RadarFirst

MY SUMMER OF AI: ACCURACY, ECONOMICS, AND SOVEREIGNTY


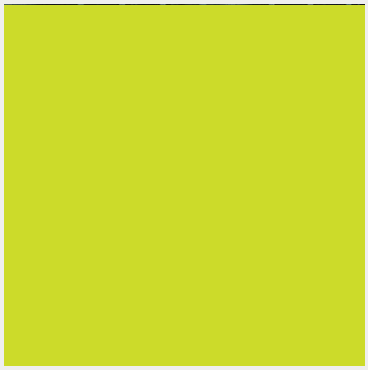



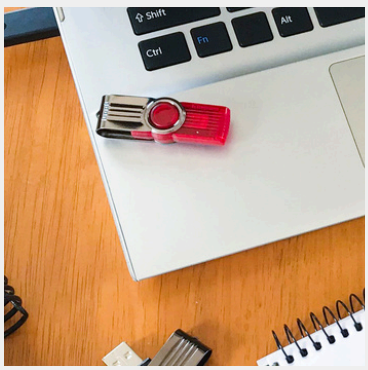
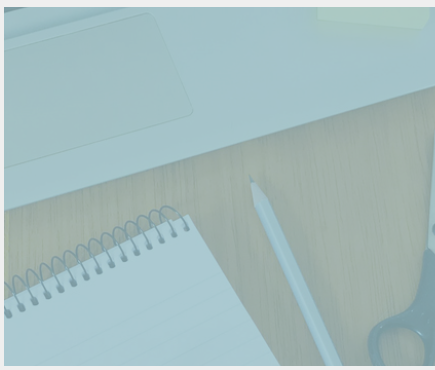

BY ZACH BURNETT,
CEO AT RADARFIRST

MY SUMMER OF AI

ACCURACY, ECONOMICS, AND SOVEREIGNTY

RadarFirst

As a child, I spent six weeks every summer at my grandparents' home outside Washington, D.C. Those weeks away from my parents became a formative part of my life.

Each summer, I was "encouraged" (notice the quotation marks) to read a book and write a report before returning home. At the time, it felt like some sort of punishment. But the lesson stayed with me: spend time being curious, then force yourself to explain what you learned.

This summer, I found myself repeating that exercise. The subject wasn't a book. It was the phenomenon of artificial intelligence, which I've come to think of as my "summer of AI."

I approached the subject from three perspectives: as a parent, an everyday user, and the CEO of RadarFirst, a privacy and AI incident management platform. I spend much of my time thinking about what happens when technology enters a real organization, where people, policies, customers, regulators, budgets, and consequences all come into play.

The more I explored, the more I returned to three questions:

- How do we know an AI system actually works?
- What does a reliable outcome really cost?
- And who controls the knowledge that makes the system useful?

Accuracy: Can We Depend on the Answer?

Like many people, AI has quietly become part of my daily routine. I now use large language models for everyday tasks: catching up on the news, finding restaurant recommendations, or checking whether a summer concert has been canceled because of the weather (which actually happened).

I also find myself speaking to my AI assistant the way a parent talks to a child: *"Are you really sure?"*

That question matters. Large language models are probabilistic. They generate likely answers, not guaranteed truths. That realization pushed me to understand how organizations can measure the reliability of an AI-generated outcome.

As the RadarFirst engineering team built our agentic layer, they introduced me to one useful starting point: Pass@k.

Pass@k measures whether an AI system can produce at least one correct answer within a given number of attempts. Pass@5, for example, asks: If the system tries five times, what is the probability that at least one answer will be correct?

A higher score suggests that the correct answer is within the model's capabilities. It does not mean the model will get the answer right the first time.

Comparing Pass@1 with Pass@k helps reveal whether a system is dependable on its first attempt or succeeds only after generating multiple alternatives. But Pass@k has an important limitation: it measures the probability of producing a verifiably correct answer. It cannot determine whether an answer that has not been independently verified is truthful.

That requires testing, evidence, and, in higher-risk situations, expert review.

This is only the beginning. The NIST approach to AI measurement and evaluation treats performance as multidimensional, not something that can be reduced to a single accuracy score. Related consistency measures, including G-Pass@k, ask a different question: Can the system produce the correct result reliably across repeated attempts?

Put simply, Pass@k asks whether the correct answer is somewhere within the model's range. Consistency measures ask whether we can depend on the model to keep producing it.

We could spend the rest of my "AI book report" on accuracy alone. For everyone's sanity, I'll leave it at this: measuring AI performance is complicated, contextual, and still evolving.

Leaders should resist asking only, "Is the AI accurate?" The more useful questions are: Accurate under what conditions? Reliable enough for what level of risk? And at what cost?

Economics: What Does a Reliable Outcome Cost?

Accuracy, however, is only one part of the equation. Even if an AI system consistently produces the right answer, organizations still have to ask whether delivering that answer is economically sustainable.

One of the first lessons I learned in software is that almost anything is possible. The real question is what it will cost.

The long-term economics of AI-enabled software fascinate me because they are changing in real time. My organization invests in access to multiple enterprise frontier models. Not long ago, some of that access resembled traditional, relatively predictable software licensing. Increasingly, the economics are usage-based.

MY SUMMER OF AI

ACCURACY, ECONOMICS, AND SOVEREIGNTY

RadarFirst

AI depends on chips, data centers, electricity, networks, storage, and highly specialized people. Every inference carries a recurring operating cost. Industry reporting on the substantial costs and uneven profitability of frontier AI companies reinforces a basic point: building and operating AI at scale is extraordinarily expensive.

Software companies spent decades building around relatively predictable SaaS economics. A seat license was finite and understandable. AI is pushing the industry toward a more complex model in which providers increasingly want to price based on realized value and, in effect, participate in the customer's return on investment.

Here is the challenge: a single customer request may trigger several model calls, retrieval operations, safety checks, and follow-up actions. Usage, and therefore cost, can rise precisely as the product becomes more useful.

If the customer expects a predictable price but the provider faces unpredictable delivery costs, the business model becomes difficult to sustain. Eventually, those costs must be absorbed, reduced, or passed along.

After my summer of learning, however, I believe this is only half the economic story.

Organizations also provide an enormously valuable input to make AI productive: their customer information, operating procedures, product strategy, institutional judgment, and years of accumulated knowledge.

To borrow Mastercard's famous phrase: priceless.

AI Sovereignty: Who Controls the Knowledge?

This led me to think about something I call the double payment.

The first payment is obvious. Organizations pay for tokens, subscriptions, computing power, and capacity.

The second payment is strategic. They provide the proprietary context that enables an AI system to understand and operate within their business.

If that context is so valuable, who controls it? Can the organization move it? Can it reuse that knowledge with another provider? Can it prevent the knowledge from strengthening a system that may someday compete with it?

ACCURACY, ECONOMICS, AND SOVEREIGNTY

RadarFirst

These questions sit at the heart of AI sovereignty: an organization's ability to control the data, context, models, infrastructure, and decision-making processes that make its AI useful.

I think about it like bringing your own bat to a Little League game. You arrive at the field, hit the home run, and take your bat home with you. The analogy may be simple, but the principle is important: the field may belong to someone else; the capability you brought to it should remain yours.

AI sovereignty will become an increasingly important enterprise decision. The right approach may involve contractual protections, private deployments, retrieval-based architectures, open-weight models, confidential computing, or some combination of them. The answer will vary by organization and use case.

What should not vary is the level of scrutiny.

Enterprises need to understand not only what they are buying, but also what they are contributing, and whether they can take it with them.

No one likes paying twice. AI should not be the exception.

The Leadership Questions Ahead

My summer book reports taught me that curiosity becomes useful only when we can explain what we have learned.

After my summer of AI, my conclusion is that leaders should evaluate AI through three connected lenses:

- **Does it produce an outcome we can trust?**
- **Can it produce that outcome at a sustainable cost?**
- **And do we retain control of the knowledge that makes the outcome possible?**

There is so much to explore around the use and governance of AI. As children go back to school and summer slips away, I find myself in relentless pursuit of accurate AI models both in my personal and professional life. Without correct outcomes, the cost does not align to value.

ACCURACY, ECONOMICS, AND SOVEREIGNTY

Cost needs to be considered for both the value of the knowledge you bring through data and pattern recognition as well as the power, people and infrastructure required by AI to process the information to get valuable outcomes.

As I finish my reflection, the onslaught of new innovations continues with daily commentary by podcasters and journalists. Bottomline, the AI story continues to be written which leaves an endless road of discovery and reflection. My grandparents would be proud I found something to replace the annual book report.

About the Author

Zach Burnett is CEO of RadarFirst, where he leads the company's vision for AI governance, regulatory risk management, and intelligent compliance. He is passionate about helping organizations operationalize trusted AI through responsible governance, intelligent automation, and enterprise-scale decision-making.

With more than 20 years of enterprise software leadership experience, Zach is a recognized voice on AI governance, agentic AI, and the future of regulatory technology.

ABOUT RADARFIRST

RadarFirst is an AI-forward incident management platform for privacy and AI. Trusted by leading organizations, RadarFirst enables teams to manage incidents with speed, consistency, and defensibility by standardizing how incidents are captured, assessed, and actioned.