INDEX

INTRODUCTION

Why AI can't see beyond the next word, how a B-movie alien egg became our economic obsession, why orientation beats prediction, and what happens when the hype and anti-hype are both wrong.

PART I: BAD STORIES ABOUT AI

A HITCHHIKER'S GUIDE TO THE AI BUBBLE

Why we're spending $16 for every $1 earned, how the missile gap outsmarts the AGI race, what happens when infrastructure outlives fantasy, and why the bubble is the wrapper.

BUBBLES POP. PLATFORMS PERSIST.

How to build moats not bubbles, how Reddit rage proves demand not failure, and what calling everything 'wrappers' reveals about your thinking.

ALCHEMY 2: ELECTRIC BOOGALOO

Why Newton spent more time on transmutation than physics, how neurons laugh at maths, why the furnace builders want $7 trillion, and how to spot when brilliant people chase impossible goals.

PART II: HOW IT REALLY WORKS

WHY YOUR AI NEVER WORKS ON THE FIRST TRY

The mathematical proof your frustration is inevitable, the law that says you'll never know if you're close, how AI turns programmers into pilots and writers into navigators, and the moment Thoughtworks admitted defeat.

THE PACHINKO MACHINE PLAYS YOU

What makes pigeons peck and humans type, why Microsoft's AI CEO thinks you're going insane, how AI got inside your OODA loop, and the difference between comic book panic and actual harm.

RACIST MATHS

How AI reveals your company's hidden values, how bias can hide in three random numbers, why Grok went MechaHitler in one beat, why killing DEI is tomorrow's smoking gun, and what owls mean for your training data.

PART III: WHO WINS, WHO WORKS

THEY PAID TO PLAY COACHELLA

Why the biggest break in music costs six figures to accept, how copyright killed the thing it was meant to protect, what Morris Levy's baseball bat taught Silicon Valley, and why the Stationers' Company would love Spotify.

BIG JOBS

The apocalypse where everyone gets hired, where productivity hides for centuries, the wrong-shaped factories, why Edison funded the electric chair to win an argument, and the invention of childhood.

THE DISAPPEARING SALARY

What Malcolm McLean's metal boxes did to Detroit, what happened when France sold off taxes, why your boss would sell your future to a vending machine, what Sam Altman can't say, and how to build an empire on crumbs.

OUTRO: THE MACHINE CAN'T PLAN. YOU CAN'T STOP.

What Engelbart's 1968 demo taught us about timing, why naming shapes perception, how iteration becomes literacy, mirrors made of maths, the four lenses. And how this ends.

INTRODUCTION

AI didn't arrive like normal technology. It landed more like a B-movie alien egg in a smouldering crater. Around it gathered the usual crowd: journalists hyping miracles, critics warning of monsters, politicians staking positions, the army circling warily, gawkers staring in a mixture of awe and disdain. Some see promise, some see peril. Most don’t know what the hell to think.

Among the crowd, two figures dominate. The landowner wants to sell tickets: AGI, productivity, AI for everything. The priest warns of apocalypse: extinction risk, mass job loss, civilisation undone.

Both are wrong in different ways. Together they create an info-smog that blinds those who need clarity most.

This book won't tell you what to do about AI. There are enough consultants selling that already. These essays dig into what's actually happening - the mechanics behind the hype, the economics behind the panic, the patterns behind the noise. Think of it as orientation, not prescription.

Quick note on terminology: I use 'AI' even though it's misleading. These are machine learning systems - pattern matching, not intelligence. As the title says, nothing artificial about the capability, nothing intelligent about the computation. When I say 'AI', I mostly mean large language models since that's what everyone's mostly using. But 'AI' is what everyone calls it, and fighting the term just makes the conversation harder. One prediction: when the hype dies, the ML capabilities will keep improving. We might see an 'AI winter' followed by an 'ML summer' - same tech, honest branding.

AUTONOMOUS, EXCEPT WHEN IT MATTERS

'Next year'.

Self-driving cars have been 'next year' since 2014. Every major tech company, every car manufacturer, every ambitious startup promised the same thing. Fully autonomous vehicles just around the corner. Next year. Maybe the year after. Definitely by 2020. Or 2025. Or 2030.

The pattern is so consistent it's become a joke in the industry. But the joke reveals something about how we misunderstand AI progress.
In 2016, Uber's then-CEO Travis Kalanick declared that by 2030, most Uber rides would be in self-driving cars. Ford promised a fully autonomous vehicle by 2021. GM targeted 2019. Tesla's Elon Musk has predicted 'next year' every year since 2014. Each deadline passes, each promise fails, and new deadlines emerge like clockwork.

The technical challenge seemed straightforward: outfit a car with cameras and sensors, train a neural network on millions of miles of driving data, and let pattern recognition do the rest. If AI could beat the world's best Go player, surely it could navigate a suburban street.

But driving isn't Go. Go has fixed rules, perfect information, clear victory conditions. Driving by contrast is made up of edge cases - construction zones that change daily, emergency vehicles that need special responses, children who might dart into the street, debris that could be a plastic bag or a rock. Each edge case spawns more edge cases. This isn't just a long tail - it's a branching maze. Each edge case multiplies, revealing new dimensions of failure. Like brain chemistry, every intervention ripples unpredictably. Fix one, distort another.
Waymo, Google's self-driving subsidiary, has spent over $20 billion and logged 20 million autonomous miles. They've achieved something remarkable: a taxi service that works pretty well in a few carefully mapped neighbourhoods in Phoenix and San Francisco. The cars drive slowly, avoid highways, and still occasionally get confused by construction cones.

This is what success actually looks like. Narrow deployment in controlled environments. Gradual expansion. Human oversight at every step. And billions poured into infrastructure, all for modest operational gains. It works, just not the way anyone imagined.

The mismatch between promise and reality created a credibility crisis. When Uber's self-driving car killed a pedestrian in 2018, it wasn't just a tragedy - it was proof that the entire industry had been lying about how close they were. When Tesla's ‘Full Self-Driving’ turned out to require constant human supervision, customers felt deceived.

The technology works - just not as a drop-in replacement for human drivers. It works as advanced cruise control that makes highway driving safer. It works as parking assistance that prevents fender benders. It works as emergency braking that saves lives. Boring, incremental improvements that nobody notices because they're not the revolution we were promised.

The self-driving car saga perfectly illustrates the AI hype cycle: wildly overestimate what's possible in the short term, completely miss what's possible in the long term. The shift comes not as a dramatic replacement but as a thousand small improvements that gradually transform the entire system.

Today's AI hype follows the same pattern. AGI is always 'just around the corner', but the real transformation happens in the boring middle ground - workflow automation, code assistance, content generation. Not artificial general intelligence, but applied machine learning at scale.

This is why orientation beats prediction every time. The prediction says 'full self-driving next year' and fails. The orientation says 'watch where the technology actually works, not where we wish it would work'. One leads you in circles - the prediction returning each year like Groundhog Day. The other helps you build real systems that create real value. AI doesn't need more forecasts.

Self-driving cars didn't fail. They succeeded - just not in the form promised. Today's AI boom will follow the same path. The AGI promises will evaporate, but the infrastructure stays, enabling transformations we haven't imagined yet.

* * *

The core limitation is embarrassingly simple: AI can't see beyond the next word. These systems everyone's either worshipping or fearing - they literally cannot plan, cannot strategise, cannot see where they're going. They pick one word, then another, then another, like a driver who can only see one metre ahead. When ChatGPT writes you a business strategy, it's not strategising - it's just selecting words that statistically follow other words. The fact that this occasionally produces coherent plans isn't intelligence, it's proof that most business writing is so formulaic you can generate it without thinking at all. That discomfort is what makes many rationalise it as a thinking machine.

This book is an attempt to clear the air. It began from the same frustration you've probably felt: the conversation about AI is broken. Hype merchants and professional doomers dominate the stage. Revered figures recycle bad logic with total confidence. Industry leaders make confident but contradictory prognostications. The result is noise, not insight. The future doesn't need fortune-tellers. It needs maps. A way of seeing AI that helps us take the next step wisely instead of running from unfounded fears and chasing mirages.

Four lenses guide the analysis. **Infrastructure**: the slow, often boring shifts that eventually change everything. **Platforms**: who controls distribution and captures value. **Iteration**: the new physics of work, where retries are nearly free and breakthroughs and dead ends multiply, reshaping how we think. **Organisation**: how we adapt in response.

These lenses don't predict AI's final form - because it's embryonic. What looks like an alien egg today will be shaped iteratively, step by step, through how societies build, govern, and respond. The point is not to blindly extrapolate in straight lines, but to orient ourselves. In John Boyd's famous OODA loop, orientation is the hinge: where raw observations are turned into usable context, and where useful decisions are made. The same applies here. AI demands orientation - clearing the noise, recognising the terrain, and being able to act without wasting cycles on fantasy and fear.

The book has three parts. *Part I: Bad Stories About AI* looks at how hype and panic dominate the narrative - the bubble debates, the 'alchemy' promises, and the cult of AGI that blinds us to what's really happening. *Part II: How It Really Works* digs into the mechanics: why nothing works on the first try, why these systems behave more like pachinko machines than minds, and why 'mathematical objectivity' is anything but. *Part III: Who Wins, Who Works* explores what follows: how platforms capture value in chokepoint economics, and how jobs don't simply vanish but multiply and reorganise in unexpected ways.

Doomsday and utopia predictions? We've already seen enough of them evaporate on contact with light. These essays do something else. They're about separating signal from noise and figuring out where we actually stand.

PART I:
BAD STORIES ABOUT AI

A HITCHHIKER'S GUIDE TO THE AI BUBBLE

> *The competition for AGI-AI that surpasses humans at all cognitive tasks-is of fundamental geopolitical importance.*

That's *The Economist*, **the other week**. Not some breathless tech blogger or venture capitalist talking their book. The world's most prestigious economic publication. Notice the framing - it treats AGI as a foregone geopolitical contest.

They're not wrong about the competition. They're just wrong about what we're competing for.

Last year I did something I hadn't done in over a decade: I wrote code again. First time in 13 years. Not because I believe AGI is coming. I think it's alchemy-level nonsense. I started because I suddenly could. Because somewhere between the $560 billion in AI infrastructure spending and the endless debates about consciousness, something genuinely revolutionary happened: machine learning became boring infrastructure.

Boring is the highest compliment I can give technology. Boring means it works. Boring means you stop thinking about how and start thinking about what. Electricity is boring. TCP/IP is boring. And now, after all the hype and terror and mysticism, AI is getting boring too.

But you wouldn't know it from reading the headlines. When former prime ministers are writing op-eds about the AGI race, you know the fantasy has captured everyone - media, politicians, markets. They're so busy staring at artificial general intelligence that they're missing the actual revolution happening at ground level.

Two stories are unfolding simultaneously. One is a spectacular bubble built on geopolitical panic and sci-fi fantasies. The other is the quiet transformation of how we build everything. When the bubble narrative pops, the buildout accelerates.

WHAT'S WORKING TODAY

I'd been building systems since the 80s - architected investment fund migrations from mainframes to networked PCs in the City, built ERP for trading firms, then spent over a decade in enterprise consulting. But I hadn't written production code since 2012.

Within weeks last year I built a serverless system processing 5 million social media posts daily, tracking topic clusters and emerging narratives in real-time. Then brand monitoring dashboards. Then a 'robojournalist' that could deep-dive any trending story. Then hardware and firmware specs for a coffee machine. Then my first mobile app.

Not toy projects. Real systems. In the time it used to take to set up a development environment.

Thirteen years away from code, and within weeks I was shipping production systems in languages I'd never used. The tools had evolved that much.

Scroll through any tech community and you'll see senior developers emerging from semi-retirement like coders coming out of carbonite. People who'd graduated to PowerPoint and architecture diagrams, who barely touched an IDE in over a decade, are suddenly shipping products.

The vibe-coding community gets this. While we debate AI's impact in boardrooms, they're already building the future on Discord and shipping it to production. Yes, they're creating security nightmares and accidentally deleting production databases. Of course they are. They're inexperienced people wielding power tools. When we gave everyone electric saws, emergency rooms saw more accidents too. That's not an argument against power tools.

While established developers debate whether AI will replace them, these kids are shipping. Developers who learned their craft in the age of pull requests and sprint planning sneer at their security failures, not realising that 'best practices' are about to flip again. The barbarians aren't at the gate. They're deploying to production. And honestly? I'm a born-again barbarian myself.

And the patterns they're creating - spec-driven development, AI-first workflows - are already being productised by big tech. The innovation is flowing upward.

I'm not alone in seeing this. As Christina Wodtke, Stanford lecturer and early web pioneer, recently noted: 'The old timers who built the early web are coding with AI like it's 1995. The same people who ignored crypto and rolled their eyes at NFTs are building again. When developers who've seen every tech cycle since Gopher start acting like excited newbies, that tells you something'.

And it's not just about code. The other week GPT suggested I make jam from some obscure regional Thai plums I'd bought. Tiny things, sour as hell, no English name. I'd never made jam before. Took a photo of the carton, got the variety identified, received a recipe calibrated for their specific sourness (even spotted this meant no pectin needed), then real-time guidance that adjusted based on photos of my pot. 'Needs 3-4 more minutes', it said, looking at the bubble pattern. It was right. This capability - expertise on demand - is transforming everything from cooking to coding.

I work in Asia and see it daily: non-English speakers using AI as professional infrastructure. The language barrier just vanished - in both directions. People composing, analysing, and creating across languages at native level. The data backs this up: over 80% of ChatGPT traffic comes from outside the US, with massive usage in India, Brazil, Japan. The economic implications are staggering.

Students aren't asking 'Is this cheating?' They're asking 'How do I build with this?' They'll spend 40 years in the workforce. By the time they retire, working without AI will seem like working without electricity.

Something real is happening. Not in the research labs or board rooms where they debate ASI timelines. But in the million small moments where people discover something new they can do. The old saying *'Be realistic, demand the impossible'* was never more true.

INFRASTRUCTURE HAPPENS

In 2018, if you wanted to use ML, you hired PhDs and bought GPUs. Custom everything.

By 2020, you could rent pre-trained models from OpenAI or Google. But integration was still bespoke. Every company had different APIs, different formats, different assumptions.

Then something shifted. The models converged on common patterns: chat-style message formats, explicit system prompts, structured output modes, and a handful of consistent parameters. OpenRouter emerged to abstract away vendor differences. AWS Bedrock unified access to multiple models. Anthropic's MCP is pushing standard tool interfaces and has been adopted by everybody. What looked like competition was actually standardisation.

Watch what happened to prices once standards emerged:

- GPT-3 (2020): $60 per million tokens
- GPT-3.5 (2022): $2 per million tokens
- GPT-3.5 (2024): $0.07 per million tokens

That is the price curve you get when a capability becomes infrastructure.

The clearest sign: how new tools are built. Cursor runs their own prompt optimisation but routes to commodity LLMs for the heavy lifting. Replit does the same. They're not trying to compete with OpenAI on model training. They're building experiences on top of commodity LLMs.

This is textbook platform evolution - exactly what Simon Wardley has been **mapping for years**. Standardisation enables scale. Scale drops prices. Low prices increase adoption. Adoption creates ecosystem.

We see it today: LLM APIs for understanding. Embedding models for similarity. Vector storage for search. Components from different vendors. What cost millions to build custom in 2018 now costs hundreds to assemble. My social media analysis system is built entirely from such commodity blocks.

The AGI crowd misses this completely. They're debating consciousness while the machine learning stack is commoditising under their feet. They're worried about 'superintelligence' while developers are treating AI as just another API to call.

The real revolution isn't making machines think. It's making them boring enough that nobody has to think about them.

THE MISSILE GAP, BUT STUPIDER

We're spending $16 for every $1 earned in AI. That 16:1 investment ratio only makes sense if the winner takes everything.

Which is exactly what everyone believes.

Rishi Sunak writes op-eds about democratic values in the AGI race. The White House issues executive orders about AI safety. China announces AI supremacy targets. The EU drafts regulations for systems that don't exist. Everyone's racing for permanent technological supremacy.

This is your bubble. Not the technology - the shared delusion that someone's about to achieve irreversible computational dominance.

The panic has a patient zero. When Geoffrey Hinton quit Google to warn about AI risk, he didn't just change jobs. He transformed a technology discussion into an existential race. Suddenly every major power faced a terrifying question: What if our enemies get AGI first?

Sam Altman knew exactly which buttons to push: congressional testimony about the need for regulation (from the company furthest ahead), warnings about AI risk, and a playbook of building in public while presenting OpenAI as the responsible actor that just needs resources to 'do it safely'.

It worked. The Stargate announcement - $500 billion for AI infrastructure - is the logical endpoint of this narrative. When you believe you're racing for permanent species-level advantage, no amount is too much. The 'long-termists' have everyone convinced we're at the hinge of history.

The curious thing about existential arms races? They're incredibly profitable for arms dealers.
Jensen Huang needs governments to panic-buy GPUs. Sam Altman needs infinite capital for compute. Microsoft, Google, and Amazon need regulatory moats only they can afford. Every warning about AGI danger is also a pitch deck.

During the Cold War, the US and Soviets would leak reports about UFOs and mind control programs. Deliberate misdirection to waste enemy resources. The AGI race has the same dynamics - except this time, everyone's falling for their own propaganda.

At least the missile gap was about real missiles.

EVERYONE'S LOOKING UP

The most interesting part of **Ed Zitron's recent 14,000-word AI takedown** isn't what he gets wrong. It's how he gets it wrong.

He spends thousands of words debunking AGI hype, then judges every AI product by AGI standards. He dismisses 'agents' because they're not fully autonomous. He mocks chatbots for not being conscious. He's so busy fighting the fantasy that he misses the reality.

He's not alone. The entire discourse has been captured by AGI framing. Critics and believers alike judge current AI by science fiction standards. It's like dismissing cars because they don't fly.

This is what Baldur Bjarnason called the '**LLMentalist effect**' (great article!) - we've projected consciousness onto pattern matching. The chatbots seem so human that we can't help but evaluate them as minds rather than tools. Even skeptics fall into the trap, spending more time debating whether they're 'truly' intelligent than asking whether they're useful.

Real revolutions happen gradually, then suddenly. In 1996, if you asked for proof the internet would change everything, what could anyone show you? Amazon selling books? Email replacing faxes? The transformative applications hadn't been invented yet because the infrastructure didn't exist.

We're in the same moment now. People demand to see the AGI-level breakthrough while missing the million small transformations already happening. My social media analysis system would have been impossible five years ago. Not impractical - impossible. The components didn't exist at any price.

Every week, developers uncover new patterns. Natural language becomes the interface for everything. Retrieval makes search contextual instead of literal. Multi-step reasoning chains that once collapsed now hold together. Not consciousness, but capability after capability that wasn't there before.

Meanwhile the hype cycle and the hand-wringing predictably miss the point.

The skeptics and believers are having the wrong argument - two sides of the same shitcoin. The C-suite fence-sitters striking a 'balanced perspective' are no better, hedging between factions that both lost the plot, serving up a smorgasbord of bad takes. The question isn't whether we'll create AGI. It's whether we'll notice that we don't need to.

WHEN THE MUSIC STOPS

The AGI bubble will pop. Not because the technology fails, but because the fantasy can't survive contact with reality.

The trigger could be anything. An AI company admitting AGI is decades away. A government realising it stockpiled GPUs for nothing. Or investors noticing that $560 billion for $35 billion in revenue isn't a business model so much as a cargo cult.

When it happens, the narrative collapse will be spectacular. All those breathless headlines about consciousness and superintelligence will age like dot-com era predictions about the 'new economy' where profits didn't matter. The Stargate project will become this generation's Webvan - ambitious, well-funded, and built on false premises.

But the doomsayers miss something crucial: the infrastructure remains. After the dot-com crash, we still had fibre optic cables, data centres, and trained engineers. The speculation died. The internet didn't.

Same pattern here. When the AGI fantasy evaporates, we'll still have:

- Models that can read, write, and analyse
- APIs that cost pennies to call
- A generation of developers who know how to build with them
- Actual products solving actual problems

The companies that survive won't be the ones promising AGI. They'll be the ones who understood early that machine learning is just really useful when available as infrastructure. Like the difference between Pets.com and Amazon - one promised to change the world, the other was building warehouses.

Medieval alchemists never turned lead into gold. But while chasing that impossible dream, they invented chemistry. They failed at transmutation but succeeded at something more valuable: understanding how the world actually works.

Same story, new century. The AGI labs won't crack consciousness. But chasing the ghost in the machine, they've built infrastructure that changes everything. It turns out the bubble was the wrapper all along.

CONCLUSION: NOW WHAT?

So what do you do with this knowledge?

If you're a developer: build. The tools are here, they're cheap, and they're getting better every week. While everyone else debates consciousness, ship products. The barbarians aren't knocking - they're already through the door with mud on their boots.

If you're a business: ignore the noise. Everyone has strong opinions about AI, and they're mostly wrong. While experts argue and vendors overpromise, focus on what works today. That boring automation, that small efficiency gain, that better interface - these compound. The companies that win won't be waiting for clarity. They'll be the ones who started with simple tools and learned by doing.

If you're an investor: you understand the moat dynamics - infrastructure players need scale, builders need distribution, data, or workflow lock-in. Want to make money on AI? Bet on the boring stuff. The companies making the tools everyone else relies on. The ones using AI to shave 3% off shipping costs. The businesses that would thrive even if we proved tomorrow that consciousness is mathematically impossible. They're building for the world where AI is infrastructure, not magic.

If you're a government: yes, sovereignty matters. You need domestic compute and models you control. But the race isn't for AGI - it's for practical ML capability. The question isn't whether to invest in infrastructure, but how much is enough. With open models improving and inference costs plummeting, the barriers are lower than the panic suggests. Build what you need, not what the arms race demands.

The hardest part isn't understanding the technology. It's seeing past the narrative.

When the bubble pops, the pundits will act shocked. How did we spend $560 billion chasing digital consciousness? How did The Economist fall for it? How did governments stockpile GPUs for a race that couldn't be won?
But by then it won't matter. The builders will have inherited the infrastructure. The vibe-coders will be running production. Your competitors will be shipping features you thought impossible. And everyone will pretend they knew all along that the real revolution was never AGI.

It was making intelligence so boring that nobody thinks twice about using it. Just like electricity. Just like the internet. Just like every transformation that actually mattered.

NOT ARTIFICIAL, NOT INTELLIGENT

WHAT AI COMPANIES DON'T WANT YOU TO KNOW

Django Beatty has spent four decades watching technology reshape business and culture.

He started in the City of London during the Big Bang, as spreadsheets replaced paper and mainframes gave way to PCs. During the dotcom era, he ran an independent record label and built large-scale media sites. Later, at Capgemini, he led enterprise web platforms before founding Fluxus, a boutique consultancy focused on AWS and AI.

He's lived through every cycle of hype and disappointment - from the internet's first broadband leap to today's AI - and learned to separate noise from signal.

You've been told AI will either save humanity or destroy it. You've been told your job is disappearing. You've been told we're racing toward artificial consciousness.

It's all wrong.

In this clear-eyed investigation of what AI actually is - pattern matching, not intelligence - Django cuts through both hype and panic to reveal what's really happening: a platform shift following the same predictable patterns as electricity, the internet, and every transformative technology before it. While pundits debate consciousness and executives chase productivity metrics that won't appear for decades, the real transformation is already underway.

From why ChatGPT sessions become addictive iterations to why the 'AI job apocalypse' keeps getting postponed, from the truth about copyright battles to what AI companies know but won't admit, this book provides the orientation everyone needs to navigate a genuinely transformative but wildly misunderstood technology.

Simon Wardley – Advisor and Speaker
"'The AGI bubble will pop. Not because the technology fails, but because the fantasy can't survive contact with reality... The real revolution isn't making machines think. It's making them boring enough that nobody has to think about them.' - Well said"

Steve Yegge – Engineering Leader
"I like this take. Cognitive exoskeleton pilot sounds about right to me."

Dries Buytaert – Founder of Drupal, Co-founder of Acquia
"Well-written... it made me think. I loved that about this essay."

Michael Nelson – Senior Fellow, Carnegie Endowment
"Useful."