SparkAI

Building vs. Buying an Edge Case Solution

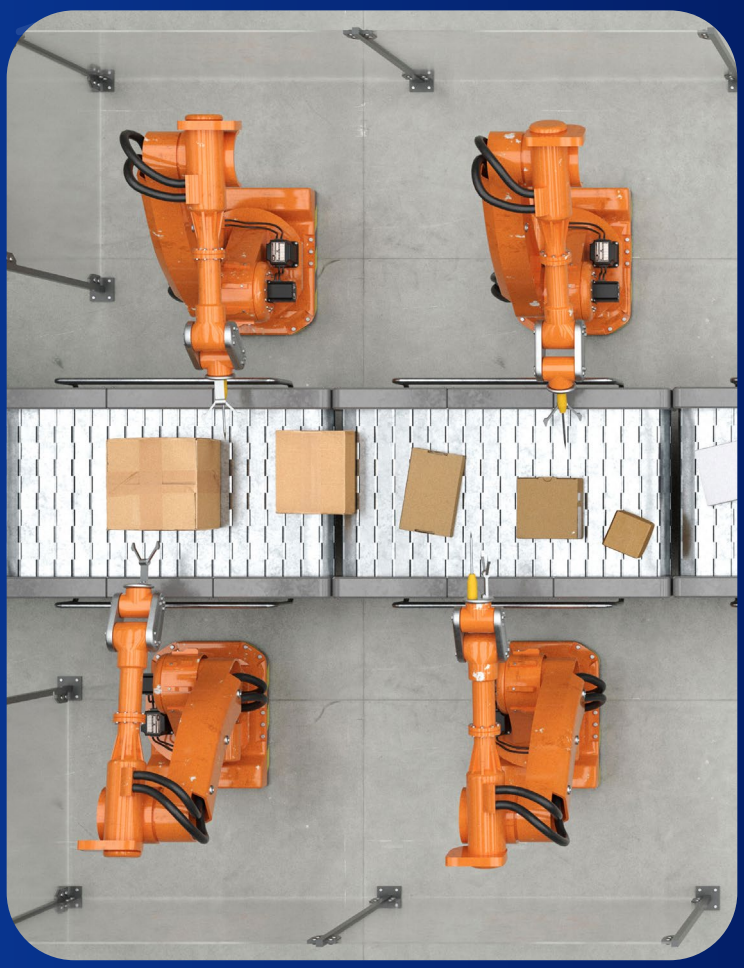

Executive Summary

AI systems are highly capable but extremely brittle. Even at their best, they are susceptible to failures on inevitable edge cases – their achilles heel.

Edge cases hamper otherwise functional AI products, inhibit uptime, slow throughput, put human safety at risk, and degrade the value of models.

Creating viable AI products requires an edge case solution.

The question this ebook answers is whether such an edge case solution should be bought from a third-party or built within your organization.

Table Of Contents

What Are Edge Cases?

AI models are trained on large troves of data. But regardless of how large a model's training dataset is, the real-world environment in which it operates is inherently unpredictable. This makes it impossible to train for the long-tail of all unexpected situations a model will encounter, i.e., edge cases. While it's true that AI is becoming increasingly capable, it's also incredibly difficult to perfect.

Edge Case

/'edʒ ˌkeɪs/ (noun)

Edge Cases are instances where an AI application cannot provide a prediction or make a decision with the confidence level that was defined as the acceptable threshold. Edge cases are often real-world scenarios that deviate from the norm — either not accounted for during model training or not well represented in training — making them difficult to impossible to incorporate into the model training process.

Related terms: AI exceptions, low-confidence predictions, anomalies, false positives

Edge Cases Impede AI Systems

Private investment in AI doubled from 2020 to 2021 – going from around $43 to $93.5 billion. Some project that the market will continue to expand at a compound annual growth rate (CAGR) of 38.1% from 2022 to 2030.

Despite such significant investment, over 85% of AI projects never yield value. Edge cases are a major contributor to the failure.

Even a slight AI misstep or incorrect prediction in response to an edge case can cause problems, especially in high-risk environments where edge cases correlate with human injury, property damage, or significant financial loss.

AI Failure Examples Across Industries

Precision Agriculture

For an autonomous tractor in the field, a cow that has wandered into the crops or a dust cloud occluding the tractor's vision sensors.

- Overuse of pesticides
- Inefficient harvesting
- Increased risk of accidents
- Financial loss

Manufacturing

For a manufacturing inspection model, mistakenly identifying a usable part as defective when the light shines on it at an odd angle.

- Missing defective products on the assembly line
- Increased risk of accidents
- Over-reliance on human intervention

Supply chain & logistics

For an industrial robot trained to move boxes in a warehouse, an anomalous box with unusual patterns, textures, or protrusions.

- Inventory mismanagement
- Decreased throughput
- Damaged goods
- Miscalculated product

Automated checkout

For a smart cart check-out system, an edge case might be products added to a cart that do not match the physical profile on record.

- Overcharging or undercharging of customers
- Incorrect product inventory

Insurance

For a model trained to detect in-ground swimming pools, an above-ground pool could be misidentified as a different property feature.

- Fraud
- Financial loss
- Delayed risk assessments

You will find plenty examples of well-publicized mishaps induced by edge cases:

1. **Zillow** dramatically overpaid for houses, resulting in the shut down of their Zillow Offers division.
2. **Tesla Autopilot** mistook the moon for a yellow traffic light, potentially endangering rider safety.
3. **IBM Watson** for Oncology recommended a patient with severe bleeding take a drug that would actually accelerate the bleeding.

But even when severe risks don't exist, an unreliable or underperforming AI system that stalls when presented with an edge case will jeopardize the commercial viability of an AI product.

SparkAI's Founder and CEO Michael Kohen sums up the consequences of the edge case problem well:

"An autonomous construction vehicle that can't recognize dusty obstacles can't launch. A bin picking robot that gets confused by misshapen boxes can't scale. A structural inspection CV model that can't tell the difference between a crack and a seam can't be trusted. As rare as these moments may be, they're enough to degrade confidence in the entire system, blocking products from commercialization, and draining immense R&D resources along the way."

Michael Kohen
Founder & CEO at SparkAI

The Downstream Impacts of Edge Cases

Edge cases render AI products unviable, which has knock-on effects at multiple levels:

Customer-Level Impacts

① **Compromised Safety**
When AI systems fail, people can get hurt and expensive assets can get damaged. Compromised safety will almost certainly delay or prevent model deployment, and it can tarnish the reputation of your brand.

② **Degraded Customer Confidence**
Customers need to be confident that your AI system will work as expected and will not disrupt their existing systems or workflows. AI that does not work 100% of the time will not persuade your customers to pivot away from their mature, stable, and trusted systems. It's an unmitigated risk to them.

Product-Level Impacts

③ **Diminished Reliability & Performance**
Customers expect a fully functioning system, 24/7. AI systems that fail unexpectedly on edge cases can't meet this expectation, and therefore need a way to be redirected towards success. Without external human intervention, they will shut down, stall, go offline, and fail in unpredictable ways. Furthermore, unreliable and underperforming systems quickly become a resource sink as their failures need to be identified and fixed. Flaws may replicate and reinforce poor model decisions over time, causing even more problems.

④ **Delayed, Premature, or No Deployment**
Deployment is a minefield for AI products. Knowing that your product doesn't work as well as it should will lead to delayed deployments that set your business back relative to competitors. Either that, or risks can be so great that you will be unable to deploy your model at all. Worse yet, you will deploy prematurely, risking potentially irreversible business level impacts.

Business-Level Impacts

⑤ **No Return on Investment**
It takes a lot of resource investment to build AI products. An AI product that fails or falters, can't launch or scale, and you can't reach ROI. Your company's efforts will go unrewarded and future AI investments will be hard to secure.

⑥ **Damaged Brand Reputation**
Poorly functioning AI systems could irrevocably damage brand reputation, especially when there are significant reliability and safety risks involved.

Why Can't AI Solve Its Own Edge Case Problems?

Even the most advanced AI systems struggle to reason through the unexpected.

The core barrier to competence is that there is a fundamental cognitive gap in how AI models are built and trained. Put differently, AI lacks cognition. And, we are unfortunately several groundbreaking, once-in-a-decade innovations away from AI being capable of solving the edge case challenge.

> "The moral of this story is clear: just because something manages to appear intelligent for a moment or two doesn't mean that it really is, or that it can handle the full range of circumstances a human would."
>
> Gary Marcus, "Rebooting AI"

How Does One Solve Edge Cases?

In contrast to AI, humans are naturally adept at dealing with unexpected situations. It's one of the many advantages that hundreds of thousands of years of evolution have conferred on our species.

Resolving AI edge cases requires a contrarian but elegantly simple solution:

Incorporating human cognition into the AI loop.

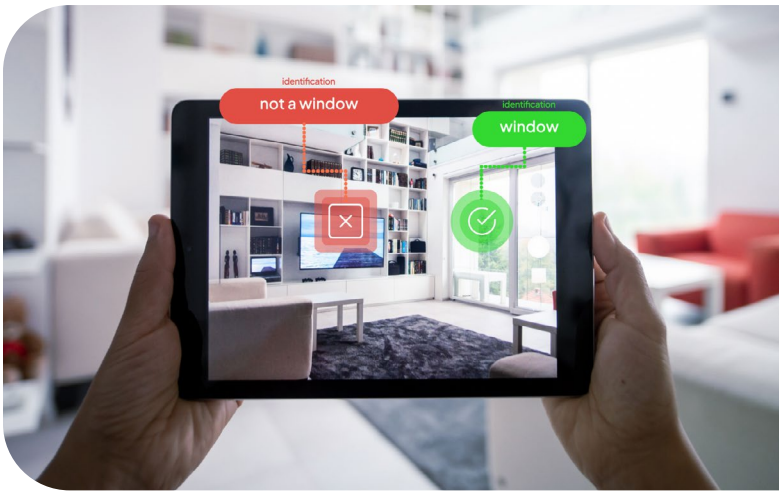

> "Deep learning is a very different beast from a human mind. At best, deep learning is a kind of idiot savant, with miraculous perceptual abilities, but very little overall comprehension."
>
> Gary Marcus, "Rebooting AI"

6 Ways to Approach the Edge Case Problem

There are at least six ways to incorporate humans into the AI loop to potentially address edge cases. However, serious speed, cost, and accuracy tradeoffs make most of them unworkable. Some constrain your ability to launch or scale, while others fail to lead to true, AI-based autonomy. Let's break them down.

Approaches that Delay Launch and Scale

Options 1 through 4 below are incomplete, imperfect, and ultimately infeasible approaches. They are workarounds at best. Let's examine each of these half-measures at a high level:

1. **Deciding to delay launch and remain in R&D mode** just pushes the problem into the future, with no plausible benefit. Because edge cases are so diverse, and because the world has a tendency to surprise us in ever-unexpected ways, the likelihood that you'll ever launch if you're always chasing that error-free horizon is slim to none.

2. **Launching without an edge case resolution plan** in place is perhaps just as ill-fated as remaining in R&D mode. Deploying without a strategy for handling edge cases leaves you vulnerable to safety, compliance, and reliability risks. The question of whether you will encounter these is not a matter of if, but when.

3. **Leaning on customers to resolve edge cases** might seem like a natural compromise. You get your product into the hands of customers early, and they get to resolve edge cases in the way that makes sense to them while informing the system of its desired behavior. Yet, customers will rapidly grow tired of shouldering the operational burden of constantly addressing individual edge cases. What this ultimately means is that you'll be blamed for providing an inferior and unreliable product.

4. **Remote piloting** is not an appealing option for AI-based robotics. By taking over for a robot in an uncertain scenario, the human operator can direct it through a challenging situation. However, this ultimately ends up being a replacement for autonomy rather than a supplement to it. In other words, this approach by design ensures that your product will deviate from the path to autonomy. This approach is also unscalable.

Learn more about each of these approaches, their benefits, and their drawbacks in this brief guide:
↳ Strategies to Launch and Scale in the Face of Critical AI Edge Cases.

Approaches that Accelerate Launch & Scale – Edge Case Solutions

In contrast to the four approaches on the previous page, the following two approaches are real-time solutions to edge cases that will help you launch and scale your AI products with confidence.

5. **Buying an edge case solution**

6. **Building an edge case solution in-house**

The remainder of this ebook compares and contrasts building versus buying a real-time edge case solution.

Buying

Buying an edge case solution is by far the simplest, fastest, and most cost-effective choice. In doing so, you gain access to all the resources and knowledge that a dedicated team has poured into resolving the edge case problem.

Building

Building your own edge case solution demands extreme engineering and operational expertise, as well as an unwavering focus on building your own in-house, human-in-the-loop team that focuses around-the-clock on solving edge cases.

Essential Requirements

Any edge case solution must have the following characteristics:

1. **Speed** — The ability to deploy, scale, and introduce new features rapidly in production so that you are solving edge cases while generating ROI and creating value for your customers.

2. **Accuracy** – The ability to ensure the uptime and reliability of your autonomous system, ensuring it makes sounder decisions on the fly.

3. **Safety** – The ability to address safety issues, such as dangerous machinery that might fail and harm people.

4. **Cost Savings** – The ability to save you money by helping you bridge commercial, engineering, and operational gaps to launch and scale.

Building an Edge Case Solution In-House

Building is no trivial matter. It requires a very serious commitment to people, process, and platform.

The Benefits of Building In-House

Control
You get 100% control over your workforce, operations, and engineering. Completely build your solution to your needs at whatever pace, and following whatever priorities you choose.

Satisfy uncommon privacy constraints
You have incredibly strict data privacy requirements that preclude working with a third party altogether. (For example, only those with a top-secret security clearance are permitted to access data.)

The Cons of Building In-House

Building a solution from scratch might at first blush seem like the obvious way to go. However, most companies do not realize the immense drawbacks to building a solution in-house.

Are you willing to invest years to develop something that is already commercially available at a fraction of the cost?

Major Drawbacks of Building In-house:

① **Inescapable cost inefficiency.**
The number of edge cases requiring real-time resolution may fluctuate unpredictably throughout a given day. To always meet this demand and ensure your AI products are supported, you need to staff your human operations to your demand peaks. This means that at all other times outside of those demand peaks, your workforce will be severely underutilized and inactive. You'll be effectively paying for people to sit around, an exorbitantly expensive labor cost. But hiring fewer workers would stretch your workforce too thin when demand peaks occur, letting your AI systems fail.

② **The operational burden of sourcing, managing & scaling a workforce.**
Once you've built your solution internally, you now have to operate it. This is where things really start to hurt. You now need to find a way to recruit, vet, train, performance-manage, and scale a distributed network of people who can handle edge cases as they arise. When someone departs, you need to replace them and train the incoming individual as quickly as possible to ensure continuity. You need to manage shifts, working hours, and payment schedules. Throughout all of this, you need to ensure constant uptime for your deployed products.

③ **An entirely non-core effort.**
Building a solution internally is immensely distracting and costly, and it diverts engineering resources away from highest value activities. Your focus is (and should be) on building the best performing AI and robotics products possible. That's what your engineers joined your company to do. Diverting their attention to build and grow an operationally intensive side project is not the best use of their time and your resources. Plus, it is unlikely that your engineering team will be excited about focusing on non-core competencies like building APIs and UIs that maximize efficiency for an edge-case resolution workforce. It is incredibly hard to hire and keep top-tier engineering talent. You do not want to shoulder them with boring work that they were not hired to do. This will hurt motivation and ultimately lead to employee churn. All of this – and so much more – is anathema to engineering organizations.

④ **Negative impact to business growth.**
Your ability to grow as a business is now tightly coupled to your ability to scale a non-core, operationally intensive solution. If you land ten new customers tomorrow, instead of focusing on servicing them, you now need to worry about how to scale up the workforce and infrastructure behind the real-time edge case solution that you manage. This risk can severely impact your growth potential and jeopardize the success of new deployments.

What Actually Goes Into Building an Edge Case Solution

Let's look at what all goes into building a minimum viable edge case solution. You can't do it piecemeal. You must build all of these things in order to properly solve edge cases in real-time.

The information below depicts a typical one-year build and assumes you will need to hire all of the staff to do so; however, your build may still take longer. The following chart does not factor in the cost of iterating and maintaining the solution after one year, which you will certainly have to do.

Operations

Recruiting workforce

- Sourcing
- Onboarding
- Workforce vetting
- Contracting

Training workforce

- Workforce training
- Ongoing re-training
- Skills building
- Knowledge continuity

Scaling workforce up/down

- Annual seasonality
- Unexpected surges
- Daily variability
- Optimum scheduling

Utilization efficiency

- Demand modeling
- Load balancing
- Idle-time overhead
- Coverage gaps

People management

- Payments
- Performance
- Incentives
- HR

Engineering

Intelligent job routing

- Job queuing
- Availability
- SLA optimization
- Context switching

Quality assurance

- Trust score
- Reporting
- Voting & audit
- Ground truth check

Workflow development

- Custom UI design
- Perf. optimization
- UX efficiency
- Workflow logic

Real-time decision tools

- Automation
- Workflow tooling
- Prediction
- Escalations

Ongoing development

- Outages
- Design updates
- Tech upgrades
- Bugs

Workforce

Injecting Human cognition

- Situational assessment
- Edge case resolution
- Contextualization

Continuous training & upskilling

- Platform training
- Business context understanding
- Use case training

If You Build It Yourself, Costs Quickly Explode

Role	Average Number	Cost
Engineers Technical lead Quality assurance lead Front-end developers Back-end developers	8	$ 1,256,000
Operations Operations managers 24/7	3	$ 420,000
Workforce Mission specialists 24/7	9	$ 351,000
Infrastructure		$ 100,000
Total First Year Cost Estimate		$ 2,127,000

An Overview of the Core Components of an Edge Case Solution

Platform

Providing fast responses requires building a customized, sophisticated toolkit. You will need to develop a platform that workers can easily use and that can handle displaying many different types of edge cases in a way that promotes rapid feedback. The need for 100% uptime and reliability necessitates placing engineers on call in case the platform suddenly breaks. You will need to A/B test many different types of platform UIs in order to determine what works for workers. AI needs to underpin the entire platform such that the power of your job routing, quality assurance, and real-time decision tools are augmented and enhanced by your own machine learning models.

Engineering

Companies that go down the arduous path of building an edge case solution in-house come to a problematic decision-making crossroads:

1. **Repurpose existing technical staff to build an edge case platform**
2. **Hire entirely new staff to build it**

Repurposing existing staff

You can potentially cut down the time and expense to engineer an edge case solution by having your existing team build it instead. This will certainly distract them from working on the company' core product, which will slow product development down significantly. As mentioned previously, this can lead to a reduction in the quality of their work, slowing of their productivity, boredom, and ultimately employee churn.

Hiring new staff

Hiring and recruiting top-level tech AI engineering talent is a massive challenge. Competition is fierce, and there's currently an undersupply of qualified engineers and tech workers.

Under ideal conditions, it will take over one year to build an internal edge case solution and the technical team to build and manage it. On average, the time cost to build out a custom software product is 4–12 months. However, this assumes that you have a team already in place. More complex projects can take even longer and throw a wrench in your product roadmap – 85 percent of software development projects go over schedule.

Before you can even begin building, it will take roughly 6–12 months to stand up an 8–person scrum team. This is the minimum team size for building a custom software product for enterprise.

Whether you leverage your existing engineering team or hire new staff, building an edge-case platform presents its own, unique challenges that will likely require new hires to fill knowledge, skills, and abilities gaps.Once finally built, there is constant maintenance effort required.

Operations: The Burden of Sourcing, Managing & Scaling a Workforce

In addition to technical staff, you need an operations team that can recruit, train, and manage a scalable workforce around the clock.

Consider also how much training you and your operations team will need to do. By definition, edge cases are unpredictable, and so each product and its environment will experience unique edge cases that must be dealt with in customized ways. Workers need to be trained in how to resolve edge cases correctly for the given context, with absolute fidelity to business requirements. Training is even more of a hassle when you have to build your personnel and platform from scratch.

Everything is harder the first time you do it, and you will make many mistakes that will cost you time, money, and employee churn.

Workforce

In addition to the technical staff and operations team, you need to hire, train, and manage a workforce that can respond to and resolve edge cases 24/7. At a minimum, you will want to ask yourself the following critical questions:

- **Do I have the capability to hire and manage such a workforce — one that is available 24/7 and can provide instant feedback?**
- **If I do have this capability, how much will it cost?**
- **How many such workers will I need to handle my edge case load, and how will I scale that workforce?**
- **Am I willing to pay these workers to have them at the ready, even when no edge cases are actively coming in as a result of natural ebbs and flows in operation or usage?**

Questions To Ask Yourself Before Building In-House

If you answer “no” to any of these questions, then you should consider buying an edge case solution instead.

- Do I have $1–2 MM laying around to fund an internally-built solution in year 1 and maintain it in subsequent years?
- Am I ready to expend the resources to do all that is required to build, maintain, and scale my own edge case solution?
- Can I build something as good or better than edge case solutions on the market, even though those vendors work on edge case resolution full time?
- Am I okay with distracting my engineering team from their current highest priorities to build this?
- Am I comfortable with not having an edge case solution provider to rely on when I need to troubleshoot or brainstorm complex problems?
- Am I ready to build, train, and manage an on-demand and scalable workforce?
- Am I willing to forego early deployment and first-mover advantage by waiting months or years to build my own solution?

Still thinking about building?
↳ Send this to your team and answer these questions before going any further!

Buying an Edge Case Solution – The Best Way to Launch and Scale

A huge benefit of buying an existing edge case solution is that the provider already has a workforce in place and is distributing the workload of many clients across it. In this way, they are not paying workers for downtime, and they can easily scale up to meet periods of peak demand.

While edge cases may differ from product to product, a purpose-built solution provides universal mechanisms by which any edge case can be sent in real-time to a trained workforce of humans that can give immediate context to the AI system.

When buying an existing solution, providers will work with you to establish criteria for resolving your edge cases and will then handle the burden of training workers to enforce these criteria.

For most businesses, buying is the only sensible, cost-effective, and quick option that lets your product launch and scale within weeks, not years.

Buying an edge case solution is by far the simplest, fastest, and most cost-effective approach to solving the edge case problem in real-time, in production. It gives you immediate access to a ready-made, fully-featured product that has already been optimized for performance, reliability, and scalability.

Buying Makes Sense When You Need...

- A viable AI product in the near-term.
- Real-time edge case resolution within weeks or months (not a year or more), so that your product can quickly generate ROI.
- First-mover advantage to secure your position in a competitive market
- To deploy with little internal engineering effort so that you can keep your engineering team focused on developing your core product.
- Zero operational exposure related to hiring, managing, & scaling a human workforce.
- The ability to seamlessly scale edge case resolution capacity up or down based on fluctuating business requirements.
- To decouple your product's ability to launch and scale from the development and maintenance of an edge case solution – and all of the engineering and operational efforts involved.

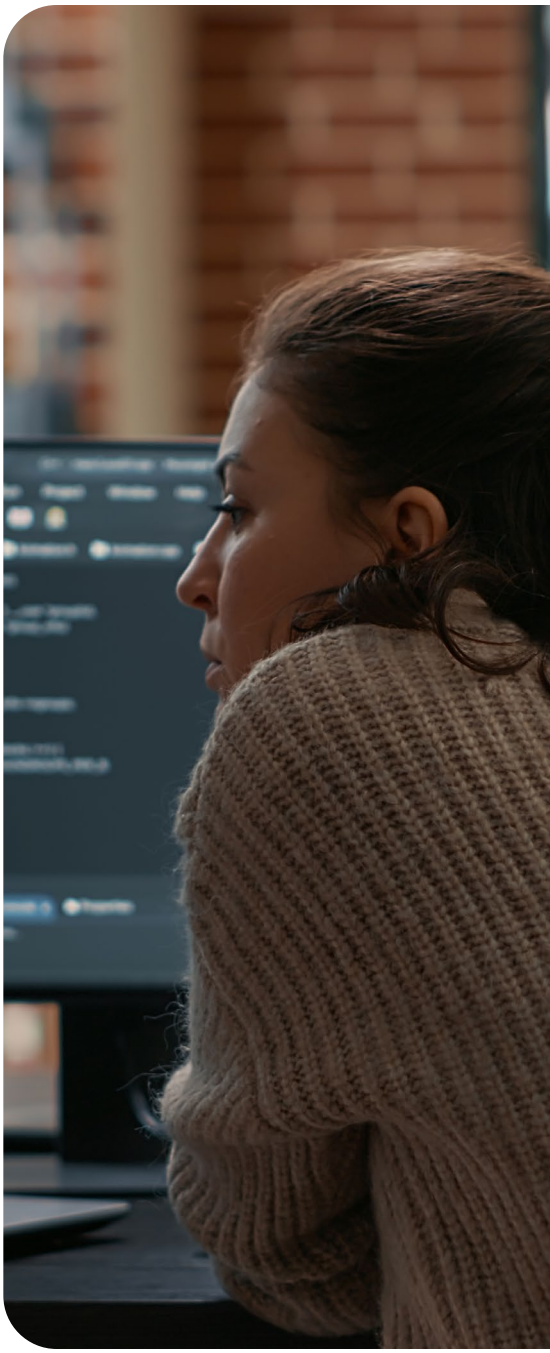

Benefits to Buying

- Zero downtime for your AI system.
- A stable, tested system developed over years by an experienced, trusted partner that works with you every step of the way.
- The ability to rapidly scale in response to demand.
- The peace of mind and freedom to focus on the core of your business: Your AI product. Not distracted by side projects (e.g., building an edge case solution).
- An enormous amount of time and money saved.
- No need to hire, train, and manage expensive teams.
- 24/7 access to an experienced, specialized workforce that is focused exclusively on quickly and accurately solving edge cases.
- Continuous model improvement as it gains exposure to new edge cases with context provided by humans in the loop.

Benefits to Building

- Control. You get 100% control over your solution's workforce, operations, and engineering.
- Satisfy uncommon privacy constraints. You have incredibly strict data privacy requirements that preclude working with a third party altogether.

Buying SparkAI — The One Solution for Resolving AI Edge Cases in Real-Time

		Building a Solution In-House	Buying SparkAI
Speed	Deploy AI products today	▮	▮▮▮
	Scale faster with real-time edge case resolution data	▮▮	▮▮▮
Cost	Avoid building & maintaining a system for resolving edge cases		▮▮▮
	Focus internal engineers on the core product & high-ROI efforts	▮	▮▮▮
Accuracy	Use expert augmented cognition that combines tech & humanpower	▮	▮▮▮

SparkAI is a complete, ready-to-deploy solution that delivers everything you need to instantly leverage real-time human cognition in your AI workflow. We combine people and technology to resolve your critical AI edge cases, false positives, and other exceptions, live in production, so that you can launch & scale automation products faster than ever.

We developed our technology and methods – which we originally perfected in the self-driving car industry – to fill the cognitive gaps in AI that obstruct your progress and cost you money. We have refined our technology platform over years of dedicated effort and it is configurable to the specific nuances of your use case.

SparkAI is the only real-time edge case solution on the market.

SparkAI Injects Human Cognition Into Your AI Workflow

The right approach for engaging human cognition in the AI workflow is to regard it as nothing more than an input – a contribution presented to the AI system for consideration. When called, SparkAI's role is not to take over control (as in remote piloting). What we deliver are the missing contextual cues that AI systems sometimes need as inputs to make confident, real-world decisions in the face of unexpected challenges.

Here's how it works: In moments of confusion, instead of breaking down, your AI system calls SparkAI. Your system sends situational data (image, video, audio, etc.) to SparkAI via API. In real-time, trained human mission specialists, augmented by our technology, resolve the difficult-to-discern details that are confusing your AI. SparkAI returns this resolution to your AI in real-time, enabling it to make a confident decision in the live, production environment.

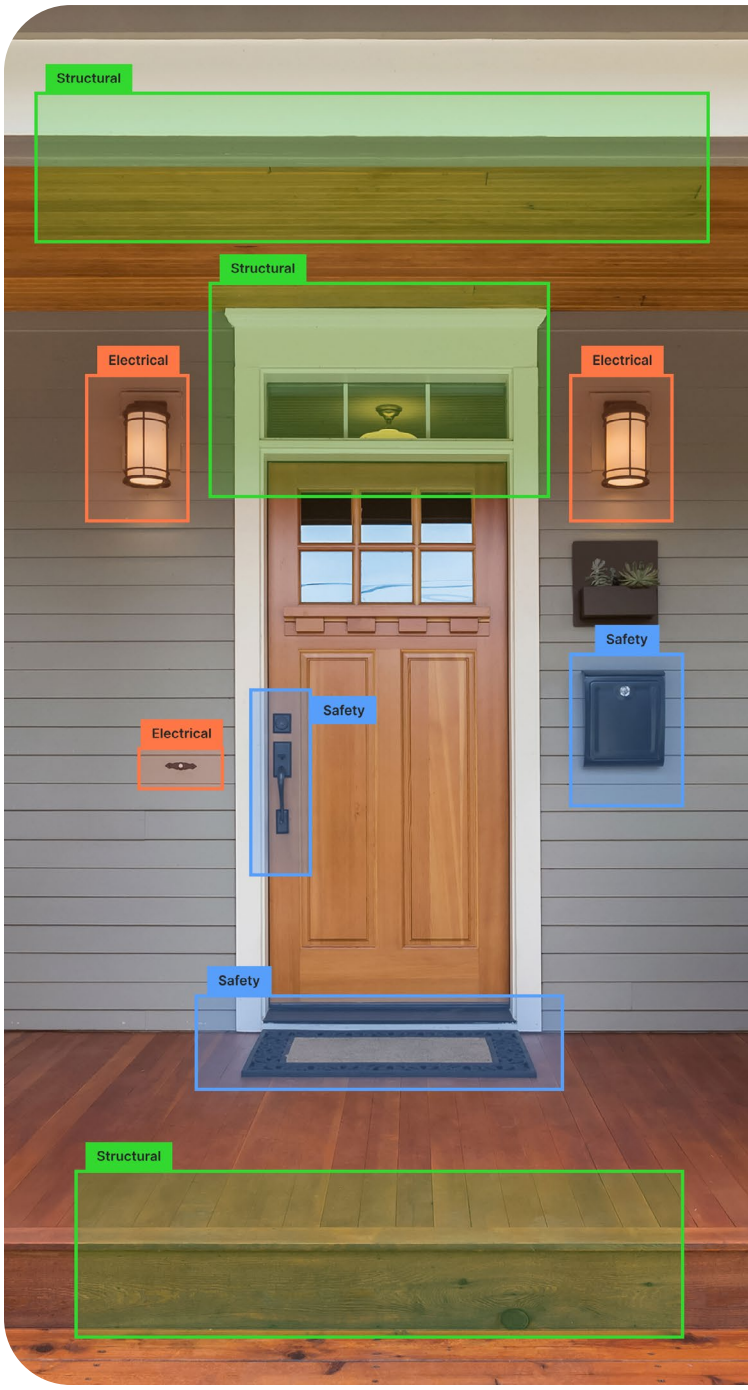

SparkAI Accelerates Your Path to Autonomy and Scale

Our approach de-escalates the high stakes race against the clock facing your product and engineering teams. It gives them the breathing room and time to develop AI products thoughtfully, without the do-or-die business pressure from the top.

In addition to a real-time decision that helps your AI system overcome live moments of confusion, SparkAI also gets you production-grade ground truth data to continually improve your models. This accelerates your ML training flywheel, fast-tracking a process that helps your data science team learn more about the situations that your models will face in the real world.

See How it Works

SparkAI Drives Outcomes

Meet our Customer

A **Fortune 100 equipment manufacturer** sought to introduce a fully autonomous tractor — a first-of-its-kind product — to enable farmers to amplify productivity and reclaim time and energy. In moments of low-confidence, we resolve difficult-to-discern details about the scene (e.g., obstacles in the tractors' paths) and return resolutions to the vehicles.

The Results

We accelerated our customer's introduction of autonomous tractors, allowing farmers to capture the full value from automation today. Following their product announcement, the company's stock price rose sharply and they garnered a broad base of recognition from agriculture and AI industry experts as an epicenter of automation innovation.

"The feeling of seeing a tractor do a job for a person without anybody doing anything is totally magical," our customer says.

24/7	21.08	1,280
real-time coverage	median response time (in seconds)	engagements per hour

Meet our Customer

This **Fortune 500 Robotics OEM** developed a next-gen automated depalletizing solution that efficiently identifies and processes large and diverse parcel loads. Using computer vision and an AI model, the system determines the precise way to pick parcels. In moments of confusion, the robot calls SparkAI, sending imagery and accompanying metadata via API. Our system and human mission specialists review the edge cases, identify potential pick points for confusing boxes in the pallet, and return the missing context to the robot.

The Results

SparkAI's cost-effective, real-time solution unlocked market viability for our customer. Their product now operates uninterrupted, in dynamic environments, at all hours. They've rapidly expanded their share of this $3.7 billion market. SparkAI's speedy implementation lets our customer focus resources on ongoing, high-value development.

24/7	<10	1,000+
real-time coverage	median response time (in seconds)	engagements per hour

Meet our Customer

Our customer is a leader in AI-powered insurance technology. Their flagship product is a smart phone application that uses computer vision to assess and underwrite property risk, instead of relying on expensive in-person inspections. This unlocks a more cost-effective and accurate assessment.

The Results

With edge cases solved, our customer is outperforming competitors and quickly capturing even more share of the fast-growing $843 billion casualty and direct insurance market. Knowing that SparkAI is here to solve their edge case problems, our customer is looking toward their future offerings with refreshed confidence and boldness. They aim to launch new AI features and products in shorter cycles, which will provide them with more immediate return on their AI investments and further bolster their market position.

24/7
real-time coverage

63
classification categories

120
engagements per day

For most businesses, in most scenarios, buying a ready-made solution is undoubtedly the best option.

Buying Is Best

The pressure on companies developing AI systems is immense. On the one hand, the pace of technological acceleration demands that businesses incorporate AI into their processes on ever-decreasing timelines.

Yet the same voices calling for speed demand more accuracy. AI models are entrusted with incredibly impactful decisions upon which entire companies, and even human lives, depend. They are naively expected to operate error-free, even though edge cases make such perfection impossible.

The only feasible solution for businesses that want to utilize AI and maintain their competitive edge is to either build their own edge case solution or buy one.

By doing so, business leaders can capitalize on the years of dedicated thought, focus, trial and error, and engineering-plus-operational follow through that a product such as SparkAI provides. In addition, they gain access to a dedicated and scalable workforce of edge-case resolution specialists that are available day and night to help resolve the complex situations with which AI systems are being faced.

In leveraging an existing solution, leaders save their organizations millions in potential costs and many multiples more of time, energy, and focus. It is difficult to underestimate the clarity, freedom, and peace of mind that comes from going all in on your business' core competencies while offloading ancillary concerns to dedicated teams that have made resolving edge cases the entirety of their business mission.

Our Mission

Revolutionary automation products shouldn't have to wait for the next AI breakthrough. Our mission is to accelerate their proliferation everywhere, by advancing companies faster along the path to commercialization. Our business is built on your success getting to market and growing there. We're here to help you realize your vision of the future, and ignite an imagination for what more you can do today.

See SparkAI in Action