensive, so it's crucial that you squeeze every last ounce of value out of your tech stack.

AI products fail time and time again, thanks to poor buy-in, needlessly complex systems, and no repeatable process for success. With this in mind, it's crucial that you implement systems that are inviting to use, easy to investigate, and promote a unified approach to development.

Adoption Takeaways

- Take a look at your existing tech stack - how can it be integrated into your AI development process?
- Consider AI platforms that manage multiple stages of the process. Crucially, you need to ensure these platforms will easily integrate with a broader AI ecosystem.
- Consider the usability of your process - does it require a lot of training to get started? Does it provide clear oversight from a strategic and performance level? Does it make the day-to-day tasks of your AI team worthwhile?

What is Your Competitive Advantage?

The market is rapidly being flooded with AI-powered product solutions, each seeking to stake its claim as the product of choice. When embarking on your product development plan, you're likely to have one niggling question: what could my competitive advantage be?

UNIQUE DATA

For enterprise AI, the answer to this question will always be data. The success and failure rates of AI-powered products overwhelmingly rely on the quality, quantity, and availability of data. Businesses that rise to the top ultimately possess unique data that clearly differentiates them within their market. Take the time to craft a unique data pipeline that sets your product above and beyond the competition.

EFFECTIVE INFRASTRUCTURE

The infrastructure that houses your product pipeline is crucial. Fail to get it right, and you run the risk of hemorrhaging cash, demotivating your people, and creating lackluster products. You'll need to carefully consider the needs of your project, the infrastructure required to scale it, and the runway time you actually have to kick things off. Further in this guide, we tackle the pros and cons of leveraging a training data platform and outline how to choose the best one for you.

EARLY ADOPTER

Is now the time to strike with AI? The short answer is yes. The route to building AI-powered products has never been more clear, exemplified by a stream of AI product releases in 2023 alone. For those with the will, expertise, and funding to do so, the impact of first-movers' advantage will be undeniable. Early adopters will have a competitive advantage over their rivals, with an accelerated development pipeline that builds better products, faster.

Who's Advancing Fastest With AI

Q. To what extent is your company looking to integrate AI technologies into its operations?

Source: PwC 2022 AI Business Survey, March 2022: Leader base of 364; Other base of 631

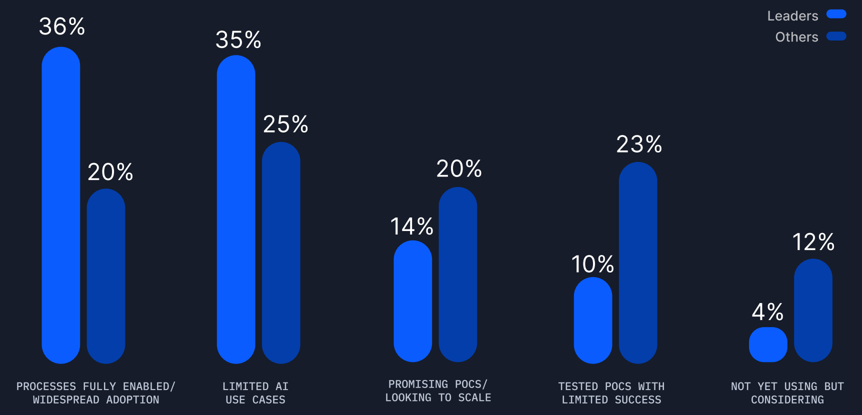

https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-business-survey.html

"Enterprises sit atop a goldmine of untapped data. As computer vision evolves, following the advancements in large language models - epitomized by innovations like Facebook's SAM - its power only increases. These foundation models are developed on generalized images and videos sourced from the web, yet enterprises possess a distinct advantage. Within your specialized domain, you hold the keys to build a top-tier AI product, thanks to your unique data that can't be found with a Google search".

Bill Leaver
Product Manager at V7

BUILDING AI

Your Enterprise Strategy

At the root of AI success is a well-considered strategy. According to Deloitte, businesses that secure high outcomes and high deployment rates are three times more likely to have a company-wide AI strategy in place.

In this section, we tackle key components of an enterprise strategy that will allow you to make the most of AI.

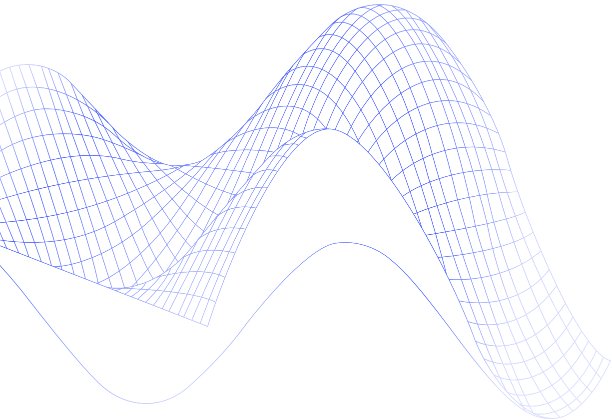

Starting with your North Star

AI poses an incredible opportunity, however, it's crucial to separate hype from tangible impact. When embarking on your project, you'll need to assess:

- How does this project align with your company's overarching goals?
- How and why is AI the best solution for your challenge?
- Where will AI bring the most value to your project, product, or process?
- How will you define the ROI of your AI development?

Without doing this, you'll struggle to create a product that aligns with company strategy and ultimately - fails to earn future investment.

"Always start with the need, before the AI. AI is a cool piece of technology and a lot can be learned by simply trying it out, but it's always better to apply AI to a problem you have right now. How many resources you put towards an AI project ultimately depends on the return on investment from those resources. Sometimes, that return on investment will just be knowledge and research and understanding - but it should always be for an explicit return."

Kevin Chang
Training Data Operations Consultant

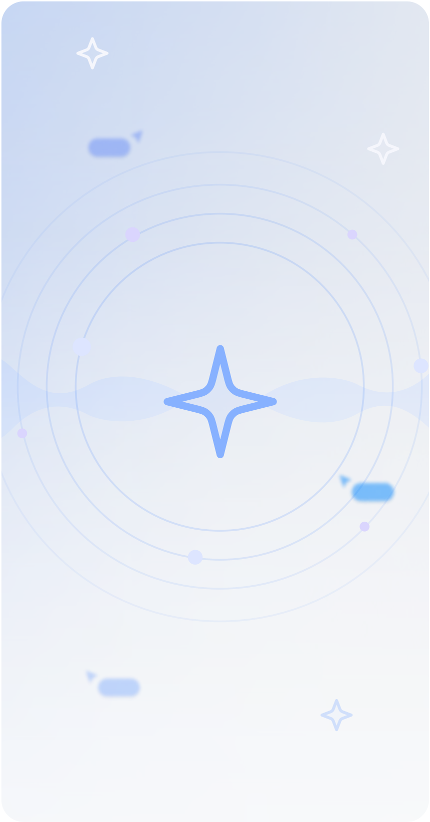

Setting Core Foundations: Analytics, Talent, Data

There's an infamous line within the AI industry, "Garbage in, garbage out." In simple terms, this refers to the importance of quality inputs, to secure quality outputs. In this case of AI, a lacking data strategy will turn up a faulty - at best - product. It's crucial you address:

- What expertise do you have internally? Do you have the data engineering talent you need to create a reliable data strategy?
- How do you measure "good"? Do you have data analytics expertise internally or do you have infrastructure that provides metrics?
- Do you have a data pipeline with sufficient quality, quantity, and reliability? Do you have a process in place to ensure your data pipeline continues to deliver?
- Are your people equipped to do a good job? Have you ensured they have the resources to be as effective and efficient as possible?
- Do you have a process in place that allows you to be optimal in how you assign data labeling tasks? Are you able to divide between expert labelers and lower-cost support?

"Once you have the need addressed, you can look at the resources you have to realize that need. In AI, those resources are always the data, the people, the processes (such as data collection), and the model that you actually want to build.

It's one thing to say you want an AI-powered product, but the model is just going to be one part of that overall product. You need to consider how your model and product will interact, how you will manage and structure data, and how you will deliver a data pipeline that considers quality, data culture, and data availability."

Kevin Chang
Training Data Operations Consultant

Interrogating Your Infrastructure

Unsurprisingly, your tech stack will be a crucial component of your AI strategy. This will span data collection, data management, data processing, and data outputs. While you can opt to build your own solution, there are many out-of-the-box solutions that manage a large sway of the development lifecycle, while integrating with third-party apps, platforms, and cloud-service vendors.

Many of these platforms successfully reduce the total cost of ownership for AI projects, while dramatically boosting the productivity of ML teams.

When it comes to your tech stack, you'll need to:

- Address how to reduce the total cost of ownership for AI projects.
- Prioritise infrastructure that promotes easy experimentation, iteration, and proofs-of-concept.
- Opt for solutions that seamlessly integrate within a wider AI ecosystem.
- Consider how to optimize data labeling, model training, and deployment.
- Reduce the time spent manually managing infrastructure.
- Leverage a solution equipped to evolve with the rate of industry change.

"This is where workflows, products, and tooling comes in. You need to make sure that your people are set up for success, and not wasting their time on manual or dull work - because it will massively lower your return on investment.

Have tooling in place, but also have processes designed around how people interact with and use that tooling."

Kevin Chang
Training Data Operations Consultant

Tackling Company-Wide AI Adoption

Disparate approaches to AI development will cause a series of issues. Development costs increase, profitability plummets, teams become ineffective, and the process overall is marred with latency and waste. This can happen for a number of reasons, which ultimately come down to your AI adoption strategy.

- Have you considered how your infrastructure will work in practice?
- Have you hired the right people to push the project forward?
- How can you ensure the company embraces a unified approach?
- Have you created project "champions"? Or in other words, empowered individuals to promote best practices across the organization?
- Have you ensured the business is optimal with its tools in use?
- Have you leveraged tools that integrate with your existing systems?

For enterprises with international offices, often a breakdown in communication results in a business paying for multiple tools that achieve the same functionality. Similarly, without a unified approach to development, businesses will often face slow development times and costly overheads. Carefully consider how you're going to set up AI development, and put the work in to ensure the business has a unified strategy.

"AI can be seen as a disruptor of existing processes: it's always better for adoption to develop it as a Copilot or Assistant to accelerate a process or task that is still human-driven. Some AI projects also don't expand as they try to do too much early on, reaching unfavorable accuracies. It's better to instead focus on well scoped relevant problems."

Alberto Rizzoli
Co-Founder of V7

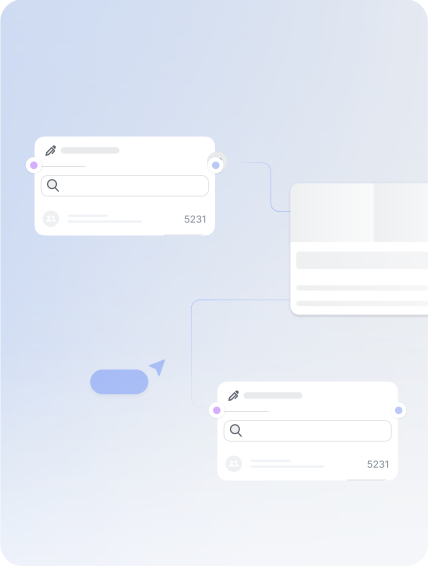

Developing with Risk and Responsibility in Mind

AI development can be vulnerable to risk, whether that's models that produce bias, or products that expose your business to reputational or legal damage. With this in mind, it's crucial you strategize how to de-risk your project from the beginning. Here you'll need to pay careful consideration to things like:

- What systems do you have in place to remove bias in the AI?
- Does your product development process make it easy to abide by regulations in your country? Or is it a black hole of development?
- For regulated industries, do you have infrastructure that aligns with your corporate governance goals?
- Is your business creating AI in a way that promotes data explainability and transparency? Do you have processes in place to make this simple?
- Does your process make it more challenging to secure compliance? Ie, for the FDA, HIPAA, ISO27001, SOC 2, GDPR, etc.
- Do you have a defined AI strategy to govern how you work, and do your people follow it?

"Mitigation of AI risk should be embedded into the way machine learning is developed, starting from the labeling instructions to the selection of training parameters.

AI leaders must identify the technical, financial, operational, reputational, and compliance risks to the business, and then formulate the strategy, process, and tooling needed to benchmark products ahead of releases, taking into account that not all data can be modeled, so AI will eventually make mistakes."

Kevin Chang
Training Data Operations Consultant

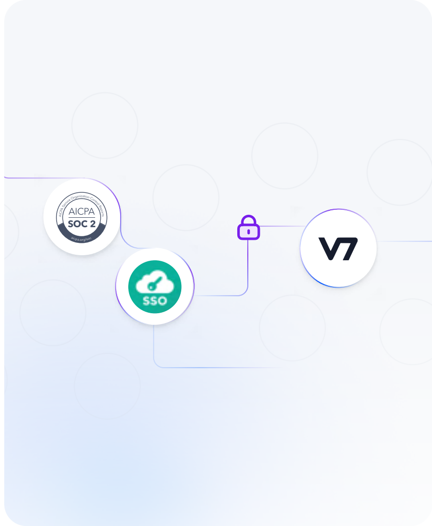

Your Toolkit Checklist

Should you Build or Buy?

When it comes to building your AI toolkit, it can be overwhelming to know where to start. New solutions crop up daily, and the AI, ML, and CV infrastructure landscape continues to evolve. This is in stark contrast to the early days of enterprise AI, where companies were often forced to build their own AI/ML infrastructure from scratch.

Overall, this approach was prone to failure, expense, and complexity - with an investment of time and resources that often limited the overall profitability of AI. Engineers would be dedicated to the upkeep of the infrastructure, the business shouldered the overhead costs of internal development, and the process itself was at the whim of market shifts that would render the components of the process obsolete.

The AI Infrastructure Alliance (AIIA), “dedicated to bringing together the essential building blocks for the Artificial Intelligence applications of today and tomorrow” recommends against this approach, citing the complexity and costs of building your own solution.

Below you can see a comparison of in-house development expectations, versus the benefits a company can expect from a training data operations platform.

	In-House Development	Platform
Initial Cost	High	Lower
Ongoing Maintenance	High	Included in SaaS Subscription
Scalability	Varies	Highly Scalable
Speed to Release	Slower	Faster
R&D/New Features	Limited	Greater Potential

Below, we’ve broken down some of the historical reasons companies opted to build, as opposed to buying an out-of-the-box solution.

Why companies decide to build

- Scarcity of external options
- Preference for keeping critical infrastructure as part of their own IP (or necessity to meet legal requirements)
- Unique use case and more control

Why companies decide to buy

- Engineering time can be better spent on their core product or other critical infrastructure tasks
- Pace of innovation in AI is higher than their team can keep up with
- ML platforms have become more flexible

“On V7’s Customer Success Team, we have the benefit of always being on the other end of the “buy” side of the build-or-buy decision. Computer Vision Teams often tell us that they made the decision to go with V7 over a custom solution for two main reasons: feature richness, and the cost of maintaining a custom solution.

I think of Meta’s SAM release as a great case study for the “buy” decision. As soon as SAM was announced we rallied V7’s Product and Deep Learning Teams to build a user-friendly integration. Our CS Team worked closely with V7’s customers to determine how SAM could be incorporated into their training data pipelines.

For our customers, this meant that engineering cycle time that would have been devoted to experimenting with SAM was freed up for other high-impact projects. We were able to focus on staying ahead of the curve so our customers could do the same.”

James Hudson
Principal Customer Success Manager at V7

The pros and cons of both approaches boil down to whether you need a solution tailored to your specific use case or prefer instant access to state-of-the-art solutions. A platform will give you the ability to focus engineering resources on other tasks.

	Build	Buy
PROS	Tailored to specific use case Sense of ownership and control over the software In some cases, working on an existing solution instead of making a painful switch	Instant access to state-of-the-art solutions Ability to focus engineering resources on core differentiators Professional support and a larger engineering team solely dedicated to improving this form of tooling
CONS	Potential scaling challenges as the team grows Requires dedicated, scarce and highly competitive talent Time-consuming and costly	It sometimes can be difficult to integrate software into existing systems Potential need for training to effectively use the software Lower flexibility in use-case customization

While building your own solution gives you more control, it can be a massive project that will drain resources. And it is a leap of faith—no one will promise you that your custom solution built in-house won’t become obsolete by the time it is ready.

There's another crucial consideration here: cost.

This includes the monetary investment as well as the opportunity cost of potential errors, falling behind with inferior capability, and distractions from core competencies. So, what are the actual costs we need to consider?

- $250,000 for a minimum viable product
- $100k to $1 million per year in cloud storage, image processing, and redundancy.
- Front-end, Back-end, and Deep Learning engineers requiring specialized skills for large-scale label rendering, inference infrastructure, and data versioning.

Despite this, there may be scenarios where building your own solution is the right move for your business. If you're trying to assess the right step forward for your business, the below checklist can help with defining this strategy.

	Buy	Build
1	Assess your team's technical expertise and available resources.	Assess your team's technical expertise and available resources.
2	Determine if an off-the-shelf solution meets your team's needs.	Identify the specific features and functionalities your team requires.
3	Research available vendors and evaluate their offerings.	Research the available open-source tools and libraries.
4	Compare the costs of purchasing and maintaining the software with the cost of building it in-house.	Evaluate the costs of building the software, including development time, ongoing maintenance, and potential opportunity costs.
5	Consider the time-to-market for a purchased solution versus a built solution.	Consider the time-to-market for building the solution in-house.
6	Evaluate the level of customization and flexibility required for your team's needs.	Determine if an off-the-shelf solution can be customized to meet your team's needs.
7	Consider the level of support and training provided by the vendor.	Determine if your team has the necessary expertise to support and maintain the software in-house.
8	Make a decision based on a cost-benefit analysis of the options available.	Make a decision based on a cost-benefit analysis of the options available.

The Best Training Data Operations Platforms

As opposed to building your own platform, there are a number of training data platforms that offer a seamless route to AI creation.

Goals of Training Data Ops Software

- Creating More Accurate AI Models
- Faster Iteration And Development Of AI Products
- Reducing Total Cost Of Ownership For AI Projects
- Boosting The Productivity Of ML Teams

Let's tackle their strengths, weaknesses, and the key points you should interrogate →

Popular Training Data Operations Platforms

To kick things off, let's take a look at some of the most popular platforms on the market today.

V7

- Automated annotation features without prior training needed.
- Composable workflows allowing multiple models and human in the loop stages.
- Dataset management that stays robust at large scale.
- AI and Human workflows that allow you to bring your own model, or integrate external ones.
- Integrated data labeling services.
- Real-time collaboration and fluid UX.
- Frame-perfect Video annotation tool.
- Backs enterprise companies like Huawei, Continental, MIT and Wanzl.

Labelbox

- AI-assisted labeling (BYO models).
- Integrated data labeling services.
- QA/QC tooling and label review workflows.
- Strong labeler performance analytics.
- Customizable interface to simplify tasks.
- Used by the likes of BueRiver Technology, Dialpad, and P&G.

scale

- ML-powered pre-labeling.
- Nucleus dataset management.
- Automated QA system with gold sets.
- Document processing features.
- Model-in-the-loop data curation.
- Used by the likes of Orchard Robotics, Harvard Medical School, and OpenSea.

Dataloop

- Model-assisted labeling.
- Multiple data type support.
- Advanced team workflows with streamlined data indexing and querying system.
- Video support.
- Used by the likes of Standard, Foresight, and Elbit Systems.

SUPERVISELY

- AI-assisted labeling.
- Multi-format data annotation and management.
- Option to develop and import plugins for custom data formats.
- 3D Point Cloud.
- Options for project management on different levels for teams, workspaces, and datasets.
- Used by the likes of Engie, Bitwise, and Firemark.

SuperAnnotate

- AI-assisted labeling (BYO Models).
- Superpixels for semantic segmentation.
- Advanced quality control systems.
- Supports various formats through image conversion.
- Used by the likes of IBM, Samsung, and Aurora.

WHAT SHOULD YOU CONSIDER WHEN CHOOSING A TRAINING DATA PLATFORM?

One of the core benefits of using a platform is its ability to manage a host of stages within your AI development pipeline. For example, a good platform will allow you to manage datasets, rapidly annotate, automate processes, conduct QA, train models, and iterate models, while providing easy-to-use dashboards that make management, efficiency, and profitability easy to control. However, if your platform fails to integrate with an ever-evolving AI ecosystem, it runs the risk of isolating your development from available tools - stunting its growth in the process.

Similarly, while platforms can provide much-needed oversight for stakeholders at all levels - if its user interface is clunky, it quickly becomes a development black hole that makes management an agonizing process.

Customer of V7

"In AI product development, Altitude Infra's most significant challenge has been to generate quality datasets quickly and efficiently. We needed a solution to accelerate our annotation process, one that better integrates in our model training workflows, boasts an intuitive UI, and enhances project management for annotation tasks. By offering easy-to-use interface and valuable resources for our developers, V7 made our model training process more effective."

Lucien Muller
Project Manager, Altitude Infra

To decide on the right platform for your business, you'll need to holistically consider your AI development process, with a focus on capabilities around integration, management, ease of use, and production power.

To get started, it helps to ask:

- Can the platform scale with your production needs?
- Is the platform easy to understand, across all levels of seniority? (Annotators, Data Scientists, Managers, C-Suite)
- Does your platform seamlessly integrate with third-party apps and cloud service vendors? Does it align with existing tools in your stack?
- Is this platform customizable to your needs? Or is it rigid by nature?
- Can you build secure and compliant AI within this platform?

To answer these questions in depth, we've taken a look at some of the most popular training data platforms of 2023, and broken them down by essential features, supported formats, and security.

	V7	Labelbox	Scale AI	SuperAnnotate	DataLoop	Supervise.ly	Roboflow	Kili Technology	Clarifai	Appen	Encord
Essential Features											
Model-assisted labeling	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Model training	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Model interface	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
CLI SDK	✓	⊖	✓	✓	✓	✓	✓	✓	⊖	⊖	✓
Python SDK	✓	✓	✓	✓	✓	✓	✓	✓	✓	⊖	✓
Mobile SDK	⊖	⊖	—	⊖	⊖	⊖	✓	⊖	✓	⊖	⊖
Annotator statistics	✓	✓	✓	✓	✓	✓	✓	✓	⊖	✓	✓
Free-hand brush and pen tools	✓	✓	✓	✓	✓	✓	⊖	✓	⊖	✓	✓
QA stages	✓	✓	✓	✓	✓	✓	✓	✓	⊖	✓	✓
Interface customization	✓	✓	✓	✓	✓	✓	⊖	✓	⊖	⊖	⊖
Vector segmentations	✓	✓	✓	✓	✓	✓	✓	⊖	⊖	✓	✓
Dataset Management											
Bulk data import/export	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Detailed training dataset filters	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Version control	✓	✓	✓	✓	✓	⊖	✓	✓	✓	—	✓
Security											
HIPAA Compliance	✓	✓	✓	✓	⊖	✓	✓	✓	✓	✓	✓
SOC2	✓	✓	✓	✓	✓	⊖	✓	✓	✓	✓	✓

Keep reading →

	V7	Labelbox	Scale AI	SuperAnnotate	DataLoop	Supervise.ly	Roboflow	Kili Technology	Clarifai	Appen	Encord
					Supported Formats						
Video support (Video Native support, not as images)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Video label interpolation	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Ultra high-resolution support	✓	✓	✓	✓	⊖	⊖	⊖	✓	✓	⊖	✓
PDF Document processing	✓	✓	✓	✓	✓	⊖	✓	✓	✓	⊖	⊖
DICOM	✓	✓	⊖	✓	⊖	✓	⊖	✓	✓	⊖	✓
Ultrasound support	✓	✓	⊖	⊖	⊖	✓	⊖	⊖	✓	⊖	✓
MRI/CT volumetric support	✓	✓	⊖	⊖	⊖	✓	⊖	✓	⊖	⊖	✓
3D LIDAR support	⊖	⊖	✓	⊖	✓	✓	⊖	⊖	⊖	✓	⊖
5000+ annotations per image support	✓	—	—	—	—	—	—	✓	⊖	—	✓
Text support	✓	✓	✓	✓	✓	⊖	✓	✓	✓	✓	⊖
					Other						
ADAS Mapping support	⊖	⊖	✓	⊖	✓	⊖	⊖	⊖	⊖	⊖	⊖
Data augmentation services	⊖	✓	✓	✓	✓	✓	✓	✓	✓	⊖	✓
REST API capabilities	✓	✓	✓	⊖	✓	✓	✓	✓	✓	✓	✓
Labeling Services	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Consensus (QA tool)	✓	✓	✓	✓	✓	✓	⊖	✓	✓	✓	✓
SLA Agreement	✓	✓	✓	✓	✓	⊖	✓	⊖	⊖	✓	✓

Looking more specifically at a simplified AI development process, below you can see the core areas that V7's platform covers:

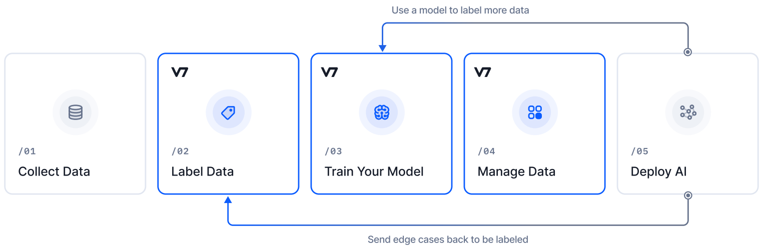

While there's a bit of groundwork in securing the right platform for your business, it's time worth spending. For those that invest in the best architecture for their use case, benefits are most commonly realized in two years or less, according to the AIIA "AI Infrastructure Ecosystem" survey.

How long did it take for the benefits you get from AI/ML to outweigh the costs of infrastructure, implementation and personnel?

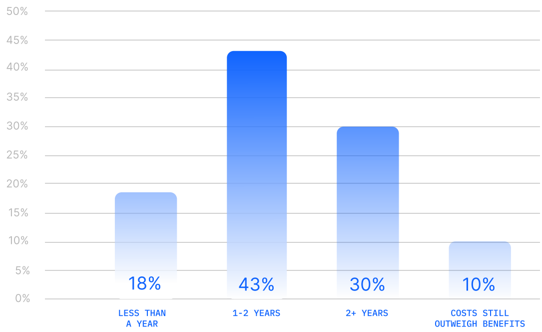

The Broader Ecosystem of Your AI

While some platforms will claim to be “end-to-end”, the reality is, they just cover a few components of the overall process. You need to be careful not to invest in a platform that's inflexible to customization, which will only serve to frustrate your AI development.

Pay careful attention to which platforms promote seamless integration, and which ones require extra work to connect to cloud storage, ML/OPs platforms, and deep learning frameworks of choice.

Prioritizing a modular approach, V7’s platform is customizable to evolving enterprise needs, including a Python SDK and API to allow bespoke integrations. Below you can see a snapshot of just some of the in-built integrations the platform offers.

This flexible infrastructure is backed by a host of resources, guides, and GitHub repositories, designed to make integration with V7 as seamless as possible.

BUILDING AI

Tackling Your Data

The precision of your model relies on one crucial thing: the scope, standard, and quality of your data. Put simply, it can be the difference between AI success and an AI flop. Below we tackle how to build quality datasets, achieve training data accuracy, and implement best practices for dataset management.

“Those companies that view data as a strategic asset are the ones that will survive and thrive.”

Bernard Marr
International Best-Selling Author
“Data Strategy: How to Profit from a World of Big Data, Analytics and the Internet of Things”

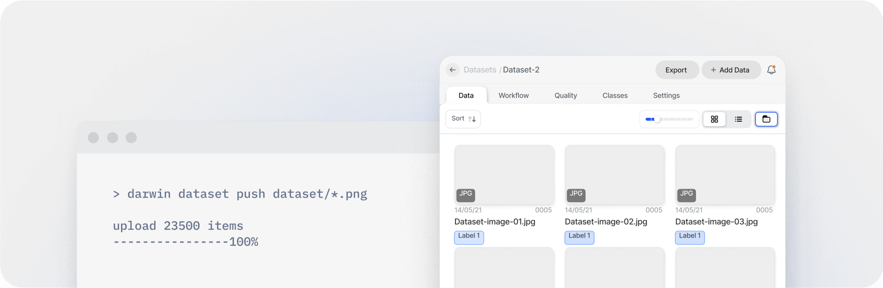

How to Develop Quality Datasets

As the foundation of your AI's knowledge, your data collection strategy needs to be carefully considered. Below we tackle four popular ways of sourcing datasets, and outline the pitfalls and pros that come with each method.

Customer of V7

"To create a successful AI product, obtaining the highest quality data possible to train your models is key. The utilization of tools that facilitate the generation of this data is essential."

Lucien Muller
Apprentice Project Manager, Altitude Infra

OPEN DATASETS AND SEARCH ENGINES

Depending on your use case, you can leverage a host of open datasets, which are free and relatively easy to use. In fact, V7 provides over 500 open-source datasets for machine learning projects, spanning from breast duct cell segmentation to aircraft detection. However, be mindful that open datasets often do not represent the unique object appearances, angles, and movements you may want to detect in your unique business problem and are therefore not enough.

Dataset search engines and forums are also viable options, providing scores of datasets to choose from.

Popular Dataset Search Engines

Google Dataset Search	Kaggle	Data.world	UCI repository
25 million datasets available	68,000 datasets available	128,278 datasets available	622 datasets available

Pros of Open Datasets and Search Engines

- Often open-source and free to use.
- Wide variety of datasets on offer.
- Numerous reputable sites (such as Google's Dataset Search) to choose from.
- Promotes open data culture and fosters innovation.
- Saves time otherwise spent on collecting data.
- Saves costs otherwise spent on data collection.

Pitfalls of Open Datasets and Search Engines

- Datasets are not unique to your use case, objects, or edge cases undercutting your competitive advantage.
- Can be difficult and time-consuming to find the right datasets.
- Datasets aren't your own, which can make relying on them long-term a bit of a trap. For example, you could be relying on an outdated dataset, when your own, new, fresh set of unique data could be better suited for your use case.
- May pollute your accuracy tests with unlikely imagery.

SCRAPE WEB DATA

Web scraping utilizes bots to collate content and data from numerous websites. This method is mostly used when a use case demands multiple sources for diversified inputs. Web scraping focuses on public online resources, such as government websites or certain social media platforms. It's worth noting, web scraping has made the headlines in recent years - for good, and bad, reasons. One of the more notable headlines, was a $277M GDPR fine for Meta back in 2022, for failing to appropriately protect users from illegal web scraping. With this in mind, web scraping can be invaluable - but ensure your strategy is a compliant one.

Pros of Scraping Web Data

- Real-time data is available, which can be particularly useful for use cases that require up-to-date data, trends, or information.
- The process can be automated, saving time and resources that would have been spent on collection.
- Ability to create very large datasets for general purpose foundation models.

Pitfalls of Scraping Web Data

- Web scraping has a learning curve, requiring internal expertise to get right.
- Handled incorrectly, web scraping can open your business to legal challenges. It's crucial to ensure your business doesn't end up scraping and using data that could get it into hot water.
- Certain sites will outright block scrapers, such as LinkedIn or Facebook.
- There is a potential for bias. For instance, if your training data is scraped from TikTok US and you are building an AI-powered TikTok app for the Asian market, it may not work as expected. Factors such as language, user demographics, and geography can all play a role.
- While scraping can provide a large amount of data easily, it is harder to verify the accuracy of the data.
- Web scraping can seldom be applied to professional use cases, hence why foundation model developers hire professional data collection firms.

OWN DATA

Depending on your use case, the above options might not be the right fit for your business. If that's the case, you'll need to leverage your own data as an enterprise. This can be particularly advantageous, as it's ultimately unique to you.

An example of this in action would be the creation of a chatbot, designed to respond to customer problems. Rather than using natural language processing datasets, data could be extracted from company documentation and anonymized records, etc.

Pros of Own Data

- Unique data provides a competitive advantage.
- Collecting this data is relatively easy and inexpensive.
- Data is often of a high standard, reliable, and fit for purpose.
- This data most accurately represents the problem or task you want to represent, leading to higher accuracies.
- Greater control around the validity of data and the legal ownership of data.

Pitfalls of Own Data

- You need to be certain you have legal ownership of the data that you're utilising. For example, under which circumstances can you use customer data? What needs to happen to the data to make it usable?
- Without a legal process in place, it will not just be challenging to collect data - it will be risky.

USE DATA AUGMENTATION

Sometimes, it's particularly difficult to either collect data or achieve the quantity needed. When this happens, you can utilize data augmentation, which essentially allows you to apply different transformations of the original data to create new data. For example, for image data, training data size can be increased by simple operations like rotation, changes in color, brightness, etc.

Pros of Data Augmentation

- Increases the amount of training data available, and can be particularly useful when collecting data is difficult.
- Saves time and costs.

Pitfalls of Data Augmentation

- Will carry the bias of the existing datasets.
- Needs systems in place to assess the quality of the augmented data.
- Now incorporated into most machine learning framework at train-time, therefore generating augmented images is becoming obsolete.

SYNTHETIC DATA

Data can also be fabricated by humans and software to represent reality. This may include extreme edge cases such as fire truck accidents for autonomous driving or rare defects in manufacturing. Services companies emerged three years ago to provide synthetic data services which may offer a stopgap solution to data availability. Today they are fewer in numbers and popularity as many use cases have proven to be challenging to comprehensively model in synthetic environments, however it remains popular in research.

Pros of Synthetic Data

- Allows the creation of edge cases where data cannot be collected.
- Comes pre-annotated as the image and labels are computer generated.
- Great for research tasks where image recognition is not key, such as robotic grasping.

Pitfalls of Synthetic Data

- Only represents the appearances modeled by the artist, omitting out-of-distribution cases.
- Still time consuming to generate despite being created digitally due to the time needed to model object varieties.
- Cannot be applied to areas that cannot be drawn or 3D modeled, such as human tissue, cellular biology, and other amorphous objects.

Data Labeling

Now that you've tackled data collection, you'll need to get annotating. We address key considerations when labeling your data, from varying strategies to cutting-edge annotation features on the market.

Let's start with a quick glossary breakdown.

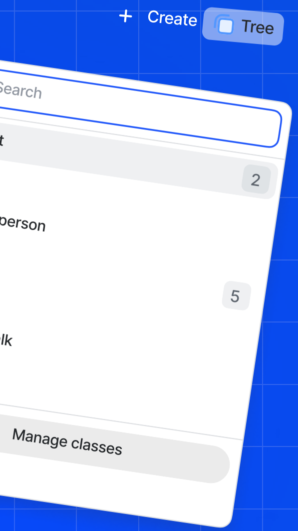

UNLABELED DATA VERSUS LABELED DATA

Your training dataset is completely dependent on the type of machine learning task you want to focus on. Machine/Deep Learning algorithms can be broadly classified on the type of data they require in three classes.

SUPERVISED LEARNING

Supervised learning, the most common type, is a type of machine learning algorithm that requires data and corresponding annotated labels to train. Popular tasks like image classification and image segmentation come under this paradigm.

Typical training consists of feeding annotated data to the machine to help the model learn and testing the learned model on unannotated data.

UNSUPERVISED LEARNING

In unsupervised learning, unannotated input data is provided, and the model trains without any knowledge of the labels that the input data might have.

Common unsupervised algorithms of training include autoencoders that have the outputs the same as the input.

SEMI-SUPERVISED LEARNING

In semi-supervised learning, a combination of both annotated and unannotated data is used for training the model.

While this reduces the cost of data annotation by using both kinds of data, there are generally a lot of severe assumptions of the training data made while training. Use cases of semi-supervised learning include protein sequence classification and Internet content analysis.

SELF-SUPERVISED LEARNING

In semi-supervised learning, a combination of both annotated and unannotated data is used for training the model.

While this reduces the cost of data annotation by using both kinds of data, there are generally a lot of severe assumptions of the training data made while training. Use cases of semi-supervised learning include protein sequence classification and Internet content analysis.

WHAT IS 'HUMAN-IN-THE-LOOP' (HITL)?

The term Human-In-The-Loop most commonly refers to constant supervision and validation of the AI model's results by a human.

There are two main ways in which humans become part of the machine learning loop:

- Labeling training data: Human annotators are required to label the training data that is being fed to (supervised/semi-supervised) machine learning models.
- Validating the model: Rare tasks and low-confidence predictions are validated by a human and the results of the validation are fed back to the model. This validated data is highly effective as it represents areas of improvements of AI.

DATA LABELING APPROACHES

There are a number of approaches to data labeling, depending on your specific use case, the time available for your project, and the size of your team. Below we tackle the most common approaches for annotating data and outline their pros and cons.

IN-HOUSE DATA LABELING

In-house data labeling achieves the highest quality labeling possible and is generally done by data scientists and data engineers within your business. High-quality labeling is crucial for industries like insurance or healthcare and often requires consultations with experts in corresponding fields for effective labeling of data.

As is expected for in-house labeling, with the increase in quality of the annotations, the time taken to annotate increases drastically, resulting in a slow data labeling process.

To combat this, training data engines will equip in-house teams with machine-based annotation capabilities - a process that drastically increases the speed of data annotation. This team is the closest to your machine learning talent and will be the quickest to react to any changes. They can be excellent reviewers of Professional Labeler data or be utilized for crucial projects.

PROFESSIONAL LABELERS (PARTNERS)

Some training data operations platforms, like V7, will offer access to a network of highly trained professional annotators, with expertise in a range of niche areas. This allows companies to scale as they grow, maintain data accuracy, and leverage platform power users.

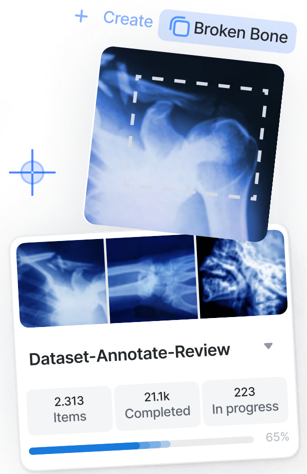

CROWDSOURCING/OUTSOURCING

Crowdsourcing secures annotated data with the help of a large number of freelancers registered at a crowdsourcing platform. These datasets will consist mostly of trivial data like images of animals, plants, and the natural environment and do not require additional expertise.

However, it's important to be careful with these platforms. As a requester, you can specify the nationality of your workers, leading many to request labelers from English-speaking countries. However, as labelers are paid per task completed, you may come across fakes. Some crowd users are not honest with their reported demographics and proficiency with English, as it affects their pay. Other misbehaviours include completing occasional tasks randomly or utilizing bots to do so, producing garbage training data. With this in mind, be careful to do due diligence when it comes to crowdsourcing your labeling.

Outsourcing is a middle ground between crowdsourcing and in-house data labeling where the task of data annotation is outsourced to an organization or an individual. One of the advantages of outsourcing to individuals is that they can be assessed on a particular topic before the work has been handed over.

CUTTING-EDGE FEATURE RELEASES IN 2023

Up next, we tackle some of the most cutting-edge features being released in 2023, and outline the impact they can have on enterprise AI.

SAM-enhanced Auto Annotate

2023 brought the release of Meta's "Segment Anything Model (SAM)", a powerful AI that can rapidly segment any object in an image or video with stunning precision. In response, V7 leveraged SAM's impressive segmentation abilities by fusing it with the V7 platform's "Auto-Annotate" feature.

Providing an engine and intuitive mechanics for the feature, V7's SAM-powered auto annotate can now annotate at incredible speeds, without compromising on quality. The quality of automated labels on V7 often matches human quality and significantly reduces fatigue which leads to much more dangerous recall errors (missed or misclassified objects).

Input image | Pre-processed with SAM

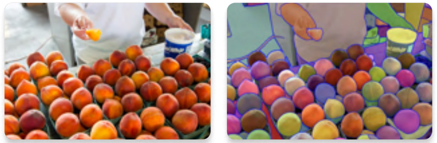

AutoLabel

Building on SAM's impressive segmentation capabilities, V7's AuoLabel allows you to detect similar objects and label multiple instances with a single click.

The benefits of this are innumerable:

- Instead of manually annotating every instance of an object, you can annotate the most representative one and the rest will be automatically segmented.
- In certain scenarios, AutoLabel's polygon segmentation accuracy is more aligned with the actual object shape than the standard pre-processing performed by SAM.
- A quick comparison of confidence scores for different objects gives you more information about your training data.

Accelerated Label Rendering (ALR)

Continuing the thread of impressive speed improvements, V7 introduced ALR (Accelerated Label Rendering) this year, delivering a 17x speed improvement in annotation action.

Working on densely annotated files, for example - those with more than 1000 annotations in a single image - can be a challenge for even the most powerful hardware. With so many annotations to render, a workview canvas can become slow and unresponsive, in turn making it challenging to add, edit, or review annotations.

With the introduction of ALR, V7 supports multiple canvas layers, in addition to supporting viewport rendering. The result? Faster rendering, smooth interactions, and an overall better experience.

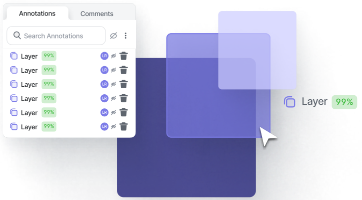

Consensus Stages

Reliable accuracy is paramount when it comes to your data and the performance of your model. However, from time to time, there are disagreements between experts. Particularly in fields such as healthcare or life sciences, nuanced data can cause problems.

In response, V7 released its Consensus Stage update, to allow multiple users or models to label the same task - in a blind test. From here, differences are calculated between the labels, and disagreements are surfaced - delivering a reliable QA process for securing the most accurate annotations. You can also use Consensus Stages to compare the overlap between AI and human or two versions of AI models.

Sampling Stages

High-quality training data is crucial for high-performing models. However, reviewing each and every data point in a dataset can be impractical and time-consuming. To combat this, V7 has introduced sampling stages, which equips users with efficient quality assurance, while minimizing review time.

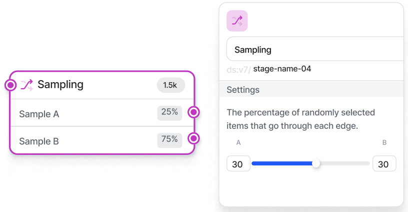

The core functionality of the Sampling Stage allows you to send a randomly selected sample of items to the next stage in a workflow. For example, you can route a small subset of material (for example, 10 to 15%) for QA, without painstakingly checking the entirety of the dataset.

The benefits of this are innumerable, from efficient QA for vast amounts of data, to improved AI model auditing, flexible QA for specific data classes, and continual measurement of your annotator's consistency.

Webhook Stages

As we've covered, integration is key. That's why V7 introduced its Webhook Stage, which is one of the many building blocks you can use to build a bespoke workflow in the V7 platform. The Webhook stage allows you to send data across different web apps via webhooks. For example, with third-party tools like Zapier, you can connect your V7 workflows to a host of other apps, such as Gmail, Trello, or Slack.

This delivers a number of benefits, from interconnectivity and automation, to monitoring and reports.

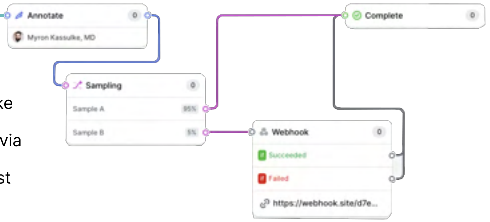

BYOM UI

In addition to allowing you to train models within the platform, V7 also allows you to bring your own model into the platform to advance your product development pipeline. This means you can automate complex tasks with your own IP models, access open-source models in an instant, and tackle unique challenges with the incorporation of custom models into V7's Auto-Annotate.

Better yet, this can all be done via the V7 GUI, meaning the process is as simple as possible.

MANAGING YOUR DATASETS

The integrity of your data matters. From where it's sourced, to how it's maintained, it's important that you consider how your datasets will be managed. In this section, we briefly cover dataset management best practices, and outline some key considerations.

Data Management Policy

Before kicking off, you'll want to put a data management policy in place. Here you'll want to address the type of data being collected, how it's stored, steps toward risk management, and who takes responsibility for each stage of management. This will help instill a unified approach to dataset management while shielding the business from unnecessary risk.

Consider Your Training Data Platform

A good ML/Ops platform should make dataset management as simple as possible. With that in mind, get interrogative. Does your platform provider provide sufficient storage, particularly as you scale? How secure is their method of storage? What steps do they take to shield your business from looming compliance changes? How simple is it to use? Does it allow you to rapidly extract information, or is it like wading through fog?

Security and Compliance

Countless models and products rely on confidential or restricted data, which comes with a series of regulations and controls. With this in mind, you need to be mindful of who gets access to sensitive data, what they're allowed to do with it, and how the data needs to be stored. Here it's important to source a platform that gives stakeholders the ability to manage data access while having controls in place to prevent a breach.

Consider Your Workflow

Ease of data management extends far beyond the day-to-day satisfaction of your project stakeholders. It directly impacts the pace, efficiency, and effectiveness of your product pipeline. Certain platforms, like V7, provide a host of dataset management features that not only uphold the integrity of your data - but allow you to maximize its value. This includes:

- The ability to search and filter millions of items, without compromising on speed.
- Assigning tasks to hundreds of labelers by priority, objects detected, or manually. This also allows you to track the progress and solve issues by communicating in real-time.
- Version control.
- Support for any image or video format, including JPG, PNG, TIF, MP4, MOV, SVS, DICOM, NIfTI, and more.
- The ability to track the progress of annotations.
- The ability to investigate the overall composition of your dataset, allowing you to troubleshoot model performance.

Customer of V7

ABYSS

"Annotation systems are playing catch up trying to project where research and cutting-edge development would be and how to make it configurable. Luckily, solutions like V7 came up, which are maturing in the industry, helping companies like ours to scale up and stay ahead with our R&D."

Suchet Bargoti
CTO, Abyss Solutions

BUILDING AI

Training and Deploying Models

By now we've covered the landscape of enterprise AI, the hurdles that lie in wait, and solutions available to streamline the process. Finally, we've arrived at the training and deployment of models.

In this section, we cover how to create a model within V7, which could be used to automate workflows, conduct QA, or enter production via the API. It's worth noting here, that V7 provides models exclusively for object detection, instance segmentation, and classification models.

Training and deploying models is a broad and complex task, with countless methods, frameworks, tools, and platforms. For the purposes of this guide, we're going to keep things simple and outline how to train a model within the V7 platform itself.

TRY V7 NOW →

How to Train a Model in V7

A reliable training data platform will make the process of training your first model a relatively simple one. Beware that these AutoML models are recommended for experimentation and annotation purposes more so than as production-ready models you may develop on Pytorch, Tensorflow, or a similar framework.

Below we break down in steps how to train a model in V7.

HOW MUCH DATA DO YOU NEED?

In V7, you can train models on as few as 100 labeled instances, though it's recommended to feed your model with enough data to deliver precision and performance.

WHAT CAN I USE MY MODEL FOR?

You can use this model to label additional data, test labeling quality and quantity, in addition to ultilizing the model in production via the API.

WHERE DO I START?

To get started, head to the model tab, pictured below.

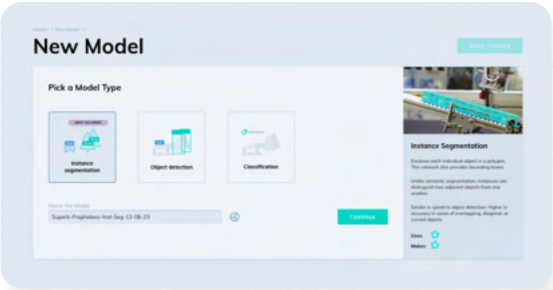

Here you'll need to select the type of model you'd like to train.

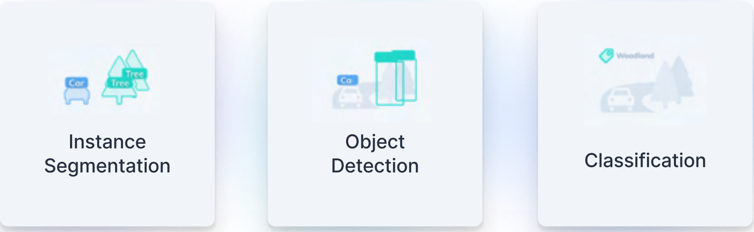

This type of model encloses each individual object in a polygon, around which will be an enclosing bounding box. Instance segmentation models are trained on the polygons in your training data, and are only compatible with annotations made in V7 with Auto-Annotate, the Brush Tool, and the Polygon Tool.

This type of model encloses each individual object in a bounding box. Object Detection models are trained on the bounding boxes and/or polygons in your training data, and are only compatible with annotations made in V7 with the Bounding Box Tool, Auto-Annotate, the Brush Tool, and the Polygon Tool

This type of model provides the most likely tag given to an image, based on its pixel content. Classification models are trained on the tags in your training data, which can be easily added to files in bulk from the Data page, or one by one in the workview of any file. At least two tag classes must be selected for training in order to train a classification model.

Choosing a Dataset

From here, you'll need to select the dataset you'll be using to train your model. In this example, we're creating a model that can detect bird species. Here, you can see that if you toggle by "Instances/Images" you can see the classes in this particular dataset, and how many instances exist of those classes. It's worth noting, that the quantity of quality of your data is the bottleneck of the performance of your model, so it's important to feed your model with a minimum of 100 instances - but ideally, the more the better.

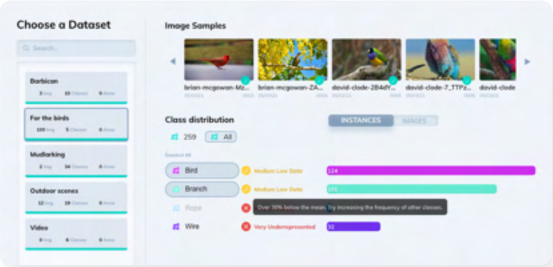

Training your model

Once you've hit continue, you're ready to train your model. At this point, it's important to remember the division between your training dataset, your test dataset, and your validation dataset. Below, you can see a 80/10/10 split, across training, validation, and test.

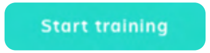

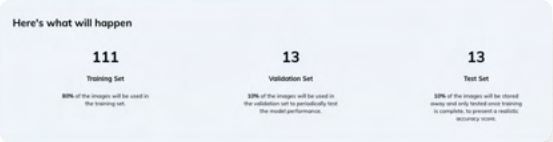

This will help you to reduce overfitting and deliver a more performant model.

OVERFITTING VS. UNDERFITTING

- Underfitting occurs when you have a high bias in your data, i.e., you are oversimplifying the problem, and as a result, the model does not work correctly in the training data.
- Overfitting occurs when your model has a high variance, i.e., the model performs well on the training data but does not perform accurately in the evaluation set. The model memorizes the data patterns in the training dataset but fails to generalize to unseen examples.

Understanding your results

Once training is complete, you'll have a number of metrics that will allow you to assess the performance of your model, before testing it out. Below we've covered this in a little detail.

LOSS: ALL MODEL TYPES

The first we'll focus on is Loss. Every type of model trained in V7 will display a loss figure and loss curve. The number that V7 displays is the loss function at the latest training epoch. The lower the loss, the fewer mistakes the model is making. Here's what we want to see:

The loss curve of a well-learned model should follow an L-shape, decreasing sharply at first and approaching a flat line over time. This is a visual representation of your model making fewer mistakes over time.

MAP: INSTANCE SEGMENTATION AND OBJECT DETECTION MODELS

Mean Average Precision (or "mAP") is measured by taking the mean of all average precisions (the area under a Precision vs Recall curve) across all IoU thresholds and for all classes. This metric provides an overall model performance, irrespective of any manually-set threshold.

ACCURACY: CLASSIFICATION MODELS

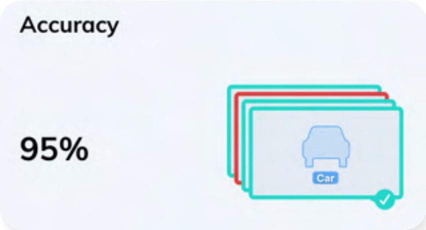

Accuracy is the percentage of predictions your model got right in the test set. It's computed by dividing the number of correct predictions by the total number of test examples.

If your metrics are looking good, then it's time for the fun to begin. You can now run your model through the API, and use it to automatically label your data.

DEPLOYING THE MODEL

Once you're happy with your model, it's ready for deployment. The process for this is fairly simple and can be broken down into two parts: generating an API key and pasting your command from V7.

The same can be said for publicly available models. If you have the model ID for the public model on V7, simply plug in your API key, image path, and the model's ID to get a response in milliseconds.

Ethics, Risks, and Compliance

Model development comes with a series of ethical quandaries, looming risks, and compliance challenges. From AI bias and model drift to data protection and regulatory limitations, you'll need to house your development in a considered strategy that de-risks your product.

In this section, we tackle core ethics, risk, and compliance challenges, and outline how to tackle them.

DEFINING YOUR COMPANY'S ETHICAL STANCE

Subject to endless discourse, debate, and - at times - derision, AI has faced a wave of controversial headlines. Perhaps the most impactful of these was an open letter demanding a pause on AI development. Cosigned by the likes of Elon Musk, the letter cited "profound risks to society" and the risk of AI-driven "propaganda and untruth".

While the letter garnered considerable interest, some researchers condemned the use of their work within the letter, while other signatories were revealed to be fake.

The letter also sparked debate regarding the apocalyptic and often inflated language around AI advancements such as LLM developments. So much so, that a deep learning engineer from Google countered with, "a 6-month moratorium on people overreacting to LLMs (in either direction)."

Between the extremes of apocalyptic AI and "a utopian technology", there's a middle ground, that favors nuance, evidence-based thinking, and ethical alignment. For those grappling with the ethical challenges of AI, solutions are rapidly coming to the fore, in the form of regulatory shifts dedicated to mitigating the harms of AI, while fueling its ability to better human life.

For enterprises wading through the ethical thunderstorm acting out in the papers, it helps to tackle your AI ethics stance from the get-go. In doing so, you set foundations to shield your business from reputational, legal, or financial risk.

To get started:

Define your company's ethical stance on AI

Where does the business stand on particular uses of AI? Which limitations does the business promote? Why is your business using AI?

Define your due diligence process for AI product development

What is the process for embarking on new AI projects? Who takes accountability? How is it documented and interrogated?

Educate your enterprise on the ethical challenges of AI

How do you convey your company's stance on AI? Are your people properly informed of the difference between unethical AI and commercially viable experiments? Who takes responsibility for ethical and legal breaches - such as non-compliant use of data?

DEALING WITH MODEL BIAS

AI can be powerful, but it's not without its faults.

For example, AI bias (when models produce data bias due to using biased data) and societal bias (when models reproduce harmful social biases that impact the end user.) One infamous example of this was Amazon's automated recruitment system, which became biased against women as a result of the data fed into it. To combat AI bias, it's important to:

- Use a diverse and representative training dataset.
- Ensure your dataset is sufficient in scope and accurately reflects the real-world use case that the model will interact with.
- Continuously monitor and evaluate the performance of your model.
- Leverage features to continually improve performance, such as RLHF, or workflows equipped to flag errors.
- Ensure cases underrepresented in reality that carry high detection importance (such as a cancer cell, manufacturing defect, or linguistic expression) are more significantly present in the dataset (a case of positive bias).

Failure to address model bias comes with a long list of risks: limited model performance, reputational damage, an erosion of trust, a loss of financial reward - the list goes on. Our advice? Define and implement a process that can reliably monitor, root out, and mitigate bias.

> *"Like all technologies before it, artificial intelligence will reflect the values of its creators. So inclusivity matters — from who designs it to who sits on the company boards and which ethical perspectives are included. Otherwise, we risk constructing machine intelligence that mirrors a narrow and privileged vision of society, with its old, familiar biases and stereotypes."*
>
> **Kate Crawford**
> Principal Researcher at Microsoft Research and Co-Founder of the AI Now Institute at NYU
> *"New York Times: Artificial Intelligence's White Guy Problem"*

DEALING WITH MODEL DRIFT

AI can also be prone to model drift, which refers to a decline in a model's performance as a result of environmental shifts. This is broken down into two types, concept drift, and data drift. Concept drift happens when patterns or relationships surrounding the data evolve, resulting in the model needing to adapt its underlying understanding. Data drift is when the data the model is trained on, shifts and differs from the data that the model is used on.

To combat model drift, it's important to:

- Continually monitor your model output, with a commitment to iterating the model throughout its lifecycle.
- Continually source new data to retrain the model.
- Implement rigorous QA processes that make model drift easy to spot and fix.
- Leverage feedback loops, RLHF, and user testing to maintain the performance of your model.
- Implement a human in the loop process to catch low-confidence results or poor user feedback.
- Implement a continuous benchmarking system, such as BenchLLM for language models.

Model drift can be particularly detrimental to businesses, particularly those leveraging AI for commercial purposes. Without continual iteration, models experiencing drift can wreak havoc on a use case, producing inaccurate predictions that can cause long-term harm.

Regulatory and Compliance Evolutions

As AI continues to evolve, governments and legislators globe-wide are racing to keep up with it. The regulatory and legal landscape for AI is continuing to form and it remains to be seen where the pendulum will finally swing.

Just this year...

→ In March 2023, the UK government released its policy paper, "AI Regulation: A Pro-Innovation Approach" which outlined a commitment to governing and regulating AI in a way that promoted innovative use cases.

→ Similarly, in May 2023, the US government defined its stance, outlining a necessity to balance innovation with placing "people and communities at the center by supporting responsible innovation that serves the public good".

→ Finally, in June 2023, European Union lawmakers voted to approve the EU AI Act, which aims to ban systems posing "an unacceptable level of risk."

It highlights three things: a unified appetite for responsible development, a race for AI innovation, and a keen understanding of AI's looming impact.

> BIZTECH LAWYERS
>
> *"The EU is the first legislator to approve comprehensive AI legislation, but many other countries have this in their sights. Governance frameworks, guidance and draft regulations have been published in many jurisdictions, including Australia, Japan, China (which already has some limited laws in this area) and Brazil, to name but a few.*
>
> *The common theme is the imposition of additional obligations on businesses developing and using AI to ensure that humans are protected from the worst risks. Rules around governance, quality control, record keeping, reporting and transparency feature prominently, with a focus on those systems that pose a high risk to key human rights."*
>
> 
> **Alison Berryman**
> Head of UK Technology for Biztech Lawyers

The Future of Development

The fear of staying behind the competition by not introducing new solutions in time is real—and we get it. Currently, Large Language Models and generative AI are revolutionizing the market, with no end in sight. From solving small tasks to potentially transforming business processes, these technologies can no longer be ignored by any organization, and beg the question about the future of AI.

Many experts predict that generative AI will have an enormous impact on the market and productivity. For example, McKinsey researchers estimated that generative AI could add around $2.6 trillion to $4.4 trillion annually across 63 different use cases.

In this section, we'll provide a quick overview of LLMs and generative AI as well as their current use cases. We'll discuss how to leverage their impact and how to prepare for future AI transformations.

Foundation Models and Generative AI: An Overview

Foundation models are pre-trained on vast datasets and can be used out-of-the-box or fine-tuned to your needs. Therefore, it doesn't take much computing power to adapt one to your company's needs. Additionally, thanks to the huge amounts of training data, the output is of much higher quality than in the case of traditional models—they can generate original content, understand causality, and more.

With the traditional approach to AI development, each model has to be trained from scratch. You have to prepare a sufficiently broad, well-labeled dataset, which takes a lot of time and resources. And the larger the dataset, the more powerful the infrastructure you need to actually train your model.

With foundation models pre-trained on large amounts of unstructured data, you can build your AI products with just minimal fine-tuning to your use case. This reduces the amount of data you need in order to build a high-quality AI product. Additionally, foundational models increase flexibility—it's easier to attune them to different use cases.

Should Your Company Adopt LLMs and Generative AI?

As always in the world of AI, things change fast, and we're still waiting for the Large Language Model dust to settle. However, with generative AI solutions taking the world by storm, sleeping on them may leave you behind your competition.

"LLMs have the potential to transform how enterprises handle their operations. They can create more intuitive and natural interactions with machines, extract knowledge more efficiently from text, and automate content creation and analysis.

This can lead to significant improvements in productivity, decision-making, and overall user experience."

Baris Aksoy
General Partner at AVB Ventures

A recent Salesforce report shows that a majority (57%) of senior IT leaders are prioritizing generative AI for their business in their business plans while one-third (33%) are making it their top priority.

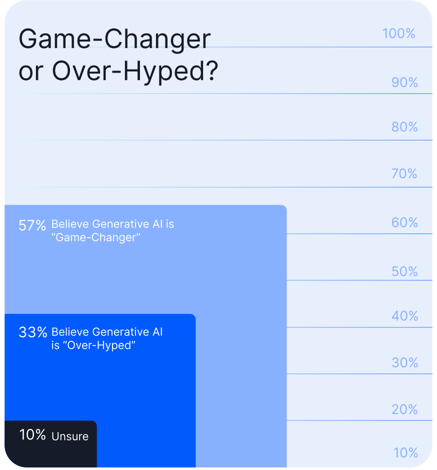

Here's what you should be mindful of before adopting generative AI into your company:

Educate yourself

This probably goes without saying, but the more you know about foundation models, and what they're actually capable of, the better you'll be able to assess what the right solution for your business is.

Do market research

Are other companies in your industry jumping head-first onto the generative AI bandwagon? Why? Does it reflect a market shift, or does it need further probing? Assess whether the risk is worth the gain.

Take full control of your data

Make sure you have solid data infrastructure in place, and that you've covered all legal grounds. Always review the data you're about to use to remove bias and false elements. You can learn how to manage your data from previous sections of this guide.

Keep humans in the loop

The majority of your models' output should be reviewed and approved by a professional. That's how you'll make sure you retain the highest possible quality of your model, ensuring its reliability, safety, and security.

Collecting user feedback and using human-in-the-loop methods for quality control are crucial for improving AI models over time and ensuring their reliability and safety. Capturing data on the inputs, outputs, user actions, and corrections can help filter and refine the dataset for fine-tuning and developing secure ML solutions.

Alberto Rizzoli
Co-Founder of V7

Top Foundation Models and LLM Use Cases

With LLMs democratizing access to machine learning solutions thanks to reduced costs and lower technological requirements, you have a chance to try out previously unavailable solutions in your business.

It's important here to seize the opportunity, however, with the technology still transforming before our eyes, proceed with caution and mitigate risk where possible.

The benefits businesses can gain with generative AI include:

- Improved efficiency through AI Co-pilots (the models can generate original outputs and adapt faster).
- Cost reduction (using foundation models reduces training data and infrastructure needs).
- Greater versatility (the models can be adjusted to new use cases easier, no need to train from scratch every time).

Let's quickly go through a few top solutions you can try to implement in your AI product development process.

GETTING INSIGHTS FROM UNSTRUCTURED DATA

LLMs make it easy to extract meaningful insights from your data and enhance information retrieval from your internal documentation, policies, and other knowledge reservoirs. You can also extract specific data to build other models.

For example, you can ask an LLM about an HR policy buried within a 100 page document, or across a company intranet. You may also ask it to summarize a large quantity of text or financial data.

FACILITATING MULTI-MODAL AI SYSTEMS WITH CONVERSATIONAL INTERFACES

LLMs can aid in designing conversational interfaces for multi-modal AI systems, improving the interaction between humans and AI. For instance, doctors can converse with a DICOM image, or analysts may chat with an LLM that has looked through the annotations of hundreds of city satellite images to survey objects.

GENERATING NEW CODE

In software development, LLMs can serve as programming assistants. They're able to provide code completion suggestions, detect errors, and come up with solutions to coding challenges. They can also assist in generating code snippets, improving readability, and preparing documentation.

AIDING IN PRODUCT RESEARCH

Generative AI can help automate data analysis, and identify trends and data patterns. It can perform sentiment analysis of users' opinions, or run simulations in pre-production to speed up the research and development process.

Toolkit for Implementing Foundation Models

Here's a list of tools we recommend you check out before you start using foundation models in your company →

ChatGPT/OpenAI API

The GPT models are able to handle content creation, coding, online communication, and other tasks surpassing the abilities of traditional NLP.

Implementing this powerful API into your business operations is cost-savvy, fairly easy (thanks to emerging no-code solutions), and flexible, making it a great experimentation ground.

LangChain

LangChain is a framework for developing LLM-based applications. LangChain provides standard, extendable interfaces, external integrations, and end-to-end implementations enabling off-the-shelf use.

PineCone

The Pinecone vector database makes it easy to build high-performance vector search applications. Vector embeddings are crucial for building long-term memory for foundation models, and Pinecone is designed to help manage the vast data.

Chroma

Chroma is an open-source database for managing embeddings in AI applications, focusing on Large Language Models (LLMs). Chroma aids LLMs the ability in building and maintaining long-term memory.

Pre-Trained Models and Frameworks: HuggingFace, LLaMA, and More

Both paid and free-to-use community-based solutions are on the market, facilitating LLM adoption for companies that don't have the infrastructure to train their own models. While not all of them are currently viable for commercial use, the market is worth monitoring—and it's worth checking how the ML community tackles even the most complex tasks.

What is the Future of AI Development?

The AI landscape is changing rapidly, and it's crucial to keep your eyes peeled. Observe, educate yourself, and adapt responsibly.

As we've already mentioned—don't introduce any solutions before a thorough business analysis and examination of your nfrastructure.

If you set out to train your own foundation models, start with diverse and representative data to minimize the risk of bias. Similarly, be mindful of privacy and security concerns—particularly with emerging technologies still undergoing regulatory shifts. And, of course, continuously monitor model performance to ensure that it is meeting your needs.

> *"In 2 years, every software product will have an AI Copilot, and AI development will become just regular software development. We strive to be the trusted data layer to these thousands of valuable AI assistants ready to learn unique business challenges."*
>
> 
> Alberto Rizzoli
> Co-Founder of V7

> *"AI Product development can cause the single biggest disruption to market dynamics since the advent of code. Ensuring fast, quality, safe product development has to be of paramount importance to every forward-looking product leader."*
>
> 
> Matt Brown
> Head of Sales at V7

BUILDING AI-POWERED PRODUCTS

Your Checklist

- [x] Identify the need or pain point that your project is solving.
- [ ] Clearly define the purpose, context, and goal of your AI solution.
- [ ] Probe into the challenges your project will face.
- [ ] Create an action plan to mitigate potential risks.
- [ ] Define your enterprise AI strategy.
- [ ] Define your company's ethical stance on AI.
- [ ] Educate your team on the strategy and ethics, and ensure a unified approach.
- [ ] Assess the existing tools, frameworks, and infrastructure available.
- [ ] Interrogate available training data platforms.
- [ ] Prioritise safe and secure integration with systems.
- [ ] Define a unified approach to development, what constitutes success, and how it can be repeated.
- [ ] Create a reliable and continuous data pipeline.
- [ ] Define and implement an efficient and accurate data labeling process.
- [ ] Boost efficiency and accuracy with automation, while keeping up to date on cutting edge features.
- [ ] Train your models, with consideration for retraining, drift, and feedback loops.
- [ ] Keep informed on shifting legal boundaries, and ensure your process is well-equipped for transparency, explainability, and the ability to pivot.

BROUGHT TO YOU BY

START BUILDING AI NOW →