

Solving the AI Edge Case Problem

Strategies to Launch and Scale
in the Face of Critical AI Edge Cases



Why Do Most AI Projects Fail?

Billions of dollars are poured into AI every year, yet most projects never make it out of the lab. With little to show for such a substantial investment, everyone from engineers and executives to investors is rightfully asking:

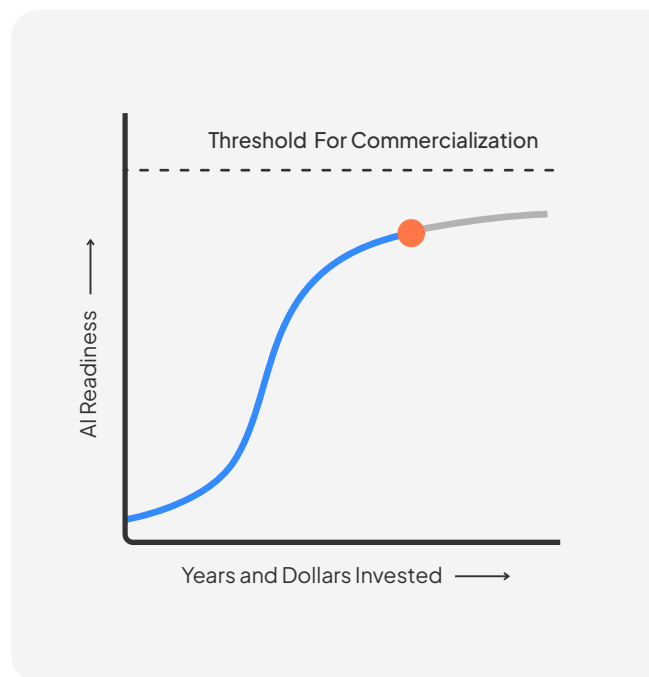
What's holding AI back?

Edge Cases – A Critical Point of Failure

The real world is infinitely complex and constantly changing. This makes it difficult to build and commercialize AI products – robotic or otherwise – that perform accurately and reliably.

Every company doing anything interesting with AI, from warehouse automation to robotic delivery, faces this same reality: AI is hard to perfect.

Though AI products are trained for a variety of circumstances, the unstructured environments in which they operate – farm fields, warehouses, sidewalks, or inside homes – are inherently unpredictable. This makes it impossible to train for the long-tail of all potential scenarios, known as edge cases.



“Your dataset has to be really representative of the environment you’re in. You’ve got sunny days, big trees, small trees, light soil, dark soil, ditches, cows, an ice cream truck that just pulled in. These things are incredibly hard to think about and train for.”

Willy Pell, VP Autonomy, Blue River Technology

Only 53% of projects make it from prototype to production

Source: Gartner 2020²

Only 14.6% of firms have successfully deployed AI capabilities

Source: NewVantage Partners 2020³



As a result, even the most advanced AI systems struggle to reason through the unexpected. At its core, the barrier is that AI is missing cognition. There is a fundamental cognitive gap in how AI models are built and trained, and we are several groundbreaking, once-in-a-decade innovations away from solving the edge case challenge. With safety, scalability, and reliability at stake, it represents a critical barrier to launch & scale.

87% of machine learning projects never enter production

Source: VentureBeat 2019¹

Commercializing AI products today is about assembling building blocks. One of those building blocks is a solution for overcoming edge cases. AI cannot survive in the real-world without it.

The right edge case resolution strategy enables you to:

- Launch & scale your products today, with AI that's still imperfect
- Overcome performance gaps preventing your product's commercial viability
- Introduce new competitive features earlier than otherwise possible
- Offset the need for operationally intensive and expensive on-site human support
- Keep your engineering teams focused on highest value activities
- Accelerate performance of existing products to deliver more value to clients

➤ What are edge cases (and other key AI terms)?



Meet SparkAI

The founding team of SparkAI came together from Zoox, UberATG, and other top self-driving car programs to address the single most pervasive problem impacting the automation industry: AI is hard to perfect. The real world often looks unpredictably different from what AI models are trained to expect. That dissonance can be confusing to AI, preventing companies from launching & scaling world-changing automation products today.

SparkAI is the first to adapt the solution that we originally developed to help self-driving cars to the wider market of automation applications. We combine people and technology in a lightweight API that resolves machine learning exceptions in real-time.

From robotic farming to computer vision-based insurance underwriting, injecting SparkAI into live, production AI workflows unlocks launch & scale for the next generation of automation applications.

Built to solve the challenges we encountered on the path to launch & scale at:



What Are the Approaches to Edge Cases?

You cannot launch or scale a product that can't deal with unexpected situations. Bringing AI products to life means filling this critical gap.

For AI innovators, then the question becomes: Which edge case approach should we use?

While there are several ways to approach edge cases, selecting the wrong one can:

- Degrade customer confidence and your brand reputation
- Delay your deployment – sacrificing revenue and ROI
- Negatively impact system performance
- Drain your resources and R&D spend
- Lead to potential safety and security issues
- Limit scale and business growth
- Create a competitive disadvantage

That's why it's so important to understand the benefits and drawbacks of each approach.

Let's look at the approaches and help you discern what is best for your AI product.



“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 seam can't be trusted.

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 along the way.”

[In AI is Missing Cognition. We Deliver it in Real-time.](#)

Michael Kohen, CEO, SparkAI

How to Tackle AI Edge Cases

- 1 **Don't launch; stay in R&D (training) mode**
Page 7
- 2 **Launch without an edge case resolution strategy**
Page 9
- 3 **Lean on customers to resolve edge cases**
Page 11
- 4 **Remote piloting to solve edge cases**
Page 13
- 5 **SparkAI**
Page 15
- 6 **Build in-house**
Page 18

“You need to understand the value of solving edge cases... Should you even solve it to begin with? Does the customer care that you solved it? Maybe you don't need to solve all these edge cases, but maybe you do. Once you've resolved that question, then you need a method to figure out how.”

Jeremy Groeteke
Head of Computational Agronomy, Syngenta Group

Stay in R&D (Training) Mode

What it is

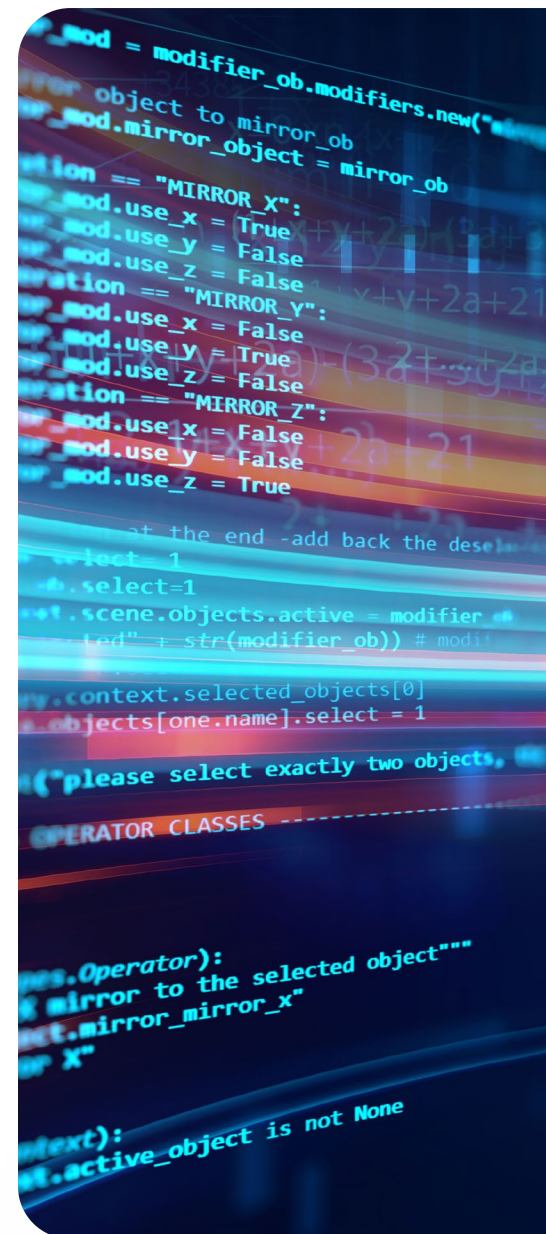
Staying in R&D mode means staying in the lab. Instead of releasing your AI system for real-world use, you continue training the model with more annotated data until eventually the frequency and severity of edge cases decreases (but are never completely eliminated).

Potential benefits

Staying in the lab lets you focus on fine-tuning your model and resolving commonly known or easy-to-foresee edge cases. This is particularly important for safety critical applications (like autonomous vehicles) operating in highly unstructured environments, in which real human life is at risk and the technology is still not quite there.

This may appeal to you if...

- You have an unlimited budget and no business pressure to launch your product
- You're comfortable with no real-world results for long periods of time
- Revenue from the AI system isn't crucial to your business (now or in the near-term)
- Waiting won't sacrifice your industry position, market share, or expectations you have already set with your consumers
- As a result of waiting to launch, your system will perform infinitely better





Success stories in data labeling have already proven that humans are a critical part of the AI equation at massive scale. While these services have helped build vast training libraries, the next million data points are not what will turn your ROI plateau into a hockey-stick.

The bottom up approach of hard, honest ML development powered by training data, needs to be paired with a top-down approach that enables teams to deliver not just the features they're currently working on, but the they hope to work on 5 years from now...

Drawbacks

Where are the results?

Staying in R&D mode means no revenue generation or ROI. Any AI business leader knows the downstream effects: Demoralizing the internal team. Losing investor confidence. Running out of operating budget. Chilling future internal investment. Ultimately, going out of business without a proven product in the market.

The lab alone will never solve the edge case problem

In a controlled environment, your model will never get the experience it needs to confidently tackle the diversity of real world edge cases. Sheltering your AI product in the lab means you're not seeing how it functions in customer environments. You're not studying how and when it fails, and you're not collecting the critical production data that you need to continue training your AI over time.

The world doesn't benefit from your product

Customers won't get to experience your innovation, and that's not what you promised them. People create AI products because they want to change the world, not sit on the sidelines.

Hand your competitors the edge

While you wait, your competitors will release similar technology before you, hurting your chances of leading the market. This threatens your investment, while also guaranteeing hefty sums of additional R&D costs to catch up.

Launch without an Edge Case Resolution Strategy

What it is

Instead of keeping your AI system behind the scenes, you launch your product without a strategy to resolve edge cases. This means that your product will knowingly fail in certain situations.

You may choose to disclaim this to your customer upfront, or you may not disclose it and just plan to deal with the customer complaints as they happen. Either way, you're launching with a degraded feature set that you anticipate will have some undesirable impact on you.

Potential benefits

This approach is all about speed. It's about prioritizing a meets-minimum approach to launch over creating an edge case resolution strategy that bridges AI gaps. You save R&D time and money, which helps you beat competitors to the finish line.

It also gives customers a chance to use your system to address their pain and problems — something that is impossible if you don't launch. With deployment, you might gain support for your product and insights into the edge cases troubling your system.

This may appeal to you if...

- Immediate revenue from your product is crucial
- A quick launch attracts press headlines and supports your market positioning
- Safety concerns aren't a limiting factor (this approach may not work for autonomous vehicles or surgical robotics, for example)
- Customers won't be overly impacted or bothered by potential failure (e.g., AI-powered ball-tracking technology that confused a soccer ball with a referee's bald head)
- Despite its flawed state, your tech is an improvement over market alternatives



Drawbacks

1. Erosion of customer confidence

When your system suffers reliability and performance problems, your customers will have to deal with the consequences. When customers don't see the value that was promised, they may start to wonder: "Why did we buy this? Why do we have to deal with major disruptions? How do I get my money back?" That same customer might bring that enduring distrust into future conversations.

2. Potential PR problems


You'll need a plan to manage the customer complaints, as well as potential PR nightmares from accidents.

3. Hesitancy to buy

Customers may decide to wait — or look for a solution that provides better reliability — a blow to your revenue. For customers who place a heavy premium on safety and reliability, your product could become a non-starter and your brand could be tarnished.

4. No chance for continued development

If you deploy your system with no edge case resolution strategy behind it, you have to maintain it in a reactive operational state. Putting out fires will take up your resources while your technology struggles to learn. Until you formulate a strategy to deal with it, your edge case problem can compound.



Lean on Customers to Resolve Edge Cases

What it is

In this approach, you rely on your customers to resolve edge cases for you (assuming you have the platform to pull it off).

This is how it works: in moments of low-confidence or confusion, an autonomous system automatically calls your customer for help. Depending on your application, your customer may be required to physically intervene (e.g. drive the robot past an obstacle) or may be able to remotely support (e.g. remotely approve a decision via an application). Either way, your customer is asked to pause what they're doing and attend to your faulty product.

Potential benefits

This approach may work if you have the right customers and set proper expectations. You can slowly expand your training data set to make modest improvements while the AI system is deployed. (Assuming your customers properly handle the edge cases.)

In some situations, leaning on customers for help can build better, longer-lasting relationships. Your users get to be actively involved in sharing what's working and what isn't. They also get to play a vital role in improving the technology they use.

This may appeal to you if...

- Your customers can and want to be part of edge case handling, and the cost:benefit of participation is low
- They can successfully use your product despite the edge cases, or the resolutions they provide
- Safety concerns aren't a factor



Drawbacks

1. Customers don't want to do this

They want a solution to their problem. This option gives the problem right back to them. Customers may become frustrated or stop providing necessary information. And, without their input, your AI may struggle to function, let alone improve. You are unlikely to acquire the data you need to resolve critical edge cases, and the outcome is often a low-performant version of the system.

2. Liability concerns

Leaning on your customers is a safety issue for many applications, like autonomous tractors, where the costs of error can be fatal. Not all customers have the knowledge or expertise to give the right level of attention to your product's edge cases. They may make hasty, haphazard, or incorrect decisions that impact the future performance of your system. Further, when safety is a concern, this could prove dangerous. The decisions your customers make will become your responsibility and impact your reputation. You could even be held liable for the outcome.

3. Disillusionment

For some customers, this level of involvement can create unexpected disappointment. The “magic” behind AI disappears when they have to spend a great deal of time reporting and fixing issues on their own. Over time, this degrades trust, confidence, and the perceived value of your technology.

4. High startup costs

There is a great deal of upfront work required before you can lean on customers to resolve edge cases for you. SparkAI was purpose-built to handle the platform and operations pieces that eliminate this concern.

“If all you're doing is saying ‘gimme, gimme, gimme,’ people start to wonder. You have to solve for the value exchange. They need gratification or that immediate answer,... They need to move on to solving their next problem.”

Jeremy Groeteke

Head of Computational Agronomy, Syngenta Group

Remote Piloting to Solve Edge Cases

What it is

In this approach to solving edge cases (primarily used for robotic AI applications), humans are used to directly control the AI system when it gets confused. There are two ways to do this: (1) physical, on-site human control, (2) remote, off-site human control also known as remote piloting. In both cases, a human is overriding the AI system and directly controlling the robots actions. This section will focus primarily on the latter of the two approaches.

With on-site control, a physical human operator is present to chaperone or otherwise support an active robot. In moments of confusion or error, the humans role is to physically solve the problem on-site in place of the AI. For example, a robot that fails to avoid a mud puddle and get stuck would be physically pushed out. An obstacle that a robot fails to classify confidently would be physically moved out of the way.

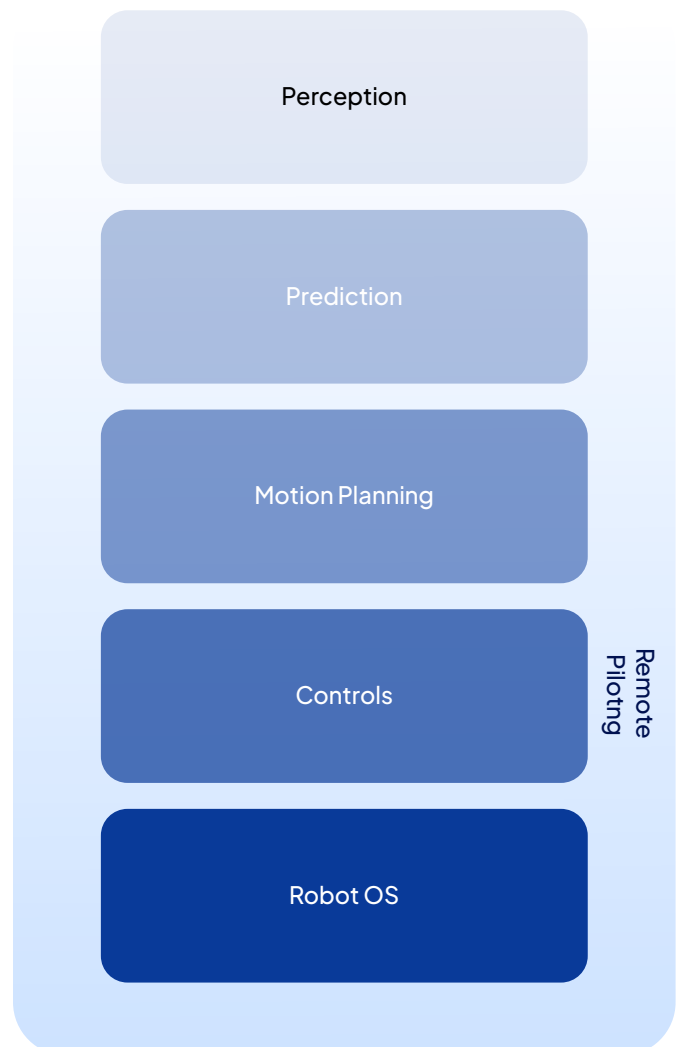
Remote piloting is similar in that it renders direct control over the AI system, but remotely. In moments of confusion, the robot relinquishes control to humans stationed remotely. Those humans have specialized equipment (like a joystick, steering wheel, pedals, etc.) that enable them to maneuver the robot or manipulate its actions, to overcome an edge case.

Across the industry, we've seen some companies champion this approach of starting with humans remotely piloting a vehicle, and then "layering in the autonomy." It usually does not work out that way.

The real world



How your robot interacts



Potential benefits

For those companies that do not aim to build autonomous systems, remote piloting is a shortcut to launching a product more akin to an RC car than an autonomous robot. In such limited cases, this approach can unlock cost and safety benefits and remove humans from situations in which they would otherwise be exposed to danger operating a vehicle “from the cab.”

Forklifts in warehouse settings are a great example of where human insight is required throughout operation and can not yet be fully automated. In this setting, remote piloting simply takes human forklift operators out of the cab and into an office.

This may appeal to you if...

- **Developing autonomy isn't your end goal**
- **You need a shortcut to launch that has some cost and safety benefits over in-person human operators (e.g., forklift operators)**
- **Humans need to be removed from danger, but you believe they still need to be in control**
- **You have the resources to dedicate to hardware, training, and personnel**

Drawbacks

1. Safety

Making the best, safest decision requires near-perfect situational awareness. Relinquishing control to a human a thousand miles away poses a serious risk. The smallest glitch in connectivity, or the slightest gap in awareness, can be dangerous.

2. Cost scalability

To deliver situational awareness to a human operator thousands of miles away, companies that have chosen this path now need to invest (heavily) in enabling it. To remotely pilot robots, human operators need specialized equipment – game-style joysticks, steering wheels, wrap-around monitors, etc. Expensive equipment like this needs to live somewhere, necessitating further investment in physical remote operations centers.

3. Operational scalability

Human operators responsible for safety-critical piloting require lengthy and specialized training. This constrains who can and cannot become an operator. Further still, each “intervention” requires a human to engage with a robot for minutes at a time. Ensuring no robot is ever left unsupported requires an oversized human workforce and a human-to-robot ratio that looks a lot more like 1:1 than the 1:many that automation promises. Beyond high costs, this isn't scalable.

4. Misalignment with your long-term mission

By leveraging remote piloting, you eliminate the discipline and urgency of architecting an AI system that learns to solve problems on its own. As you encounter challenges in development, it's tempting to shift more responsibility away from AI and onto human operators. This growing dependency diverts attention and resources, moving you further from developing genuine autonomy.



Real-Time Edge Case Resolution with SparkAI

What it is

SparkAI has generalized technology and methods – that 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 combine people and technology to resolve your critical AI edge cases, false positives, and other exceptions live in production, so you can launch & scale automation products faster than ever.

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) – AI systems are intelligent enough. What we deliver are the missing contextual cues that those systems sometimes need as inputs to make confident real-world decisions in the face of unexpected challenges.

How it works (simply): in moments of confusion, instead of breaking down, your AI system calls SparkAI. Your system sends situational data (image, video, etc.) to SparkAI via API. In real-time, trained human mission specialists, plus our tech, resolve the difficult-to-discern details that are confusing your AI. (For example: What is this obstacle? Which box should I pick? Is it safe for me to proceed?) SparkAI returns this resolution to your AI in real-time, enabling it to make a confident decision in the live production environment.





Capture new revenue

Launch new products with imperfect AI



Innovate for less

Bootstrap new models for evaluation



Deliver client value

Accelerate performance of existing products



Run with confidence

Audit & QA model outputs in real-time

Potential benefits

This is the only solution on the market for dealing with critical AI edge cases in real-time. The outcomes are momentous. Launch & scale your products today, with AI that's still imperfect. Instantly bridge performance gaps to commercial viability. Introduce new competitive features earlier than otherwise possible. Offset the need for operationally intensive and expensive on-site human support.

This approach also de-escalates the high stakes race against the clock facing your autonomy engineers. It gives them the breathing room and time to develop AI models 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 your models.

This may appeal to you if...

- You need real-time edge case resolution so your product is reliable and high-performing
- You need a fast, cost-effective, and simple solution that deploys via one line of code
- Launching today gets your system to consumers first and secures your industry position
- You want zero operational exposure (i.e. hiring, managing, & scaling a human workforce)
- You want to seamlessly scale capacity up (or down) based on your business needs

Drawbacks

1. People aren't perfect

When a human is involved in any process, the potential for errors exists.

SparkAI leverages multiple solutions to mitigate this risk. Consensus-based voting ensures that every human response is validated by multiple individuals. An algorithmic QA system ensures each individual human is consistently performing at or above baseline. Additionally, a SparkAI ML estimator makes its own prediction, which is compared to the human responses. Functional safety is also best addressed at a system level, whereby the SparkAI contribution is one of multiple inputs the system weighs when choosing an action.

2. Cloud connectivity required


Your system must be capable of communicating with SparkAI remotely. Live, continuous connectivity or a high bandwidth connection are not required. However, in moments of SparkAI engagement, your system must be minimally capable of calling SparkAI, sending a small data payload, and receiving a response from the cloud.

3. Response time

SparkAI response times are measured in single digit seconds. That's fast, but certain applications may require millisecond responses.

4. Low-lift initial setup

Initially, SparkAI enablement teams work with your team to design the workflow: What does the handshake look like between your AI system and SparkAI? In what scenarios will your system call SparkAI? What does SparkAI's response look like?



Building a SparkAI-like Service In-House

What it is

You could buy SparkAI and work with us. Or, you could (attempt to) build something like it internally. Building something like SparkAI in-house requires a very serious commitment to people, process, and platform. You'll need to build out an intelligent dispatch algorithm, real-time QA systems, workflow and decision tools, a robust and highly available API, a whole new operations team with a scalable workforce, modular GUIs and workflows, an efficient means for staffing and payment... And that's just the tip of the iceberg.

This may appeal to you if...

- You don't mind diverting your AI engineering team from building your core product in order to build something entirely non-core: a platform and operations team to chase the edge case problem
- Humans are a core part of your product value proposition, e.g. robotics security
- You need to iterate extremely rapidly on how edge cases are resolved
- Your edge case resolution requires niche experts with high levels of education, e.g., doctors, lawyers, etc.
- You have a specialized edge case resolution workforce [that you can easily scale up or down at anytime
- You have abnormally strict data privacy requirements that force all of your data to stay in-house



Drawbacks to Building an Edge Case Resolution Solution In-House?

In-house drawback	SparkAI solution
Immense setup cost and time to build	SparkAI is a complete, ready-to-deploy solution, delivering everything you need to instantly leverage real-time human cognition in your AI workflow. Our technology platform has been refined over years of dedicated effort and configurable specifically to the nuances of your use application.
Entirely non-core effort	This is SparkAI's core product. We are singularly focused on making our solution more performant, accurate, and cost-efficient, and continuously pushing our product forward. Doing this well requires the level of obsession that is central to our business.
Operational burden of sourcing, managing & scaling workforce	SparkAI already has its own, directly-managed workforce of humans (called Mission Specialists). Mission Specialists are trained specifically for your use case. Further, we've built and are continuously refining systems to automatically measure performance, provide feedback, manage schedules & continuity, and direct training. This enables us to manage a distributed workforce efficiently on a massive scale. Doing so is far from trivial, which is why SparkAI's core team is composed of industry veterans with decades of experience building and managing massively distributed human operations functions. All of this means you have zero exposure to the burden of human operations.
Inescapable cost inefficiency	SparkAI enables you to unlock fractional operations. By servicing demand from multiple customers centrally, we can realize efficiencies of scale that no company operating alone can access. Our Mission Specialist workforce is trained on multiple customer applications. If any one customer is less active in a given period, those resources are allocated to another customer's real-time needs. An intelligent dispatch algorithm continuously balances supply and demand, ensuring unbroken coverage for all customers (plus a healthy margin of safety). The result is that you only ever pay for what you need.
Impact to business growth	Like other infrastructure you buy, SparkAI scales seamlessly up (or down) with the demands and seasonality of your business. You have zero exposure to operational management, enabling you to focus on growing your core business unimpeded.

Simple Build vs. Buy Breakdown

While building an in-house platform (coupled with an operations team) might sound appealing in limited scenarios, it will always be an extremely challenging approach that involves enormous investment in people, process, and new technology. Here's what that tradeoff looks like:

Build

This is not a part of your core value proposition, so why would you drain your resources building it?

Plan

- Prioritize over core product

Budget

- Dedicate money
- Dedicate people
- Dedicate time

Build

- Split engineering focus
- New infrastructure development
- Custom interface design
- Ongoing maintenance
- Intelligent dispatch
- Real-time QA
- Workflow development

Recruit

- Sourcing
- Workforce vetting
- Onboarding
- Contracting
-

Schedule

- Demand modeling
- Idle-time overhead
- Load balancing
- Coverage gaps
- Annual seasonality
- Daily variability
- Unexpected surges

Train

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

Manage

- Payments
- Incentives
- Product performance
- People performance
- Outages
- Bugs

Scale

All of the above, but more of it...

Buy

Instantly activate a real-time edge case solution that scales seamlessly, with zero operational exposure to you.



Is it Time to Launch Your AI?

The next step is to choose the option that best meets your business requirements and AI product goals.

If you're ready to launch a new product or accelerate the performance of an existing product, then ask yourself these questions to determine the best way to launch your AI faster:

- What's the value of solving edge cases (to our customers, for our product)?
- How will we overcome the edge cases standing in the way of our AI launch?
- Does it need customer feedback to improve?
- Do we have the time, talent, and resources to build and maintain a real-time edge case-resolution system ourselves?
- If we wait to launch, what are the potential consequences?

We're here to help answer these questions and more.

We built SparkAI because we fundamentally believe in the promise of AI to advance human progress. Our mission is to accelerate its proliferation in everyday life, and our business is built on you successfully getting to market and growing there. We're here because we believe in you and the future you want to create – let's make it real, today.

[See SparkAI in Action](#)