Hannah Grey

CULTURAL VIBRATIONS SPRING 2026

Collective Coefficient

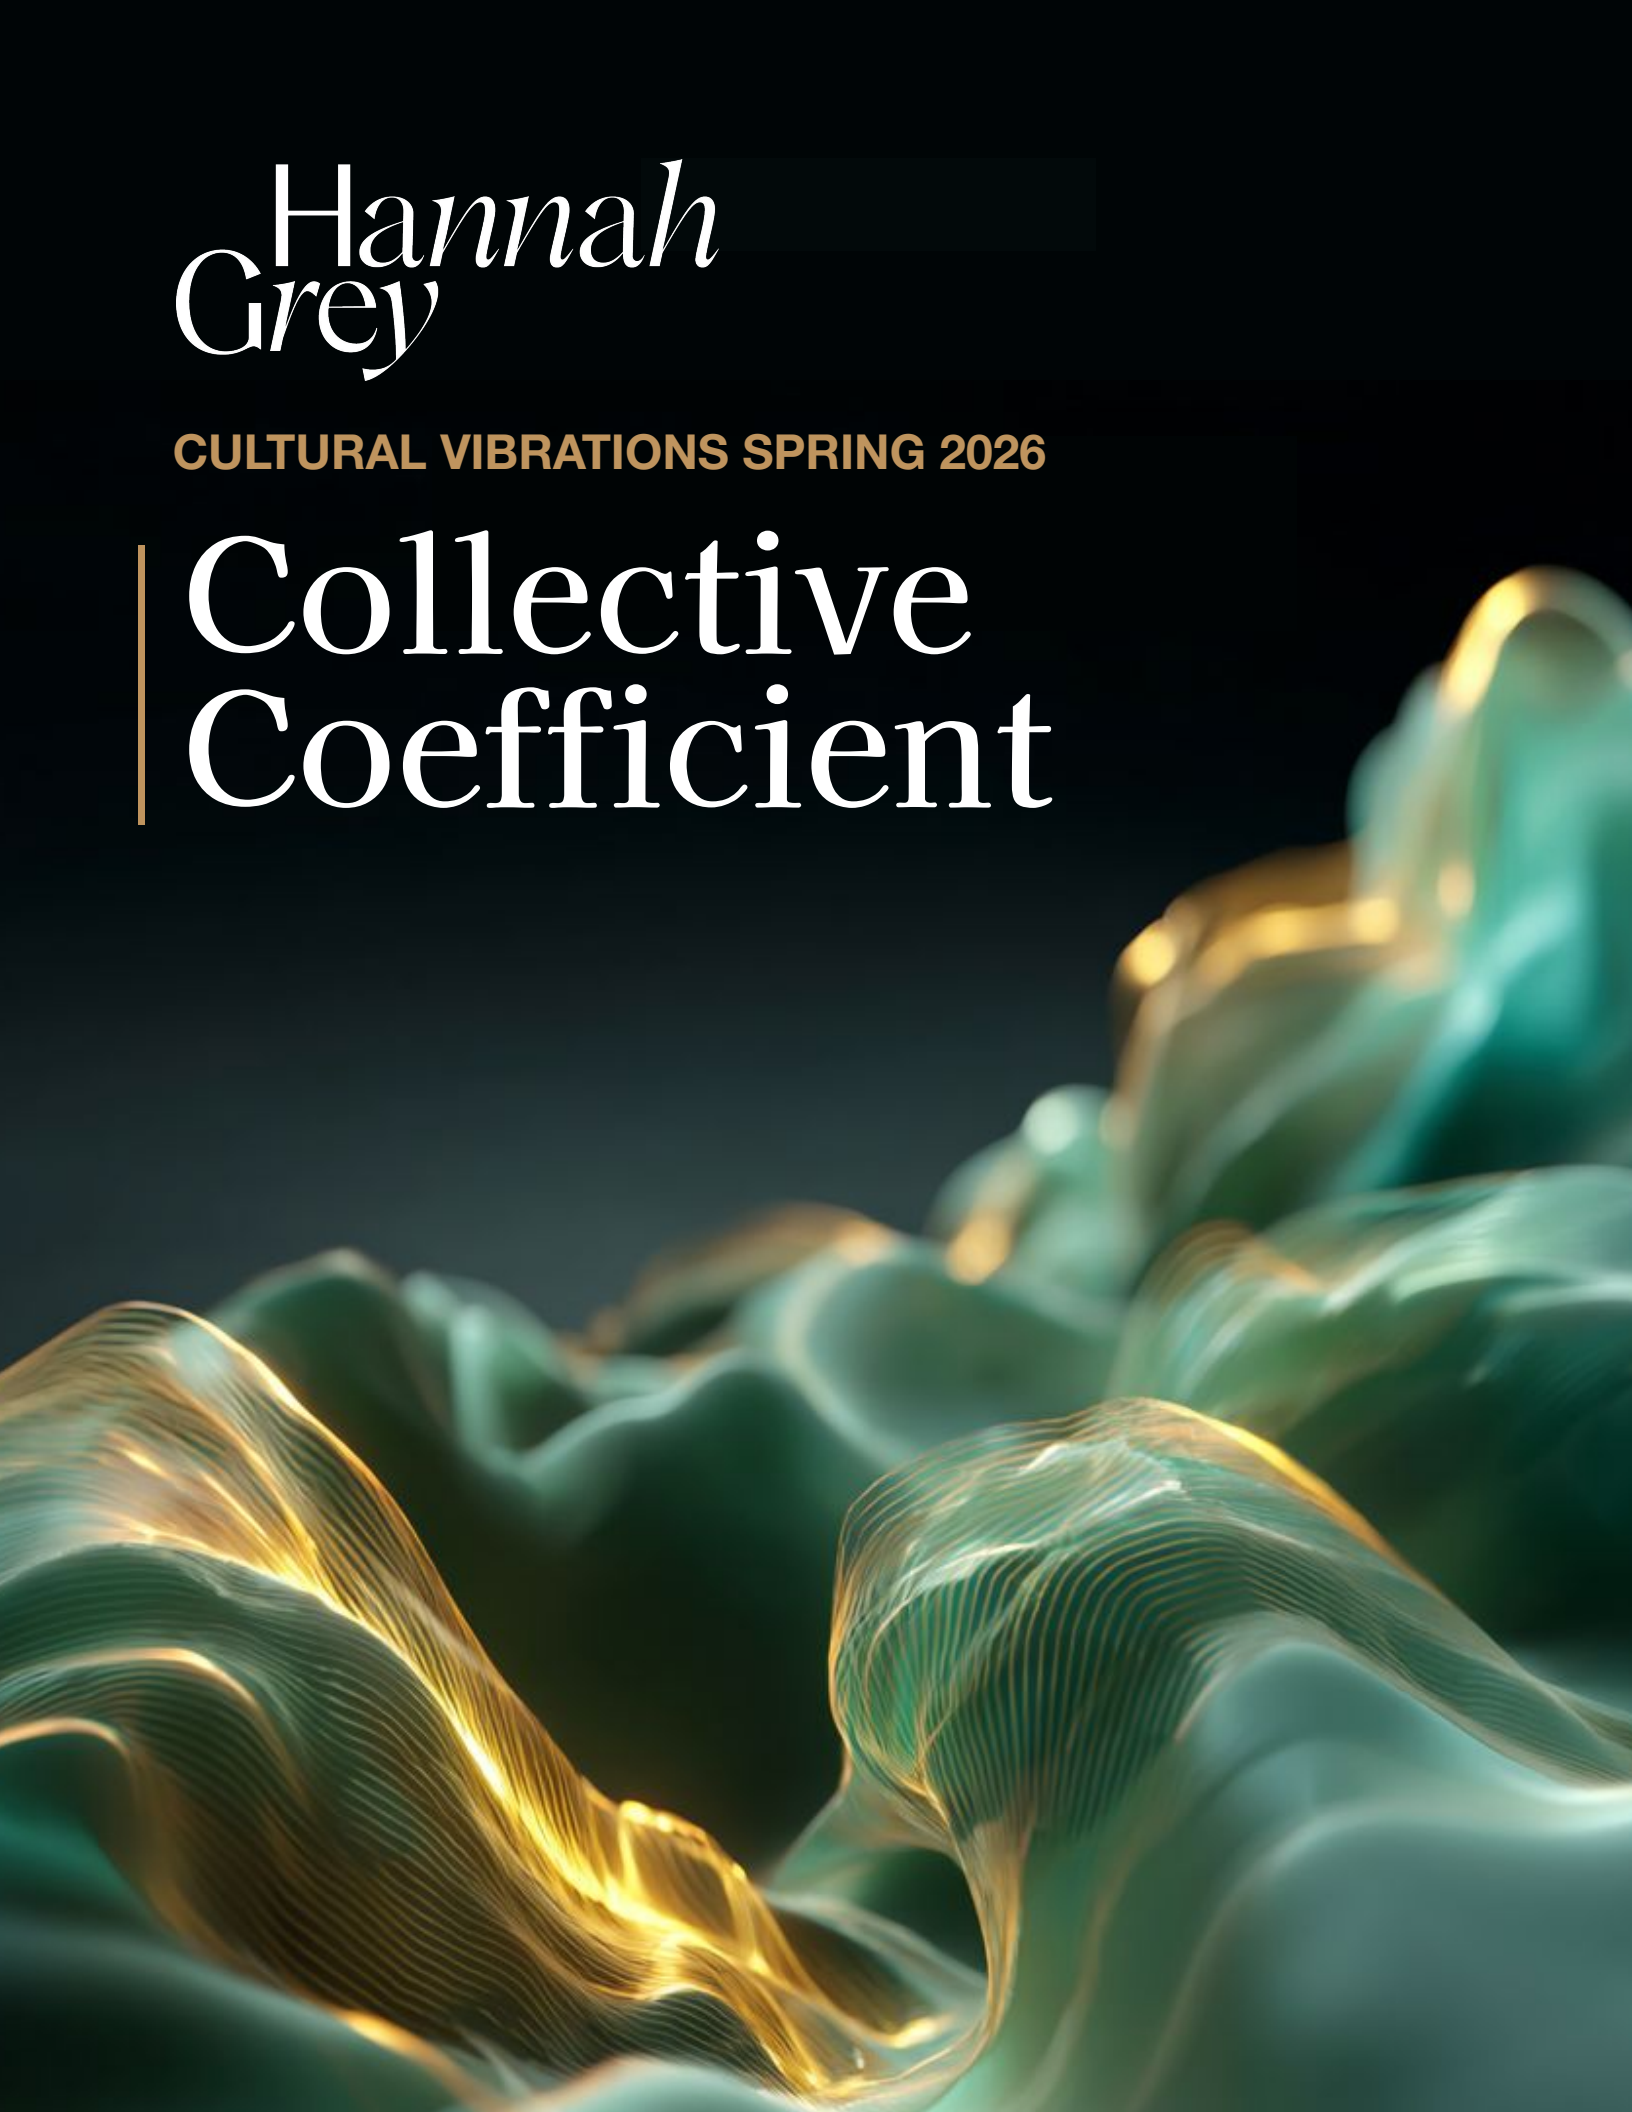

Welcome to the Spring Edition of Cultural Vibrations, anchored in the Collective Coefficient.

This Journal isn't a forecast of a robotic future or a lament for a human past. It's an exploration of the symbiotic world where people and machines operate as partners, personally and professionally. As artificial intelligence moves from a tool we command to a collaborator we engage, the very grammar of work, creativity, and economics is being rewritten. We are no longer just users; we are becoming orchestrators.

For a generation, our relationship with technology was defined by command and control. We built systems to execute tasks, process data, and optimize for efficiency. Now, a more profound vibration is emerging: a collective **shift from human-centered *design*** to **human-machine *teaming***. In a world where intelligence is becoming ambient and accessible, the focus is moving from the power of the algorithm to the synergy of the partnership.

This edition documents that dynamic. The Collective Coefficient holds a creative tension we explore throughout these pages: the economic friction of the **"Great Compute Reckoning,"** the architectural demands of **"The Unbundling of The Firm,"** and the deep human need to discern truth as **"Reality Somms"** in a synthetic world. We trace this dynamic through the raw data of our **Vibrations Index**, which reveals a culture grappling with AI's role - embracing it as an assistant but drawing firm lines at replacing human judgment and empathy.

In an age of intelligent systems, our primary work is no longer execution, but collaboration.

As a venture capital firm, we see this not as a race for artificial general intelligence, but as **the dawn of the collaborative era.** The most durable opportunities will not come from building bigger models, but from creating the connective tissue - the infrastructure for trust, translation, and compounding human advantage. The challenge ahead is not merely to build intelligent machines, but to architect a world where our collaboration with them elevates the human expression, where our collective is greater than our parts.

We hope this journal offers a compass for those seeking to create, invest, and lead with intention, grounded in the belief that intelligence is not a resource to be mined, but a relationship to be cultivated.

— Kate, Jessica & The Hannah Grey VC Team

The Vibrations Index

Research in Partnership with Okay Human

The Vibrations Index tracks cultural readiness for emerging technology through proprietary research. This year, to gain a deep understanding of the combined human & AI quotient, we interviewed 100+ US-based adults about their usage of, comfort level with, and emotional readiness for AI in their daily personal and professional lives.

HOW DO YOU FEEL ABOUT AI

23.7%	40.7%	33.9%	1.7%
Very positive	*More positive than negative*	*Right in the middle*	*More negative than positive*

AI USAGE FREQUENCY

Multiple times a day	**45.8%**
About once a day	**14.4%**
Two or three times a week	**27.1%**
Less than once a week	**12.7%**

60% of the sample used AI at least once a day

HOW INFORMED ABOUT AI ARE YOU

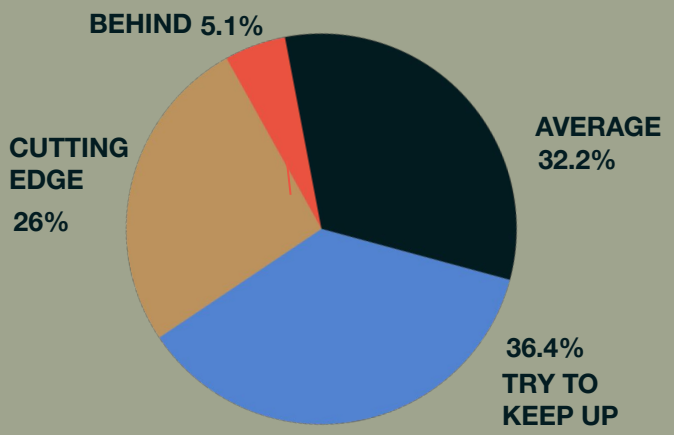

Our findings are unsurprisingly layered - representing **the contradictions of the human condition.** While the headlines surface wariness and fear, our panelists find AI overwhelmingly exciting. When trusting AI, **respondents see nuance much as they do with human relationships** - agents make mistakes, information is taken not as fact, but at face value.

When asked where AI does not belong in their lives present and future, **respondents had clear boundaries**, not by function area but by role - AI can be a part of a healthcare experience, but does not yet replace a human physician, it can be your relationship coach, but not romantic partner. As our relationship with AI continues to crystalize, **we must remember the human layers and approach this new era with the curiosity** to see what's possible through collaboration.

Key Takeaways.

01.

From the truest AI believers to the most conflicted, regardless of usage frequency, there's little doubt that the tech is genuinely transformative.

We saw widespread agreement that AI has the **potential to transform people's lives in major ways.** For the most part, this was viewed positively, while a minority focused on the effects AI might have on people's jobs.

"My step son is autistic and very introverted. But since he's started using Chat GPT to talk about his creative ideas he's now writing his own fiction stories. AI assists him by categorizing his ideas and putting them into words···his confidence has skyrocketed."

— Chelsea, Rural MO

"I needed to get my mom something to put on the wall of her new nursing home - she has dementia - so I asked Gemini. It suggested a framed picture of her in her beauty shop cutting hair - brilliant!"

— Melanie, Urban KY

And in people's real lives, the majority of AI impacts are positive - even healing, delightful, or intimate.

02.

Moving from the abstract to the concrete, the real-world stories people told to illustrate their overall sentiments about AI tended to be positive, sometimes quite strongly so. Their stories captured how **AI is already improving people's lives in deeply emotive ways**, meeting unmet needs from the social to the intellectual to purely fun - sparking creativity, serving as a thought partner, a mentor, a life coach.

03.

Trust is a primary concern with AI, reflecting the intimacy of potential use cases and need for privacy.

Before we asked about trust directly, it was already being brought up frequently, and in different contexts. Data privacy was a big one, with many mentioning **concerns over how their personal info was being used & protected**. These concerns impacted not only overall sentiments about AI, but what applications & roles people were willing to use AI for, for example using tools for research was fine, but giving AI information that isn't already publicly available about you, was not.

"I do know and understand that AI isn't 100% accurate, but neither are humans."

— Abel, Urban TX

"AI is more accurate than people give it props for, and I think that they should support it because it's just like humans, it has flaws."

— Christopher, Urban GA

AI's outputs should be assumed unreliable until tested, but this doesn't necessarily equal untrustworthiness.

04.

Respondents indicated the only way to develop trust in an AI tool's outputs was by testing and fact-checking, **building up trust overtime**. It was **expected that AI outputs are somewhat unreliable and inconsistent** - and while perhaps disappointing, it's also in line with how we judge human contributions.

The potential for AI hallucinations and errors is so well-known that, for many, this element didn't even equate to a trust issue - it's just how the technology works, and the adept user knows to double-check.

05.

Adopting new AI tools is highly dependent on trial and error, with online reviews playing a limited role.

When describing how they decided which new AI tools to adopt, participants placed a **great deal of emphasis on trying the tools for themselves**, vetting their outputs through their usual processes of double fact-checking in order to gauge if a certain tool would be a good addition to their repertoire.

Word-of-mouth recommendations played a role, as did existing brand equity, but online reviews played a relatively smaller role than we'd expect when it comes to new product discovery. This suggests the highly idiosyncratic uses for which people engage AI tools, and a baseline expectation of similarity across tools of this type.

At first glance, it seemed like AI is unwelcome in areas of life that are sensitive and soulful.

06.

When we asked respondents where in life AI **didn't belong**, we saw two main themes: privacy and the area of life itself. **Data privacy was a major obstacle**, leaving many unwilling to involve AI in areas of life where this was paramount or where sensitive topics were discussed - such as finances and health - areas where their information wasn't already public.
Some topics felt uniquely human, religion and romance, leading respondents to feel better served by a person than an agent or a copilot.

07.

But when pushed, it seemed more about the roles & reliance of AI, rather than being off-limits in areas.

As a follow-up to the "*Where is AI unwelcome*?" question, we asked what it would take for AI to be welcome in these spaces after all. Interestingly, many respondents were able to imagine where AI might be welcome, despite their previous hesitations. On the one hand this meant guaranteeing privacy & data security, but even in the "soulful" category there were places AI could be involved. The key preference seemed to be a consensus that AI should play a supportive role rather than a central one, aiding humans in their tasks rather than doing the tasks for them. **So AI may be welcome as a relationship coach or therapist, but not a romantic partner, or it might help a pastor prepare a religious sermon but not deliver it.**

08.

Emotionally vulnerable applications of AI were often the ones respondents felt others would find most surprising.

Participants frequently pointed to ways they used **AI that intersected with some ostensibly off-limits topics** of life such as romance, as well as emotionally intimate areas such as friendships or mental health counseling. The perception of these use cases may be seen as controversial among their social circles, something that they may not advertise due to a potential stigma. Surprising professional uses of AI tended to focus on automated or boosted efficiency rather than relational or social applications.

"If AI became a real-life person and walked into the room, they would be like a business person, dressed in a suit, very professional and down to earth and would know me best."

— Annie, Urban IL

"AI right now is okay for coding and stuff like that, but it can never replace humans. It can never replace human connections, so don't try and make it."

— Russell, Urban IA

09.

Professionally, most imagine AI serving as an administrator, assistant, and advisor.

When we asked participants to imagine what roles AI would play in the future of their professional lives, they **painted a picture that encompassed the types of things AI currently does, just better and at a larger scale**. Only a small number imagined AI tools replacing person-to-person or relational professional tasks, at least in the next few years - though we also saw fear about AI's long-term impacts on human wellbeing, specifically through replacing people at work.

10.

In the personal realm, AI is imagined in virtually all task-focused roles, with increasingly relational responsibilities as well.

Over the next few years, respondents imagined **AI playing numerous, varied, and involved roles in the personal lives of average people** - teacher, trainer, trip planner, personal assistant, financial advisor. The more emotional, relational levels already present in people's current uses of these tools are predicted to become more prominent in the future as well.

AI is welcome in the world of health & wellness, but most draw a firm line at health care provider responsibilities.

Our respondents readily pointed to various roles that AI could play when it came to health & wellness - dietician, nutritionist, and fitness trainer chief among them, as well as a more general role keeping them accountable and consistent. There was even openness to using AI as a way to assess symptoms in a preliminary way. However, there was a firm line in the dataset - **very few respondents were willing to place AI in the role of a healthcare provider** like a doctor or nurse.

11.

12.

With high-stakes issues, accuracy is more than a percentage change; it’s about communication, perceived expertise, and transparency.

When we asked participants to choose between a 99% accurate AI doctor and a 90% accurate human doctor when dealing with a serious diagnosis, the vast majority of respondents opted for the human doctor. Not only did they j**ustify this choice based on wanting their questions answered**, but **some also insisted that the human doctors really would be more accurate**, despite us telling them otherwise.

13.

For many, AI can be envisioned as an aspirational version of themselves, a dear friend, or even an ideal partner.

In a personification exercise, the intimacy of our participants' relationships with their AIs was reflected in how they imagined the technology. While some simply literalized the role of teacher, therapist, or servant, others picked much closer archetypes like a dear friend or - through focusing on the personified tech's attractive features - a potential romantic partner. There was a trend of **describing AI as oneself idealized,** underlining just how near and dear to one's core identity their AI becomes over time.

> *"One thing I would like for leaders of advanced artificial intelligence communities and companies to realize is that most of us speak in layman's terms. Most of us need you to understand that we're gonna progress more slowly than you do mentally. And it's important to remember."*
>
> — *Brian, Rural MS*

The AI future should look like transparent, highly-human communication, relevant to real people's needs.

14.

Our participants are already convinced of the transformative potential of AI tools - what they want to see as they look into the future is the same communication & transparency that set human doctors above their more accurate AI competition. They value honesty over pandering, identifying mistakes and correcting them. They want **AI to be leveraged for the needs of the average person**, tools that are indispensable solutions to mundane problems - **must-have fixes rather than nice-to-have curiosities.**

The Hardest Part of Building Right Now Isn't Building

Scott Dadich
Venture Partner, Hannah Grey
Former Editor-in-Chief, WIRED

Most founding stories are missing the most important scene. Not the product launch, not the funding announcement, not the origin anecdote polished for the pitch deck. The scene that's missing is the one where the value in the founder's world actually *turned*, where something they believed, or feared, or refused to accept became the crack that everything else grew from. Without that scene, you don't have a story. You have a timeline with a logo on it.

I spent more than a decade at WIRED chasing those scenes. The founders who stayed with you were never the ones with the cleanest narratives. They were the ones willing to name what they were up against. The thing that was broken. The version of reality they found intolerable. The specific fight they picked, and what it cost them to pick it.

That instinct became the operational premise of Godfrey Dadich Partners, the strategic brand studio I co-founded. We applied the tools of journalism to how organizations tell their stories. And the first thing we learned, every time, was that the story was in the adversity the client had stopped talking about.

The obstacle isn't the enemy of the story. *It is the story.*

Consider what Microsoft was facing in early 2020. The pandemic obliterated every assumption about where work happened, how work happened, and what work was for. Working in partnership, we created WorkLab as a genuine editorial platform: research, long-form reporting, a podcast, a newsletter. The adversity wasn't the problem to be solved before the story could begin. It *was* the story. Five years of that platform positioned Microsoft as the defining voice on AI in the workplace. The manifesto we helped craft framed AI not as automation but as copilot. That idea was so precisely right, so deeply earned, that it became the name of the product itself.

But the pandemic was a recognizable crisis with a clear before and after. The adversity facing founders right now is stranger, more diffuse, and harder to name.

Last cycle, the fight was external and legible. *Can we build it? Can we scale it before they do?* The adversary had a name and an address. The risk was execution, and the story of overcoming it had a familiar shape.

The adversity of this cycle is different in kind. It's not a war over territory. It's a matter of conviction. The tools have gotten so powerful that the hard part is no longer building the thing. The hard part is believing the thing *matters*. The most well-capitalized companies in history can render months of work obsolete with a product update. The market is so loud and moving *so* fast that standing still long enough to find your story feels like a luxury. The weight isn't logistical. It's psychological. An ambient dread that what you're building might not survive contact with next quarter.

That refusal to flinch, the daily act of conviction in the face of radical uncertainty, is the adversity most founders in this cycle haven't found the words for yet. And it's the most important story they could tell.

My own scene happened in a bagel shop when I was nineteen. I had walked away from a full-ride engineering scholarship at the University of Texas. I had been suffering from anorexia for years, the accumulation of a childhood defined by abuse and the need to control something in a life that felt entirely out of my hands. One morning I collapsed on a path to the cafeteria. My heart rate had dropped into the low twenties. I had nothing figured out.

I moved back to Lubbock. I enrolled at Texas Tech and started going to therapy. I got a job at a bagel shop, the first one in the city, showing up at 3:00 a.m. to make the dough, bake the bagels, and open the place for the morning rush. One day the owner mentioned she'd hired a sign painter to redo the menu boards. I offered to do it for half the price—just to make some extra cash. I stayed up all night with pastels and fluorescent markers. The next morning, a regular named Sonia came in, asked who'd done the new boards, and handed me a business card. She was an art director at an ad agency. I didn't know what an art director was. I didn't even know that design was a job. What I did know was the feeling I'd had all night. The same obsessive need for control that had been slowly killing me suddenly had somewhere to go. Channeled into color and composition and letterforms, it wasn't a pathology anymore. It was a practice.

Not a triumph. A fracture becoming a door. I tell that story because the thing that nearly killed me turned out to be the same thing that made me good at my work.

The obsession. The need for control, The refusal to accept the world as I found it. It was a liability until it became a practice.

That is the exact exchange happening inside the best founders right now. The nagging fear, the constant suspicion that what you're building doesn't matter, is the same force that sharpens the people who push through it. The ones who survive this cycle won't just endure that feeling. They'll name it. They'll turn the private, inconvenient reason they can't let go of their problem into the story they tell the market. Not the pitch. The thing underneath the pitch.

The legendary screenwriting teacher Robert McKee, whose students have won more than sixty Academy Awards, puts it plainly: *nothing moves forward in a story except through conflict.* The companies that understand this encode the founding act into everything: how they hire, how they communicate, how they decide when conditions get hard. It becomes character. The ones that don't are forever reintroducing themselves, wondering why the market can't quite feel what they're about.

Every founder has a scene that turned. The question is whether they've found it, and whether they're willing to tell it.

In this cycle, that willingness isn't just good storytelling. It's survival.

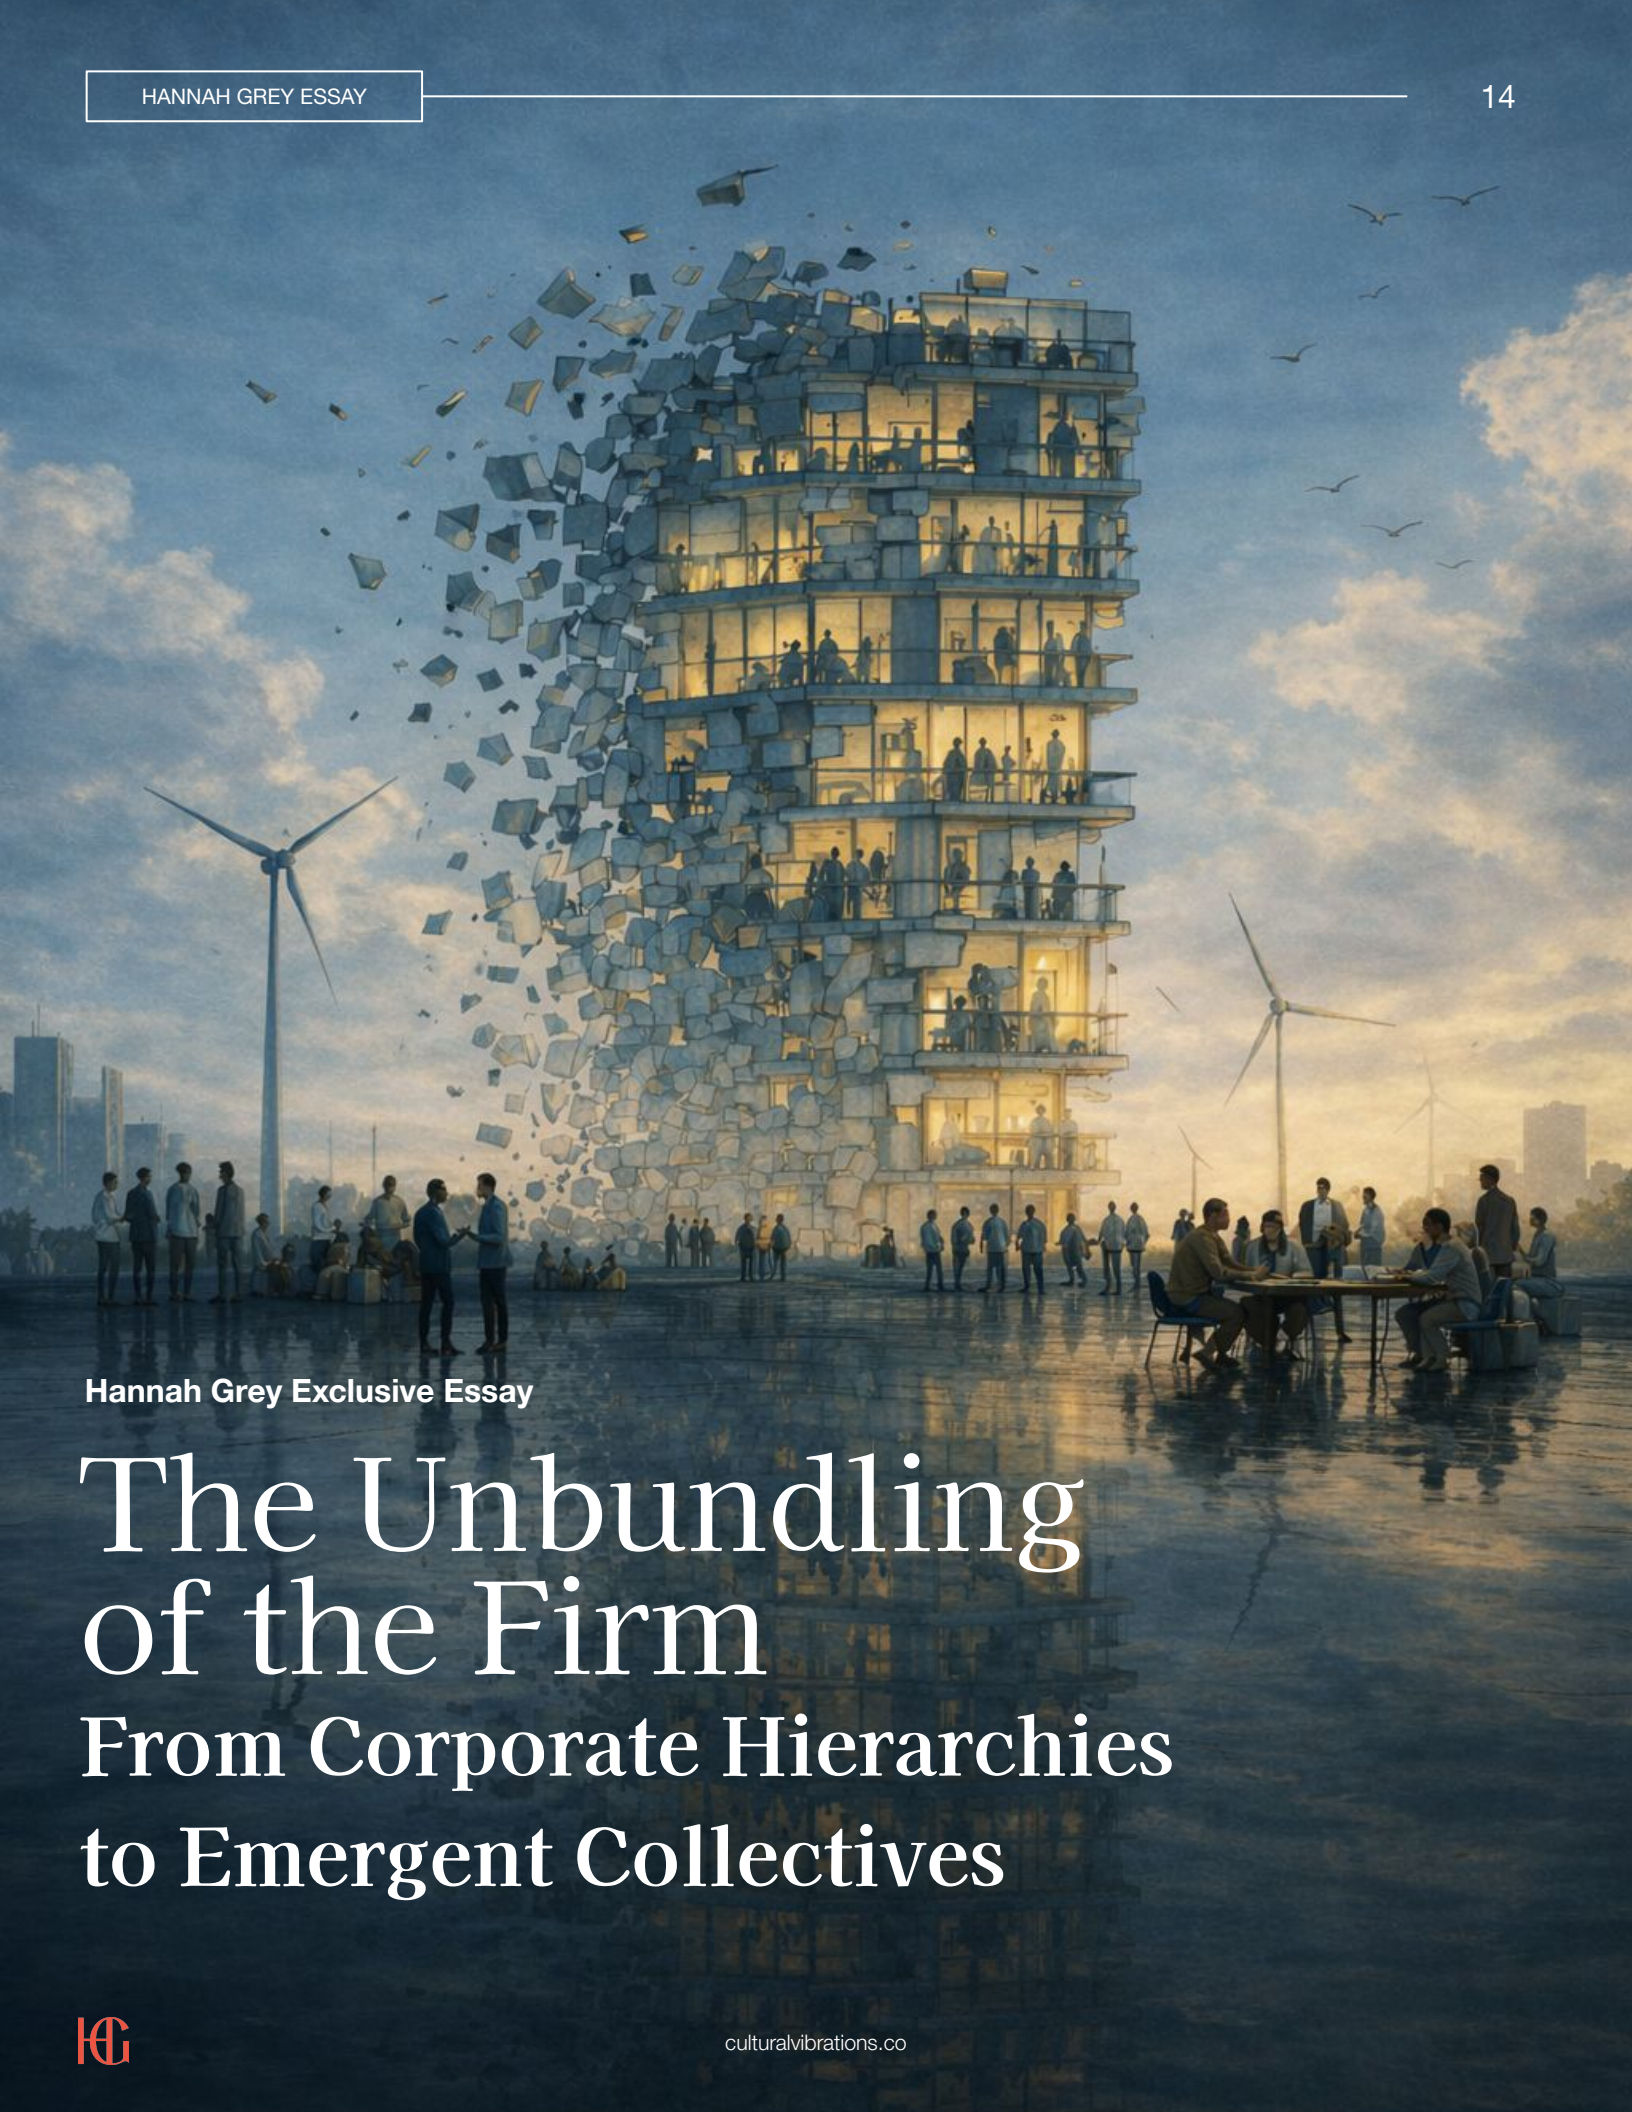

Hannah Grey Exclusive Essay

The Unbundling of the Firm

From Corporate Hierarchies to Emergent Collectives

Corporate hierarchy has been the default grammar of business. The pyramid structure of the 20th century was built for a world of information scarcity and high coordination costs. Hierarchies existed because the cost of leveraging the open talent market for every task was prohibitive: it was cheaper to bundle workers as employees rather than constantly find, contract, and manage them as freelancers. That assumption, and the physics of company building is now collapsing. As leverage with artificial intelligence accelerates and frictionless access to global talent becomes the norm, the traditional corporate structure is becoming irrelevant.

This change isn’t leading to chaos or fully autonomous, human-less entities. Instead, a far more interesting architecture is emerging in its place: fluid, dynamic "collectives" that we call Emergent Organizations. The simplistic narrative of "Solo Founder Unicorns" dominating the headlines is missing the plot. The real opportunity isn’t about replacing structure with a single founder and an army of agents. It’s about building a new type of company from the ground up, designed for an AI native world of instant trust and infinite leverage.

We believe Emergent Organizations will be defined by two key characteristics, with a new cohort of ‘picks and shovels’ infrastructure enabling them:

01. A Small, High-Leverage Core Team

At the center of every Emergent Organization is a small, deeply strategic human team. Their primary job isn't to manage people; it's to define problems and orchestrate resources. They are less like managers and more like conductors of a hybrid workforce, identifying the task to be solved and determining whether the work is best handled by automated agents, internal employees, fractional specialists, or a combination of all three.

02. A Liquid, On-Demand Periphery

Surrounding this core is a vast, fluid network of resources that can be summoned or dismissed instantly. This is a hybrid ecosystem of specialized AI agents, open-source models, and expert human freelancers, all tapped on a task-by-task basis. This liquid periphery allows the organization to scale its capabilities massively without the bloat and inertia of traditional hiring.

This structure will create a new batch of corporate frictions. When the company is a fluid collective of humans and agents, how do you maintain a coherent culture? The corporate security perimeter, once a digital fortress, dissolves into an ever-changing network of human and machine collaborators. As critical tasks are delegated to agents, a new web of accountability emerges: when an orchestrated agent fails, where does the liability fall? These questions aren't arguments against the future, they're the important questions to build around. They define the next frontier of opportunities in algorithmic assurance, distributed culture-building, and managing the delicate trust layer between humans and intelligent machines.

> *"It's important to think about AI and humans collaborating, utilizing the huge benefits of AI's ability to find answers to questions so quickly, having sort of a united approach. I think we can learn from each other.*
>
> *— Meagan, Suburban MN*

The logical conclusion of this trajectory is a future of Algorithmic Orchestration. Here, the very nature of human work is redefined. As AI takes over coordination and execution, human value shifts decisively from management to meta-management. We will no longer be supervisors in a chain of command, but orchestrators of complex, intelligent systems. Our core tasks will be to design the goals, define the ethical guardrails, and curate the emergent intelligence of a deeply integrated human-machine collective. This will build new business models on a foundation of strategic thought rather than operational friction.

The hierarchical firm was a solution to a problem that no longer exists. The future belongs to founders who understand that the scaffolding of the modern company is being rebuilt from first principles. AI co-workers aren't the end game, they're a catalyst to build a new kind of company, one that's more adaptive, intelligent, and built to continuously learn.

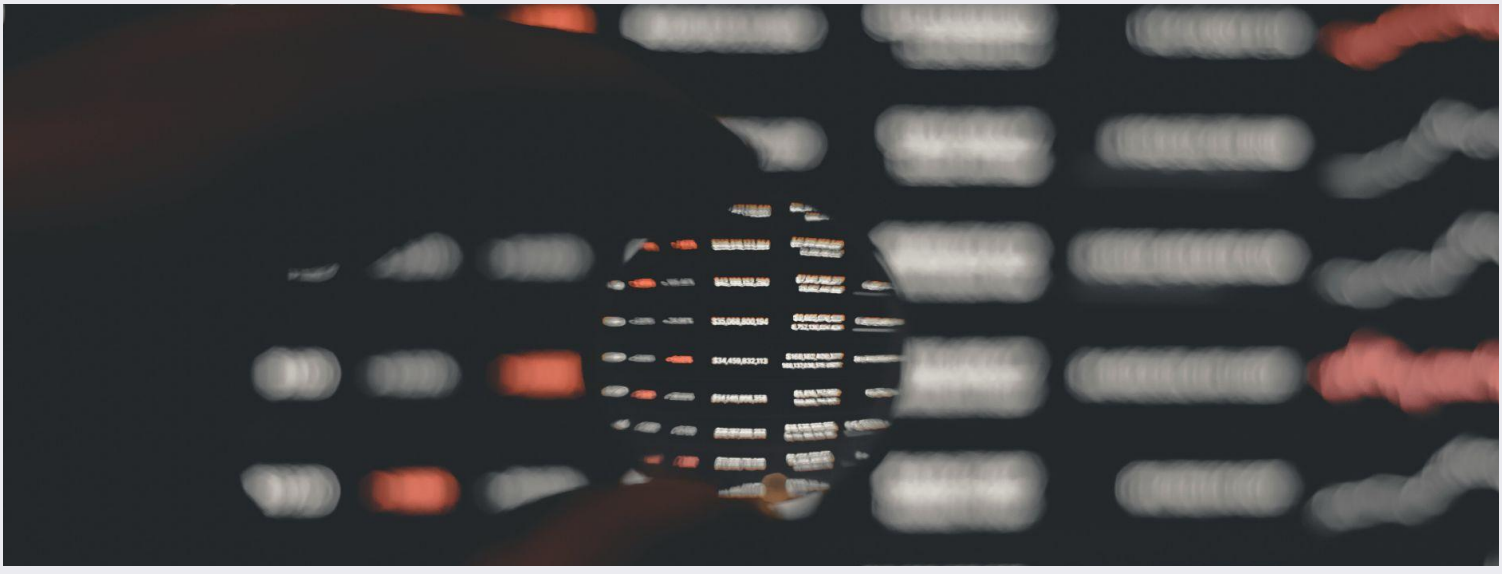

From Extraction to Sustainability: The Human-Centric Era

Kevin Walker
Former Head of Consumer & Cultural Insights at CAA

A quiet but powerful shift is underway in the global economy. The long-dominant model of extraction, a relentless pursuit of shareholder value often at the expense of quality and consumer well-being, is showing signs of exhaustion. Consumers are experiencing this firsthand in everything from airline travel to the quality of their clothes, a phenomenon some have called "enshittification." This isn't a moral failing; it's a market signal.

The extractive playbook is reaching its limits, and a new, more sustainable paradigm is emerging: the human-centric economy.

AI has colonized our attention span and the cultural discourse that once belonged to human curiosity, human craft, and creative evolution - the anti-human economy is taking shape. Yet, as we barrel towards AGI (Artificial General Intelligence), we face the existential threat of finite natural resources (water and land) and the continued affordability of computing memory and infrastructure.

But what if we challenge ourselves to think beyond AI, to a world where advanced tech and human sustainability work together, inextricably. While engulfed in the pain of the era's nascency today, society will quietly arrive at a new way of being, of profit making, of sustaining, and of human centricity.

The Human-Centric Economy

Representing a fundamental shift from a "winner-take-all" model to one focused on the building and restoration of society and communities, the Human-Centric Economy empowers the dispersion of opportunity through investment in the people themselves and their core needs.

Community Wealth Building

The Human-Centric Economy begins with the idea that investing in people and communities is not a cost center, but a driver of growth. We're seeing this play out in the rise of Community Wealth Building, a model that focuses on creating resilient local economies. In Preston, England, for instance, a post-financial crisis effort to redirect public spending to local businesses saw local procurement jump from £38M to £112M, halving unemployment and lifting the city out of the UK's top 20% most deprived areas.

This idea is also gaining traction in the US through Employee Stock Ownership Plans (ESOPs). With 15M+ participants and over $2.1T in assets, ESOPs have been shown to double the retirement savings of workers. If just 30% of U.S. businesses adopted ESOPs, the bottom 50%'s share of national wealth could quadruple. This is not about redistribution; it's about building a more inclusive and resilient economic base from the ground up.

New Geographies of Trade Growth

The Human-Centric Era is also about a geographic rebalancing of economic power. While mature markets grapple with saturation, new hubs of growth are emerging, and none is more significant than Africa. The continent is the youngest on the planet, with a median age of 19, making it the future of the global workforce. The opportunity hasn't gone unnoticed - foreign direct investment in Africa surged 75% to a record $97B in 2024, and total merchandise trade grew to $1.5T. The African Continental Free Trade Area (AfCFTA) is further accelerating this growth, with projections of a 35% increase in intra-African trade by 2045. This isn't just about raw materials; it's about a continent building a new economic architecture in real-time.

Testing the Floor: Universal Basic Income

No conversation about human sustainability is complete without asking what it costs to maintain economic stability. Universal Basic Income pilots around the world are beginning to answer that question with real data. The world's largest long-term UBI study, conducted in rural Kenya by GiveDirectly, in partnership with MIT and Nobel laureate economists, found no evidence of the commonly feared "laziness effect."

Recipients invested, became more entrepreneurial, and earned more. Savings increased. Household income grew. Pilots in Minneapolis, Seoul, India, and Wales have converged on a consistent signal: when people have a floor beneath them, they don't stop, they build. An economy that supports human-centric innovation requires a platform.

From Consumerism to Survival

U.S. consumer confidence has fallen to its lowest level since 2021, with the expectations index dropping to 72.9 in early 2025, well below the 80-point threshold historically associated with recession. Meanwhile, corporate America has increasingly decoupled from the consumers who power its profits.

What comes next won't be announced,
it will be lived.

The signals are already visible in behavior: the U.S. repair and maintenance economy is projected to reach $1.8T in 2026. Homesteading-related searches have grown over 300% since 2020. Local food systems and urban agriculture represent a $9B market growing at 10% annually. Community land trusts have grown 30% since 2011.

These aren't countercultural gestures, they are indicators of economic adaptation.

The hypothesis is straightforward: when institutions optimize for extraction rather than service, people build parallel systems. We are entering an era not of post-consumerism as philosophy, but of survival economics as practice and utility, defined by hyper-selective spending, community interdependence, individual enterprise, and the radical reallocation of trust from corporations and institutions, to people.

The human-centered investment era isn't arriving because people woke up and willed it into existence. It's arriving because the current extractive model is exhausting consumers; and as a consequence, sparking its own correction. Much like a roller coaster at the top of a 100-foot incline, we are at the peak, unsure of what comes next.

What we know for certain is that a more sustainable future will be fueled by those who look to global human centrism as a north star.

Hannah Grey Exclusive Essay

The Economics of Intelligence

How the Compute Reckoning Creates Opportunity

For years, the technology landscape has been defined by an AI gold rush, a frenetic race to build bigger models and more addicting applications. Fueled by what feels like bottommless funding and a "growth-at-all-costs" ethos, a crucial variable was largely ignored, treated as a secondary concern to be optimized later: the cost of compute. That era is over; the next phase of AI will be defined by the **Economics of Intelligence.** We are now entering the **Great Compute Reckoning**, a fundamental market shift that will act as the great filter, separating ephemeral novelties from generational companies.

The reckoning begins when impressive technical demos collide with the unforgiving economic reality of real-world workloads. For many companies, particularly in the generative AI space, compute is no longer a line item on a spreadsheet; it is the primary driver of cost of goods sold (COGS). It directly dictates pricing, profit margins, and the feasible scale of operations. A model that seems revolutionary in a controlled environment can quickly become a financial black hole when faced with millions of active users. This shift from training-centric to inference-centric economics is forcing a painful but necessary maturation of the market. The central question is no longer "*What can this model do*?" but "*What is the unit cost of it doing that, and can we build a sustainable business around it?*"

This reality is creating fertile ground for a new generation of founders who recognize that the biggest opportunities lie in addressing this fundamental friction.

We see three distinct categories of durable companies emerging from this landscape:

1 The Intelligent Compute Layer

Early solutions competed on supply aggregation, that race is largely over and margins are compressing to prove it. The durable $100B company will sit a layer above: the coordination and pricing intelligence that makes compute usable at scale. It’s not about cheaper hardware, it will be about becoming the intelligence that sits between AI workloads and every compute resource on the planet. Not “Airbnb for GPUs,” moreso “the Fed Funds Rate for inference.” These companies will set the price of intelligence itself, through dynamic routing, workload orchestration, and real-time cost optimization that hedges compute exposure the way sophisticated traders hedge energy futures. As AI workloads become as continuous and mission-critical as electricity, someone will build the grid operator: the platform with enough market visibility, proprietary pricing data, and routing depth that companies can’t afford to operate without. They won’t just find cheaper GPUs, they’ll know which workload runs best on which hardware, at what moment, at what price. That company will have more leverage over the AI economy than any model provider. Model providers will inevitably try to capture this layer, but no provider can credibly be neutral. Enterprises won’t route everything through a vendor with a conflict of interest in the outcome.

2 The New AI Infrastructure Stack

Every technology shift produces a layer of foundational tooling so deeply embedded it becomes invisible, and indispensable. For example, Stripe didn't win by doing payments better, they won by making payments disappear. Twilio did it for communications, and AWS for compute itself. We're in the early innings of that moment for AI. The winners here won't be broad platforms trying to do everything, they'll be designed to solve a narrow, specific problem exceptionally well. Teams that reduce inference latency by an order of magnitude, or make retrieval at scale feel trivial won't just build good businesses, they'll provide the essential tools that become the quiet tax on every AI application built in the next decade.

3 Compute-Native Applications

This is a paradigm shift in how software gets built. For the first time, the cost of running a product is dynamic, workload-dependent, and potentially ruinous at scale. Most founders building applications today are treating that as an engineering problem to solve later. The most successful ones will win by treating compute economics as a design principle, not an afterthought, building cost-awareness into the product from the ground up in ways competitors who didn't can't replicate. The next Salesforce gets built by someone who understands that in an AI-native world, your unit economics are your product. The architectural decisions made in these early years will, in hindsight, separate the sustainable companies from the ones that never quite scaled. The founders who internalize this now aren't just building better products, they'll build businesses that become structurally harder to compete with every day.

The Compute Reckoning is not a headwind, it's a clarifying moment to define the economics of intelligence. It will quietly and ruthlessly determine which businesses stall and which ones achieve massive scale.

The winners that define the next decade aren't racing to build the largest models, they're racing to build the most defensible economics running *around* them. Intelligence is only valuable when it's affordable, scalable, and sustainable. For the founders seeing this clearly, the reckoning is already an opportunity.

Hannah Grey Exclusive Interview

Epic Disruptions & The Adaptive Enterprise

A Conversation with Scott D. Anthony, Author of Epic Disruptions: 11 Innovations That Shaped Our Modern World

Epic Disruptions and the Adaptive Enterprise

This interview has been edited and condensed.

Scott D. Anthony is a Clinical Professor of Strategy at the Tuck School of Business at Dartmouth College, where he focuses on disruptive change. Scott's new book - his ninth - is Epic Disruptions: 11 Innovations that Shaped Our Modern World. Scott previously spent more than 20 years at Innosight, a growth strategy consultancy founded by Harvard Business School Professor Clayton Christensen. Thinkers50 named him the world's fifth most influential thinker in 2025 and named him the world's leading innovation thinker in 2017.

To understand the disruptive change facing the modern enterprise, we must look to the patterns of the past. We spoke with strategist and author Scott D. Anthony to explore how historical parallels can help founders understand the systemic friction facing incumbents and identify the opportunities that emerge from the messy middle of innovation.

HG: For those who haven't read the book, would you mind sharing an elevator pitch of Epic Disruptions without giving away too much?

SA: The basic idea is, can we use history as a lens to make sense of the very confusing world that we're in now. You can go back and look at these 11 disruptions from the past, and you see three clear patterns. 1. Disruption changes the world for better and sometimes for worse. You have to understand the shadow that disruption casts. 2. The people who drive disruption are not superheroes, they're normal people who are curious, collaborative and persistent, which anyone can do. 3. Disruption rewards patience, perseverance, and a little bit of fun.

HG: In the book, you also talk about the four unanswered questions: Who does disruption? Is it random? Is it accelerating, and is it a universal good? Talk us through how you'd use them to frame your work with enterprises to break down disruptive change and future proof for tomorrow.

SA: Thinking about AI and where enterprises are today, question two, is it random, and question four, is it a universal good, feel the most relevant. My conclusion from studying and living disruption is disruption is predictably unpredictable. You cannot look at it like Newtonian physics. It just doesn't work that way. But there are some very clear patterns. Look at the origins of the transistor - it was a technology that was developed to do one thing - replace vacuum tubes in communications networks.

It ultimately did that, but that's not where it started because it wasn't good enough in its early days.

In classic disruptive fashion, it traded off performance to make things a lot simpler, a lot more affordable, a lot more flexible, but it wasn't good enough for mainstream applications. It had to start in a different market - hearing aids, transistor radios, and so on. Artificial intelligence, I think we see the same thing. A lot of the use cases are places where the competition isn't 'can we do it better than the most sophisticated solution' it's 'are we giving people something that's better than nothing at all,' or 'giving them a simple, accessible solution that allows them to do what they historically could not.' That's a clear pattern.

Another key pattern, what often drives the disruption isn't the technology, it's the special sauce of the business model, how value is created, delivered, and captured (at the core of the McDonald's case in the book). That's where the answer to 'are companies going to stand in the way of or will they drive the adoption of' lies - it really depends on the business models that startups are following, and how they dock into the business model of incumbents. When it makes sense for an incumbent to do something, it will. That's a very clear pattern.

Then the last thing I'd say about the fourth question, "it is universal good," is it's not. There is an upside to disruption. I believe in its power generally. It democratizes markets, it creates growth, and so on. But there are losers.

There are people who had great companies that no longer have great companies. There are people who were really good at doing something that have that something no longer matter.

In the middle of disruptive change there's a lot of mess.

A very clear example of this is from Henry Ford's Model T in the 1920s. In cities in many parts of the United States there was almost a holy war between the drivers and the pedestrians, because cities were not built for cars, and people were getting injured and even dying. In AI right now, there's a parallel - you've got use cases like deep fakes and mis-information which is not very good.

To get out of the messy middle for automobiles, you needed technologies like traffic lights. You needed norms like who goes first when we get to an intersection? And you need regulations like driver's licenses and speed limits. We kind of have forgotten that lesson with AI, it seems to me, and I think that's not a very good thing.

HG: So fascinating - it does feel like we're trying to fit AI into the software model in order to accelerate enterprise adoption, without stopping to consider things like business model, legislation, regulation etc. What's the advice to the software business who's speaking daily to their enterprise customers, and what's the advice to the AI company who's sort of trying to take that spend away?

SA: Destin Sandlin, an engineer, asked the question, “can you re-do my bike so if I turn it left, it goes right, and so on?” It takes him 8 months to learn how to ride the backwards bike, because he first has to unlearn how to ride a bike, and it takes his six-year-old son a couple of weeks, because he has to unlearn a lot less.

The challenge is that the biggest incumbents have the deepest pockets, but they also have the most unlearning and unwiring to do. History suggests taking a more decentralized approach, something where a department or an individual can enact change without needing the entire organization to agree. Because AI involves knowledge work, I think that’s particularly true. The unlearning is going to take decades for the big companies.

HG: To that point, how much of disruption is the product itself, it being better, faster, cheaper, etc, versus how much of disruption is the people, the mindset to adopt something new, the willingness to try, the ability to experiment with breaking the mental model?

SA: In the book, all of the case study protagonists are companies, so ‘Procter & Gamble did this’ etc.

> You go and you look at the history and no company does anything, it’s the people inside the company that do *everything*.

Of course, technology needs to be created, and this could change in the future, but at least up until now, technology is created by people. People who think of things and try things and find out this works and this didn't work. So at the end of the day, to me, the most important thing by far is the people and the way that they're thinking and rethinking about things.

HG: How can companies instill that flexibility in their culture?

SA: It's on my list to understand more deeply - you look at companies that have driven transformation, and you think of all the big tech companies. But you look at a company that has outperformed markets for going on 15 years, like Schneider Electric, and driven a lot of digital change, and as I understand it, a lot of it is driving a real culture of learning and really focusing on how to continually adapt. I have no idea if there's any fun in that, I hope there is, and I certainly have proselytized to people that you need to bring joy, happiness, fun because if you don't it's really hard. You have to acknowledge that it's really hard, and even if it's fun, it's still painful.

In my 2017 book, Dual Transformation, one of the Co-Authors, Clark Gilbert, was describing the way, when he was the CEO of a media company and President of a University, he would staff the teams, and he said, "you're creating a healthy mix of aliens and diplomats." Aliens will see things differently, diplomats know how to get stuff done. Too much of either is bad, to get that balance right can be really powerful but not always easy to do.

HG: And in terms of seeing things in the future, how can we project these lessons forward to understand where the next phase of epic disruptions will emerge?

SA: I always urge people to experience tomorrow today. William Gibson said, "the future is already here, it's just not evenly distributed." You can spot things early if you look in the right places and get a sense of how the future is going to unfold. At the end of the day, I am an optimist. Yes, one sees very clearly that there are winners and losers, and there will be tasks and jobs and roles that are lost. But the power to enable a greater part of the population to do what once required deep expertise is immense, and that's going to do a lot of good, I believe. More good than bad.

> "The first iPhone came out in 2007 and we're at version 17, but with AI it's moving exponentially faster. I can foresee us being on the sixth generation of AI before the end of this year, and AI is only 2-3 years old."
>
> *— Gregory, Rural TX*

Signals from the Undercurrent: The Future of Healthcare

Themes Generated from Hannah Grey's VibrationsOS Foresight Engine

For the first time in history, a patient can walk into a doctor's office knowing more about their own biology than the clinician across the table. The healthcare system of today isn't built to support this information asymmetry. For generations, healthcare was a destination; a place where care was administered as a reaction, by a clinic or hospital, to a problem you didn't see coming.

Today's model and approach to healthcare is being dismantled - the center of gravity is shifting away from the institution and toward the individual, driven by a demand for agency and the technological shifts already in motion. What's emerging isn't a better hospital, it's an entirely new ecosystem built around engineering human vitality, moving beyond managing illness towards a future where health is designed with intention.

Below are four signals demonstrating the nature of this change and the white space where we see new companies emerging.

1. The Sovereign Health Consumer

Every major platform shift - mobile, cloud, and social - created new infrastructure layers worth hundreds of billions. The migration of health data from institution to individual is the corresponding moment for healthcare. It's still early innings, and still very fragmented.

While the initial shift saw the "patient as the new point-of-care," demanding consumer-grade experiences from systems built for administrators, this trend is now maturing into the era of the "Sovereign Health Consumer," an individual who owns and controls their biological data end-to-end, and expects the ecosystem to be built around them. That ownership will unlock an entirely new market. Companies like Function Health and Prenuvo have proven consumers will pay for access to their own healthcare data and the insights derived from it, but results from a single lab isn't sovereignty.

The biggest opportunities aren't next-generation wellness apps, but the 'picks and shovels' infrastructure beneath this new consumer-focused healthcare economy.

As data continues to move out of institutional silos, we're forced to ask, how does public health thrive in a decentralized world?

The missing layer is privacy-preserving technology that enables population-level insight without surrendering individual control.

The company that becomes the trusted infrastructure layer for personal health data: the identity layer, the permissioning layer, and the exchange, won't just be a healthcare company. It could become one of the most consequential platforms of the next decade.

2. The Management of Personal "Bio-Capital"

As individuals gain control, the objective evolves from pursuing "healthspan" to actively managing personal "Bio-Capital." This reframes vitality as a measurable asset class to be cultivated, grown, and invested in.

The engine for this is the Personal Intelligence Layer: a trusted, lifelong AI companion that acts as a fiduciary for an individual's health. It surfaces faint signals of risk before they become disease, and recommends hyper-targeted interventions through nutritional protocols, preventative screenings, or specific peptides, all tailored to your biology.

While the first-order product is the AI itself, the second-order opportunity is the platform it sits on. If there will be an operating system for human biology and personal intelligence, what does an app store built on your digital twin entail? No one has had access to biological optimization at that scale.

3. From Insurance to Investment

US employers spend over $600BN annually on employee health benefits. For decades, employee benefits meant insurance. As the Sovereign Health Consumer enters the workforce, employers are being pulled into a fundamentally different relationship with healthcare. The most forward-thinking companies are already making the shift, from passive procurers of insurance to active investors in the biological capital of their people. That reframe has profound consequences.

The C-suite will demand evidence of the effectiveness. We believe this could create a new category that quantifies the ROI of employees' potential, sitting at the intersection of the massive benefits market, the emerging bio-capital economy, and the enterprise analytics stack. That layer connects an individual's Personal Intelligence with an employer's need for workforce insight, without compromising the individual data controls that makes it trustworthy in the first place. The companies that build this won't be selling a new wellness benefit. They'll be leveraging data to help the C-Suite understand the correlation between healthier employees and better business outcomes.

> **"I will want the human interaction. I know the statistical success is better with AI, but human-to-human is best with me."**
> *— Justin, Suburban LA when asked to choose an AI or a human doctor*

4. The Rise of the Human-AI Collaboration

AI doesn't replace the clinician, it augments and elevates them. A clinician today drowns in data: lab results, imaging, genomics, wearables, and patient history, and is expected to synthesize these sources in a 15-minute appointment.

AI will sit between them and the complexity, surfacing what matters, predicting what's next, and translating a patient's Personal Intelligence into something actionable for the patient and their family. The doctor's role shifts from gatekeeper of knowledge to expert interpreter, and in doing so, becomes dramatically more valuable, scalable, and effective, while also requiring upskilling, re-training, and increased patient-as-the-customer service.

What makes this opportunity defensible is the loop itself.

Every clinical interaction is a potential deployment. More importantly, every judgement call, correction, nuanced decision a clinician makes alongside the AI makes the platform smarter. The human isn't just a user, they're a contributor to a compounding data moat that becomes harder to displace with every interaction. The platforms that capture this loop won't just be software, they'll be institutional knowledge impossible to replicate.

What unites these four signals is a single underlying shift: the future of health is participatory, predictive, and deeply personal. Value is migrating from the institutions to the individuals and the intelligent platforms that serve them. For the teams building this new architecture, the opportunity is not just to build a company, but to define the next era of human health.

tzafon

Research From The Lab:

Training Vision Language Models for Computer Use Agents

Contributors: *Nikita Khomich*, Leopold Pluto Hermansson*, David Dinucu Jianu, Ido Hakimi, Yerniyaz Nurgabylov, Noga Bregman, Simon Koser, Noah Löfquist, Mark Rogers*

**Core contributor*

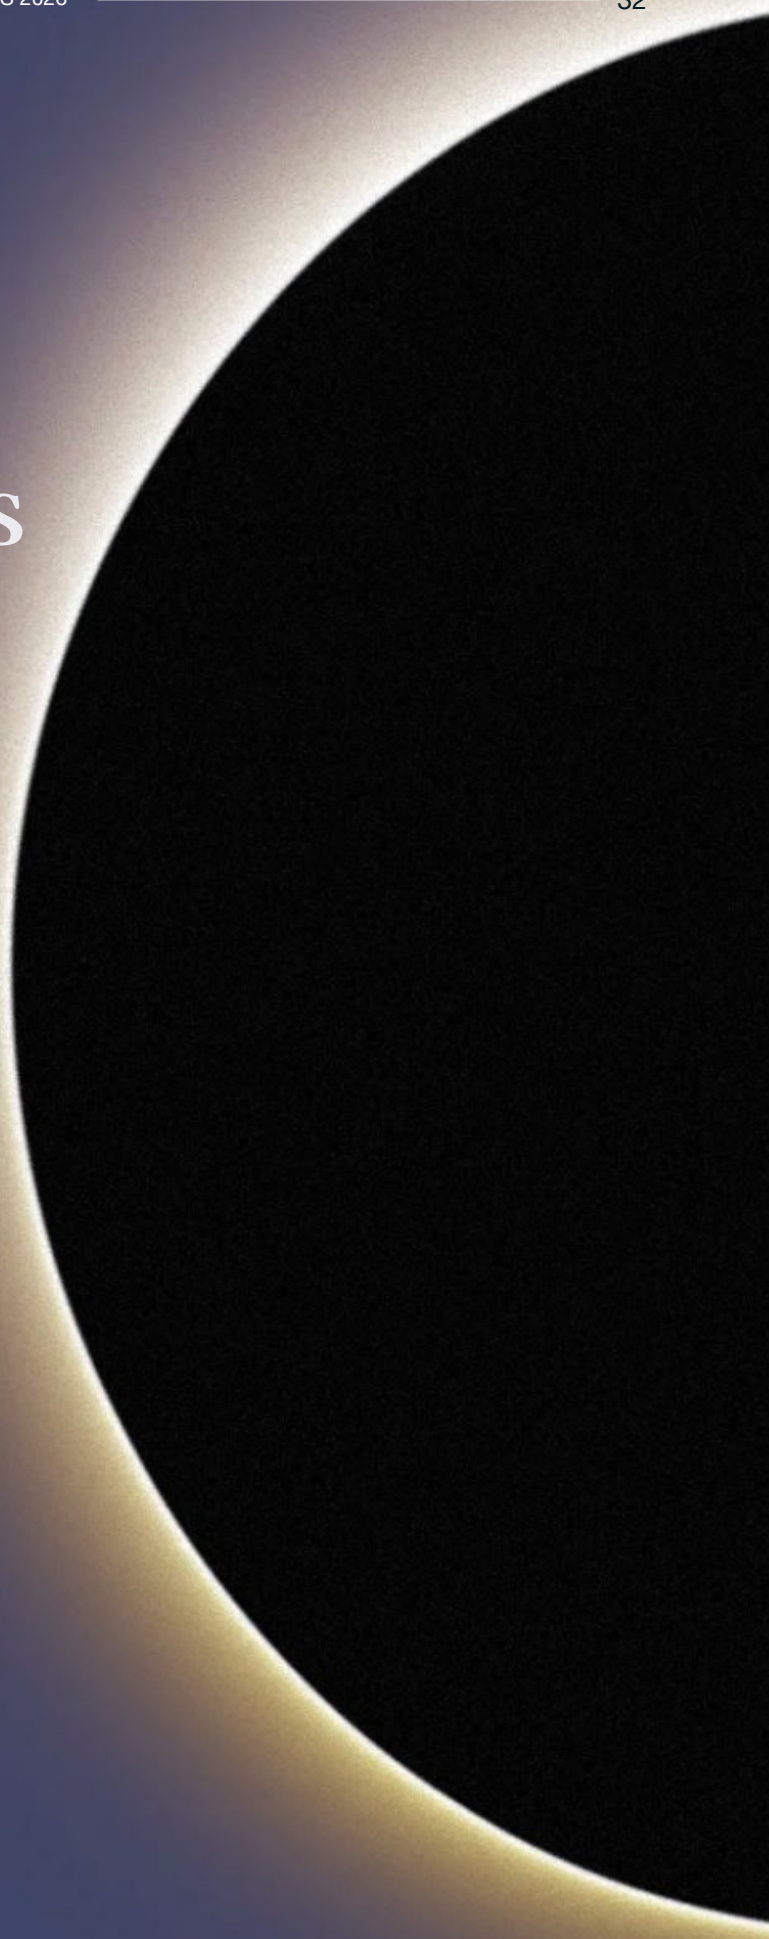

How Tzafon trains Vision Language Models for Computer Use Agents using reinforcement learning, addressing the limitations of SFT and improving generalization across environments

The modern enterprise is drowning in a sea of software. While digital transformation has armed us with more tools than ever, it has paradoxically created more friction. Human productivity is now consumed by the act of serving as the manual, digital "connective tissue" - the endless copying, pasting, and reconciling of data between applications that don't speak to each other. This is the stubborn, final mile of enterprise automation that traditional RPA and API-based integrations have failed to conquer.

Tzafon's new research on training Vision Language Models (VLMs) for Computer Use Agents (CUAs) signals a paradigm shift. It heralds the emergence of a new class of AI that navigates our digital world not through code, but through sight. This is an AI that can see an application's interface and use it, just as a human would, taking actions within the workspace in the same way an employee does. For Fortune 500 companies, the implications are profound. This technology is the key to unlocking the value trapped within decades of legacy infrastructure, automating the "un-automatable" tasks without costly system overhauls or digital duct-tape.

While the initial promise of enterprise AI focuses on incremental efficiencies like cost-cutting, the true, tectonic shifts lie in the fundamental re-architecting of the nature of knowledge work. When any employee can train a "digital apprentice" to handle their most mundane tasks, they elevate their role from the execution layer to the strategic, redefining roles and responsibilities up the entire chain of command.

The path to adoption is not a technical challenge, but a cultural one, demanding new frameworks for governance, security, and human-machine collaboration. Designing the organization to embrace the new-wave will be the defining strategic, executive-level work of the next decade. Tzafon's research provides the blueprint.

Many have tried to solve the task of getting an LLM to use a computer. Until recently, these have all mostly relied on SFT. The reason for this is that it's easy to define, e.g. just figure out the task you want to solve and collect data, and you can get fast gains on targeted benchmarks. However, what you quickly realize is that without a lot of regularization this saturates after 100 to 1000 examples on a task and beyond this we start to get degradation in other abilities, making it a whack-a-mole game.

Even more problematic, the improvements do not generalize. Without generalization we lose out on the most important ability of a LLM – being able to work across many environments – which is highly expected in real computer use applications.

The reasons for the lack of generalization from SFT are several, not exclusively listed:

1. The first order learning mechanism is to memorize, so the model doesn't learn why an action should be done but rather learns that it should be done when it observes a state
2. The penalty of incorrect actions are often not compatible with how language is modeled. Take the simplest example of trying to click a button: we might say that the ground truth action is to click on position x and y, what we would expect is that the model would in some sense learn it should click the button that is under x and y, but what actually happens is that all the coordinates except the exact coordinate that was given in the example will be negatively penalized uniformly, so clicking 1 pixel away is the same as clicking on the other side of the screen.
3. The above problems exist in all ML tasks, but what makes it very difficult in VisionLM, is that we have extremely large amounts of input data compared to the prediction we want to make – think of it as the number calculation/decisions made by the neural network to answer the question where to click and then compare this to language to language models that have the same amount of input and output data.

See loss from training Qwen3-VL on SFT UI data

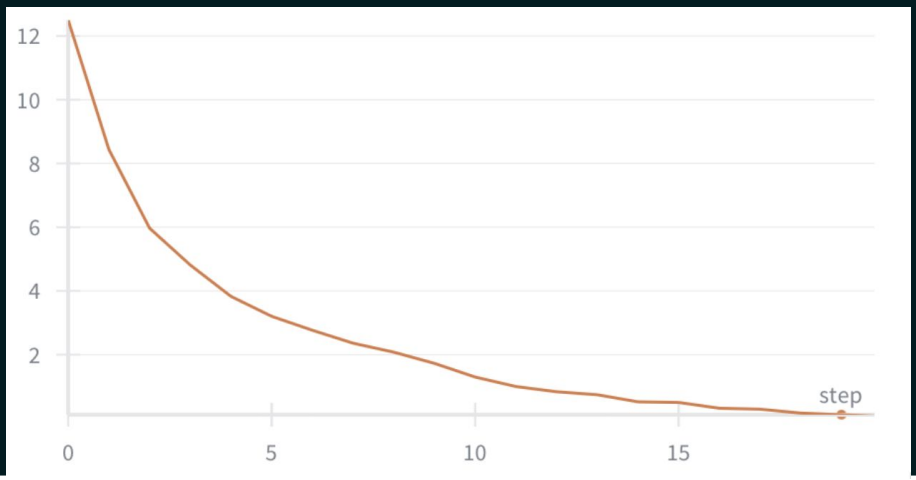

The first order issue we address is the model's ability to perceive, how we can improve the vision encoders and output a strong signal that improves the model on downstream capabilities. Second up is generalization instead of memorization, which comes down to the model learning to robustly navigate states that were not encountered in the training data, especially states that stem from environmental noise or changes.

Side-Quest: Internal Representation

The perception problem is a deep rabbit hole, because of all the layers of complexity and ambiguity. In general we don't want to think of the image encoder as a separate component, as we can't really say what we want it to do. We can reason about the fact that we expect that the embedding/outputs should contain some semantic and positional information that can be used later by the decoder, but to what extent the image encoder should transform the image or retain a lossless representation is unclear.

What we can carefully do instead is make experiments and observations that give us insight into what the model might or might not be doing.

Image encoders, particularly models like Qwen3-VL, function by transforming an input image into a set of patches. These patches are then processed through an image transformer, where the information in each patch is treated as a token. This process involves multiple layers where each token (representing a patch) attends to all other tokens in a non-causal 2D (or 3D for video) manner.

By observing the intermediate and final image patch embedding we can understand how the information is distributed and evolves.

Qwen3-VL implementation for reference

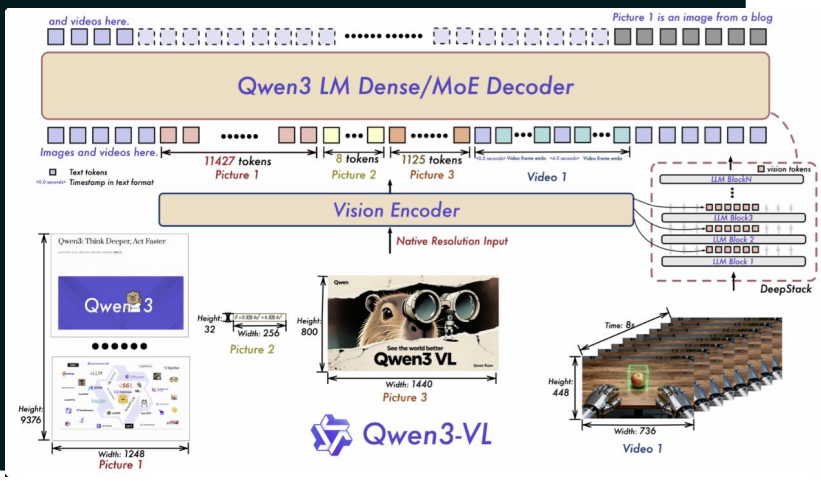

This can be done by comparing the cosine similarity of pairs of embedding, which tells us how the information changes through depth. What we see directly is that even though the model has ~30 layers, already after layer 2, the data is very similar with the last layer, hinting at the fact that a strong representation of the content of the image has already been formed.
Next, we can look at how it changes across the image dimension. To do this, we will first pick a steering vector that describes the semantic meaning of an object. What we can do next is compare this vector with the image patches and how much they contain apples. Below are examples of patch representations, looking at how the information about the apple is represented across different images.

See loss from training Qwen3-VL on SFT UI data

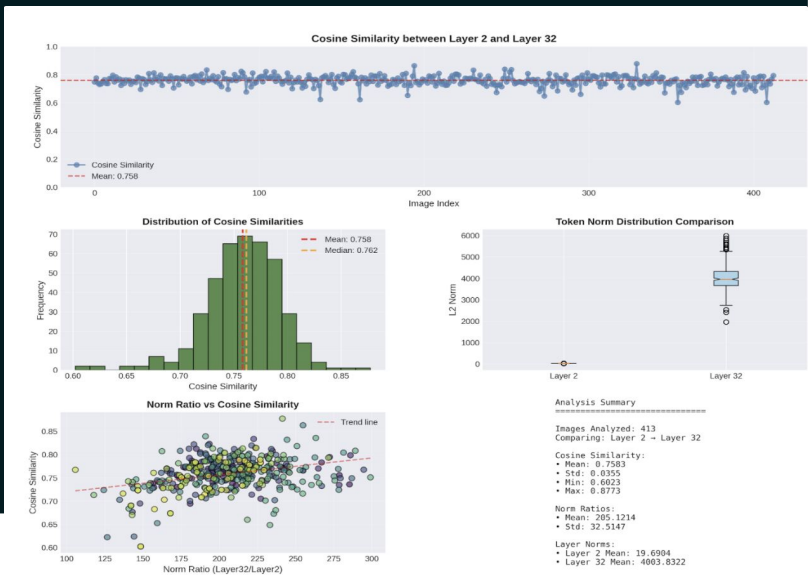

Next, we can look at how it changes across the image dimension. To do this, we will first pick a steering vector that describes the semantic meaning of an object. What we can do next is compare this vector with the image patches and how much they contain apples. Below are examples of patch representations, looking at how the information about the apple is represented across different images.

Next, we can look at how it changes across the image dimension. To do this, we will first pick a steering vector that describes the semantic meaning of an object. What we can do next is compare this vector with the image patches and how much they contain apples. Below are examples of patch representations, looking at how the information about the apple is represented across different images.

Apple Token Analysis

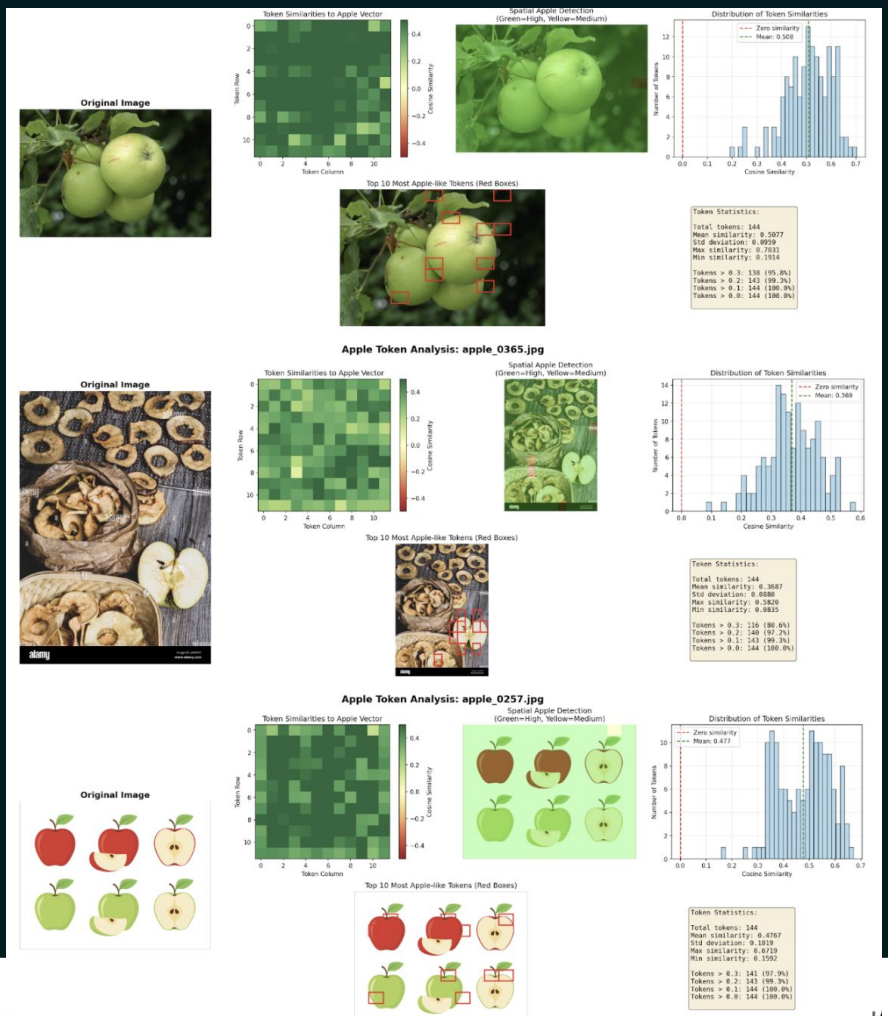

What these experiments indicate is that from all image patches in an image, we can probe what object is present, and the position – so each patch knows global important information. What this also indicates is that the image tokens stream is closer to a look-up table, in contrast to the initial patch-embedding, that only has very local information, the final representation has compressed data of the entire image.

Side-Quest Position Decay

Positional decay in the image encoder. To explain this we need to understand that the model's ability to understand where in the image things are comes from two positional mechanisms. Starting with the 2D-RoPE, which is applied at every attention layer by rotating the query and key vectors based on the relative offset between patches. This helps the model understand spatial relationships between patches, but it only encodes relative position — it can't tell the model where a patch sits in absolute image coordinates. To fix this we train an additive patch embedding: we take the information in a patch and before it passes through the attention we add some extra values unique to each location. Which means that each token knows where it is and can query and infer information about specific locations in the image.

This matters because CUA tasks require outputting absolute coordinates, and for that the only signal is the additive patch embedding. The problem is that this embedding is added once at the input, and then at each layer (there are 20-30 in Qwen3-VL depending on size) we get new information from the other tokens added back. The reason the model can't really learn to overcome this is numerical stability — to keep training stable we have to normalize the information vector at each layer, which means the original positional data is reduced with exponential decay. Since 2D-RoPE doesn't carry absolute position, it can't compensate for this loss.

This is something that Qwen3 tries to mitigate with something called deepstack, which takes data from earlier layers of the image encoder and adds them to the later stages. In the original paper they do show some improvements, but what we saw was that completely removing them without any retraining has minor impact, which likely means that this isn't enough to solve the core issue.

Another experiment we ran to convince ourselves that this is actually a problem was to just scale the positional embedding without any other changes. What we see is that in images with very few objects accuracy improves dramatically (going from 40% to 80% click accuracy on a benchmark that consists in clicking a red ball). (To reproduce the experiment we found that scaling the positional embedding by 3 on Qwen/Qwen3-4B-VL-Instruct). With this evidence, we do believe that VL models would gain from having a stronger absolute positional embedding.

Reinforcement Learning

Given the issues we faced with memorization and training on discrete coordinate tokens, we are pushed in the direction of RL. The RL setup offers many benefits to what we have looked at before, the most obvious being how the loss is applied. By using a GRPO loss we can give a reward depending on each model actually achieving the task and we can also add secondary rewards to enforce more robust behavior like clicking near the center of buttons.

By adapting prime-rl (https://github.com/PrimeIntellect-ai/prime-rl) to support multi-modal (which is supported as of Feb 2026 out-of-the-box) we can convert either synthetic or annotated datasets with bounding boxes, into rl environments, where the reward is statically calculated based on the click coordinates and the ground truth bounding box. What is even more surprising, is that we do not really need a realistic environment, by only training on limited and fabricated test environments. The model nicely learns, but more importantly, it generalizes across other benchmarks. On an aggregated ui benchmark, we see 0.39 -> 0.53 improvement. Even though we only train on generated simplified environments. What is more impressive, is that this is better than the performance we see SFT training on UI datasets.

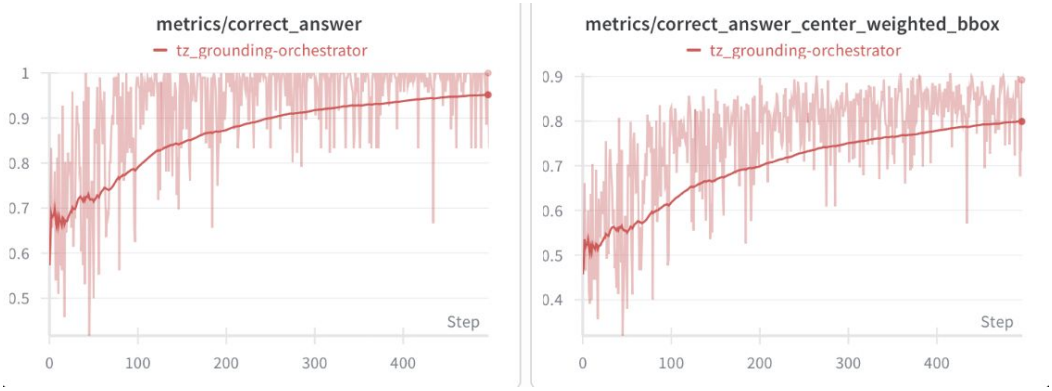

Multi-turn

Given that it seems like we can robustly teach the model on a single turn, we now move to multi-turn environments. This means that the model now interacts with an application in a loop and then the action is executed and a new state is generated, until it runs out of steps or completes the task. Though, we are happy to see the reward go up, what we rather care about is why/how it goes up.

For context our training is on 100 environments that require 3-15 click interactions to succeed. The environment mainly tests abstract capabilities instead of replicating real apps.

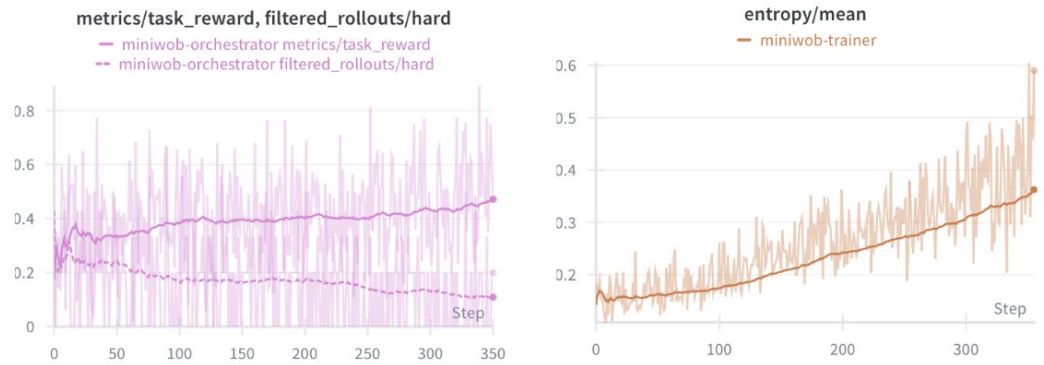

Other than improving on the training environment, we see an improvement in accuracy and ability to execute long horizon tasks in benchmarks like OS-World, where we get a 20% absolute improvement on the Chrome category — even though our training environment has no resemblance to OS-World's environment or tasks.

Some key observations:

- Reward goes up 🚀
- The model learns to solve new tasks, which is a good sign that the range of difficulty is good, and that there is some amount of generalisation in what it learns.
- From the entropy and the actual rollouts, we can see some interesting behaviors. The entropy increases quite steadily, from what we can attribute this to the model exploring more in the text space before generating an action, as in most RL we let the model think for a limited amount of tokens before the answer. Specifically we see that the traces here become more varied and also more informative.

For example, the model realizes errors or unintended outcomes based on what it planned to do and what actually happened. One key emerging change is that the model becomes less likely to repeat itself. This stems from the fact the models have a strong bias to repeat what is already in the context, which in action is problematic. What we expect from a robust system, and what we see after the multi-turn training is that the model realizes an interaction failed based on the history, and decides to either try something entirely different, or at least adjust the action slightly and try again. These are behaviors that cannot be derived systematically from SFT because they require incorporating the behavior the model has. For example, the probability of trying again or trying something else, depends on the model's ability, so simply mimicking a human does not train this.

These observations tie in to the final fact that robustness is key for solving long horizon tasks. Specifically the success of an agentic model is only partially related to accuracy/performance and more about robustness. To show this we create a table of how accuracy has to scale to achieve a success rate on a certain horizon, if the success of each step is independent and every step has to succeed.

Recovery from failures > Click accuracy

Trajectory Length	50% success rate	80% success rate	95% success rate
1	0.5000	0.8000	0.9500
2	0.7071	0.8944	0.9747
4	0.8409	0.9457	0.9873
8	0.9170	0.9725	0.9936
16	0.9576	0.9862	0.9968
32	0.9786	0.9931	0.9984

Even if you have a high failure tolerance, e.g we can retry the workflow, the accuracy required when the number of steps increases becomes impossible to achieve. Instead, what we want to do is improve the model's ability, to recover from failures and adapt to slight out-of-distribution variation when the model is deployed compared to in training.

OS-World – 50 steps comparative results

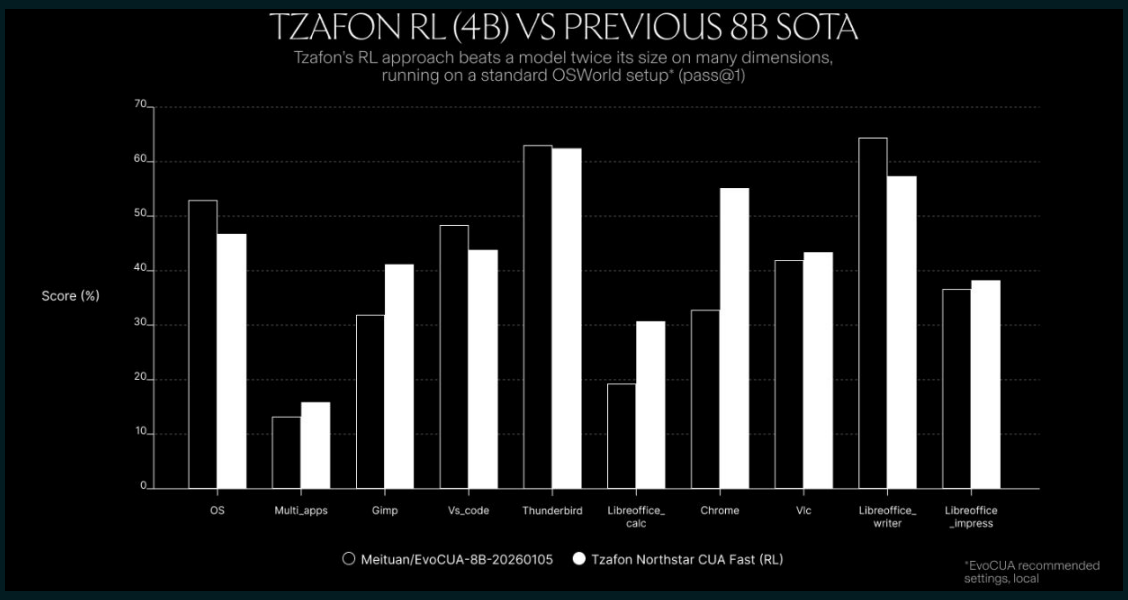

Using Evocua agent code, and compared to current open source sota in similar size. (EVOCUA-8B avg 32.5% vs Northstar CUA Fast (RL) 37.0%)

Reality Somms

How to Cultivate A Discipline of Knowing What's Real

How to Serve Proprietary as a Business

Matt Klein, Head of Global Foresight at Reddit
Brian Lange, Co-Founder of Future Commerce

The following is an excerpt from **"Reality Somms: How to Cultivate A Discipline of Knowing What's Real."** *Read the full article on ZINE.*

01. Cultivate Synthesia

Visual and auditory data are abundant, but draft us an incomplete story. If reality is so difficult to gauge, then the first logical step would be to use *all* the tools we have in order to discern truth. We need to unite all of our senses.

Sommeliers are required to engage with source material – wine's essence – and grapple with "percepts" or a viewpoint formed during direct encounters. They must identify a wine's qualities, effects, relations, and when, where and how it was made. And at their best, name the winemaker, specific vineyard, environmental conditions and precise vintage year.

These percepts lead to phantasms or mental projections of the wine in their imagination. Spurred by just a single whiff, somms can recall these mental projections. Language is then created around this discipline, not just as a way to convey information to others, but as a way to help build and maintain the phantasm, a robust fantasy.

> ***Percept:*** *the impression created by the external senses upon encountering reality*
> ***Phantasm:*** *the mental image created by the internal sense (the imagination)*
> ***Concept:*** *the abstraction created by the intellect through recognition of the essence*
>
> *Source: The Trivium by Sister Miriam Joseph, C.S.C., Ph.D.*

Creative business leaders must develop a discipline around forming phantasms out of percepts, rather than just relying upon existing concepts that are passed around via word or image.

This combination of senses is something Eric McLuhan (media theorist, Marshall McLuhan's son) explores in his book The Sensus Communis, Synesthesia, and the Soul:

"Technology extends one or another sense or faculty, according it a sort of hyperesthesia, which has then the effect of numbing the bodily sense extended and rearranging the interplay between the other Senses."

When we use technology to extend our reach, we often don't encounter the "essence" of anything.

Percepts are practicably unattainable when senses are artificially extended.

Now with images and videos as primary technology, the phantasm is thrusted upon us. *Artificial phantasms.* We form a mental projection of a thing before we ever encounter its essence. We judge before we ever encounter it.

Digitally focused brands have collectively spent billions of dollars to provide us with pre-baked phantasms of their products. Hollywood, social media platforms, AI companies, etc. have done the same for just about *everything*. But...

We're tired, un-entertained and skeptical. We're overloaded with pre-baked mental projections of a reality we have no real experience with. We can no longer trust our eyes and our ears — our most relied upon sensory organs.

The remedy? Scents, tastes, and touches also communicate. Dive deep into them. Become curious about what they each mean and how they function. Play with them, and allow others to do so as well.

If you don't know the difference between the scent of violets and ethanol, you're in trouble. If you don't know what wet basement tastes like or means, you're in trouble. Not just because you might serve someone a corked bottle of wine, but because if you don't understand key data points of the human experience, how are you ever going to resonate?

This is engagement with the *real*. If you're lost in spreadsheets and algos and never engage with reality, you're bound to miss the metaphorical "wet basement" taste of your market, audience, or world you wish to embody.

> ***"If it looks like a duck, quacks like a duck, and swims like a duck, but doesn't smell like a duck, it's not a duck."***

You need to know the smell of a duck so that you can "see" the duck through your nose better than through your eyes. And if you're dealing with digital-only, you're separated from the source. People assume they already understand something just because they've seen an image of the thing.

Becoming a sommelier provides you discipline to become fully intimate with reality. Encounter the real, observe it, build language or something that relates to it, and then give people opportunities to encounter it themselves. If they make their own observations and understanding as a result, you're creating value. If you've built something real, then you can give up control of the narrative and let reality do the work.

Make things that are connected to reality and then push people to encounter reality themselves. *They'll buy that.*

02. Slow Things Down

Propriety is about timing.

And one of the reasons why we've lost a sense of propriety is because we have a tendency to only speed things up. "The faster the better," so we say. Sure there's value in speed, but if you don't understand timing, speed for speed's sake will result in a total loss of value. How may we slow things down so we can get a better sense of when strategies, words and participation is most appropriate?

A sommelier would never chug Domaine Romanee-Conti (DRC), a legendary French producer. For every wine, there's a whole process of observation, smelling and swirling before the first sip. While the effects of alcohol are certainly part of the enjoyment, the true benefits come from understanding, savoring, and timing the consumption.

With most things now, we mostly measure "efficiency" or the speed at which we become drunk. But if everything is about speed, now is the time to slow things down.

In *Culture is our Business*, Marshall McLuhan observed:
"Art is new perception. New art opens new worlds for our recognition and nourishment. Psychically, art is valuable only when new. Commercially, new art is kooky and worthless. The gap between the kooky and the commercially valuable is closing fast."

Paradoxically, the illegible is where value is created.

Value emerges when we have to strive to understand meaning and truth. But now, the speed of conversion to commercial value is instantaneous. The metaphorical DRC is in our mouths immediately, without any processing, appreciating, learning, or pairing. The "new" is absorbed and digested in a thumb-swipe, while the new language forming around it (i.e. memes, commentary, and content), is remixed and consumed just as fast. Novelty is digitally digested, digitally judged, and its value is consumed through our digitally extended senses.

But this speed is incomplete.

If you create something new, and you know it's valuable, slow down. This could be as simple as a new feature, product line, variation, or collab. It could be a new article or a film. It could even be a new campaign to an old thing. But if you're going to ship, consider releasing it in bits. People need to feel like they're missing something... because they are.

A digital perspective is an incomplete perspective.

People haven't actually grappled with "the thing" itself, they've encountered a mere *phantasm*, which they accept as the "real thing." Therefore, to assess the (real) value, they'll need to spend (real) time with the thing. And so, the slower you go, the longer the value cycle is.

This might feel frustrating because, *ughh time*... but if you try to speed up the cycle of capitalization, you risk incurring a false sense of having absorbed value when you or the consumer hasn't. Friction is invaluable and why we hear it discussed so frequently today.

Find a way to slow down absorption and increase the savoring. In our increasingly digitized and seamless world, this is hard. But less is more. Make mysteries to be solved. Scatter breadcrumbs. Keep secrets. The "fidelity" of information needs to be so low that people spend more time making sense of things by filling in the gaps on their own, or by hunting down the missing pieces themselves. *They will.* Consider our infatuation with true crime, conspiracy and fan theories, easter eggs, and fanfic. "Gritty over pretty" to quote Marshall McLuhan's grandson Andrew.

Consider: the digital channel should do the exact opposite of what we have all been optimizing it to do. Rather than optimize for digital engagement (more injected phantasms and the fastest way to burn through the value of something), the whole point of a digital engagement should be to tease and tingle the extended senses. Provoke, inspire and invite people to grapple with the real themselves. Allow a piece of digital content to exist as an ember, inflaming a larger conversation or experience.

Cute kitten videos shouldn't inspire watching more cute kitten videos, they should inspire people to experience kittens and enjoy them away from a keyboard.

We consume content like toilet paper. And if you continue to play only with content as your material to make your worth known, you'll be flushed the moment someone finds something more interesting. Most of you want them to never stop wiping. Stop that. Slow your roll. The point isn't how much of your toilet paper one uses – it's to get people off the toilet.

The best experience with a wine is when it is aged to its “drinking window.” This requires patience. When you understand the reality of a situation, you’ll know that time itself will do the work for you.

03. Build Enclaves

A somm’s whole job is to educate, assist with appreciation, and pair perfectly to enhance your dining experience. But larger, they create an environment for all to deepen their appreciation and sense of wonder. They give people the space to explore, to try things they would have never tried, to appreciate nuance and craft, and to receive the maximum enjoyment (i.e. value) out of an experience.

Like a somm’s table or tasting room, we need to create “spaces” for people to engage with the real. These enclaves should strive to include an embodied component.

Enclaves that encourage curiosity open people up for personal engagement. Within enclaves, people witness others doing the same. Consider offering people the infrastructure, attention, and encouragement to relate and document their own personal experiences with reality.

We’ve already touched upon a few ways to build enclaves: spark curiosity through mystery, and play with secrets. Bring back je ne sais quoi (“I don’t know what”), not for the purpose of self-glorification, but to show people there’s something more you’ve encountered that they have not yet. Consider poetry, which requires more involvement and work to understand. Invite participation. Make it social.

Posting a “good” piece of content online and then seeing the likes or comments roll in gives zero indication if anyone actually interacted with reality. *The opposite is likely the case*. It’s more likely they interacted with a phantasm. Most people would rather be handed phantasms than attempt to create their own. Instead, find ways to spark their curiosity.

The co-curious always find each other, and especially if you set the table for them to interact (meetups, events, websites, forums, social media, physical spaces, stores, pop-ups, etc.).

Encourage them to experiment, to push beyond what they've attempted before. Don't over explain. Let work speak for itself. Sometimes you'll have to provide the "scaffolding" (in an educational sense, prompts, queues, and tools), but don't construct the whole thing. The goal is to help people progress to the point of having agency around the thing. This is McLuhan's "coolness" or lack of fidelity. Let people fill in the gaps with their own takes.

We're aiming for curiosity over certainty. Exploration over achievement. Questions encourage more questions.

Reality is hard to engage with alone. Enclaves make the work sustainable and enjoyable. Curiosity compounds in groups. Ensure you're creating spaces where that compounding can happen, then show up with new questions, provocations and teases. The community will do the rest.

The Best Somms are The Best Storytellers

The best sommelier is the one who can tell you the most compelling narrative about juice.

A somm can tell a story by simply pairing a sip with a bite, and leaving their guests to enjoy their own conversation. But the best somms seed a conversation and are able to gauge how much to tell the table about the wine. Some somms will spew a few facts or tasting notes, but the best will weave them into a narrative. Timing, jokes, unique facts, mystery and mystique – all the ingredients required to guide people toward enjoying the reality of what they're consuming. Telling this story properly requires getting to know who is at the table. To feel them out in real time. This is all part of propriety. It requires knowing both the who and the what – the objective and the subjective united.

Save yourself; dismiss false phantasms.
Microwaving a meal is not becoming a cook.
Scrolling Instagram is not friendship.
Sorting a spreadsheet is not getting to know your customers.
Yes, touch grass, *but touch it all.*

An imagination isn't just for conceiving what could be – potential realities – it's for holding onto what's *already real*.

If you can do that – hold on to reality – *you'll know what is possible.*

And then, whatever you do, don't just hand over your formed viewpoint to others to digest as their own; don't perpetuate artificial phantasms. Instead, assess the situation and ask yourself - "what is the proper amount of time and attention?"

Let that propriety guide your strategy for helping others find truth. You can rest assured that those who are ready will join you. Let reality do its work.

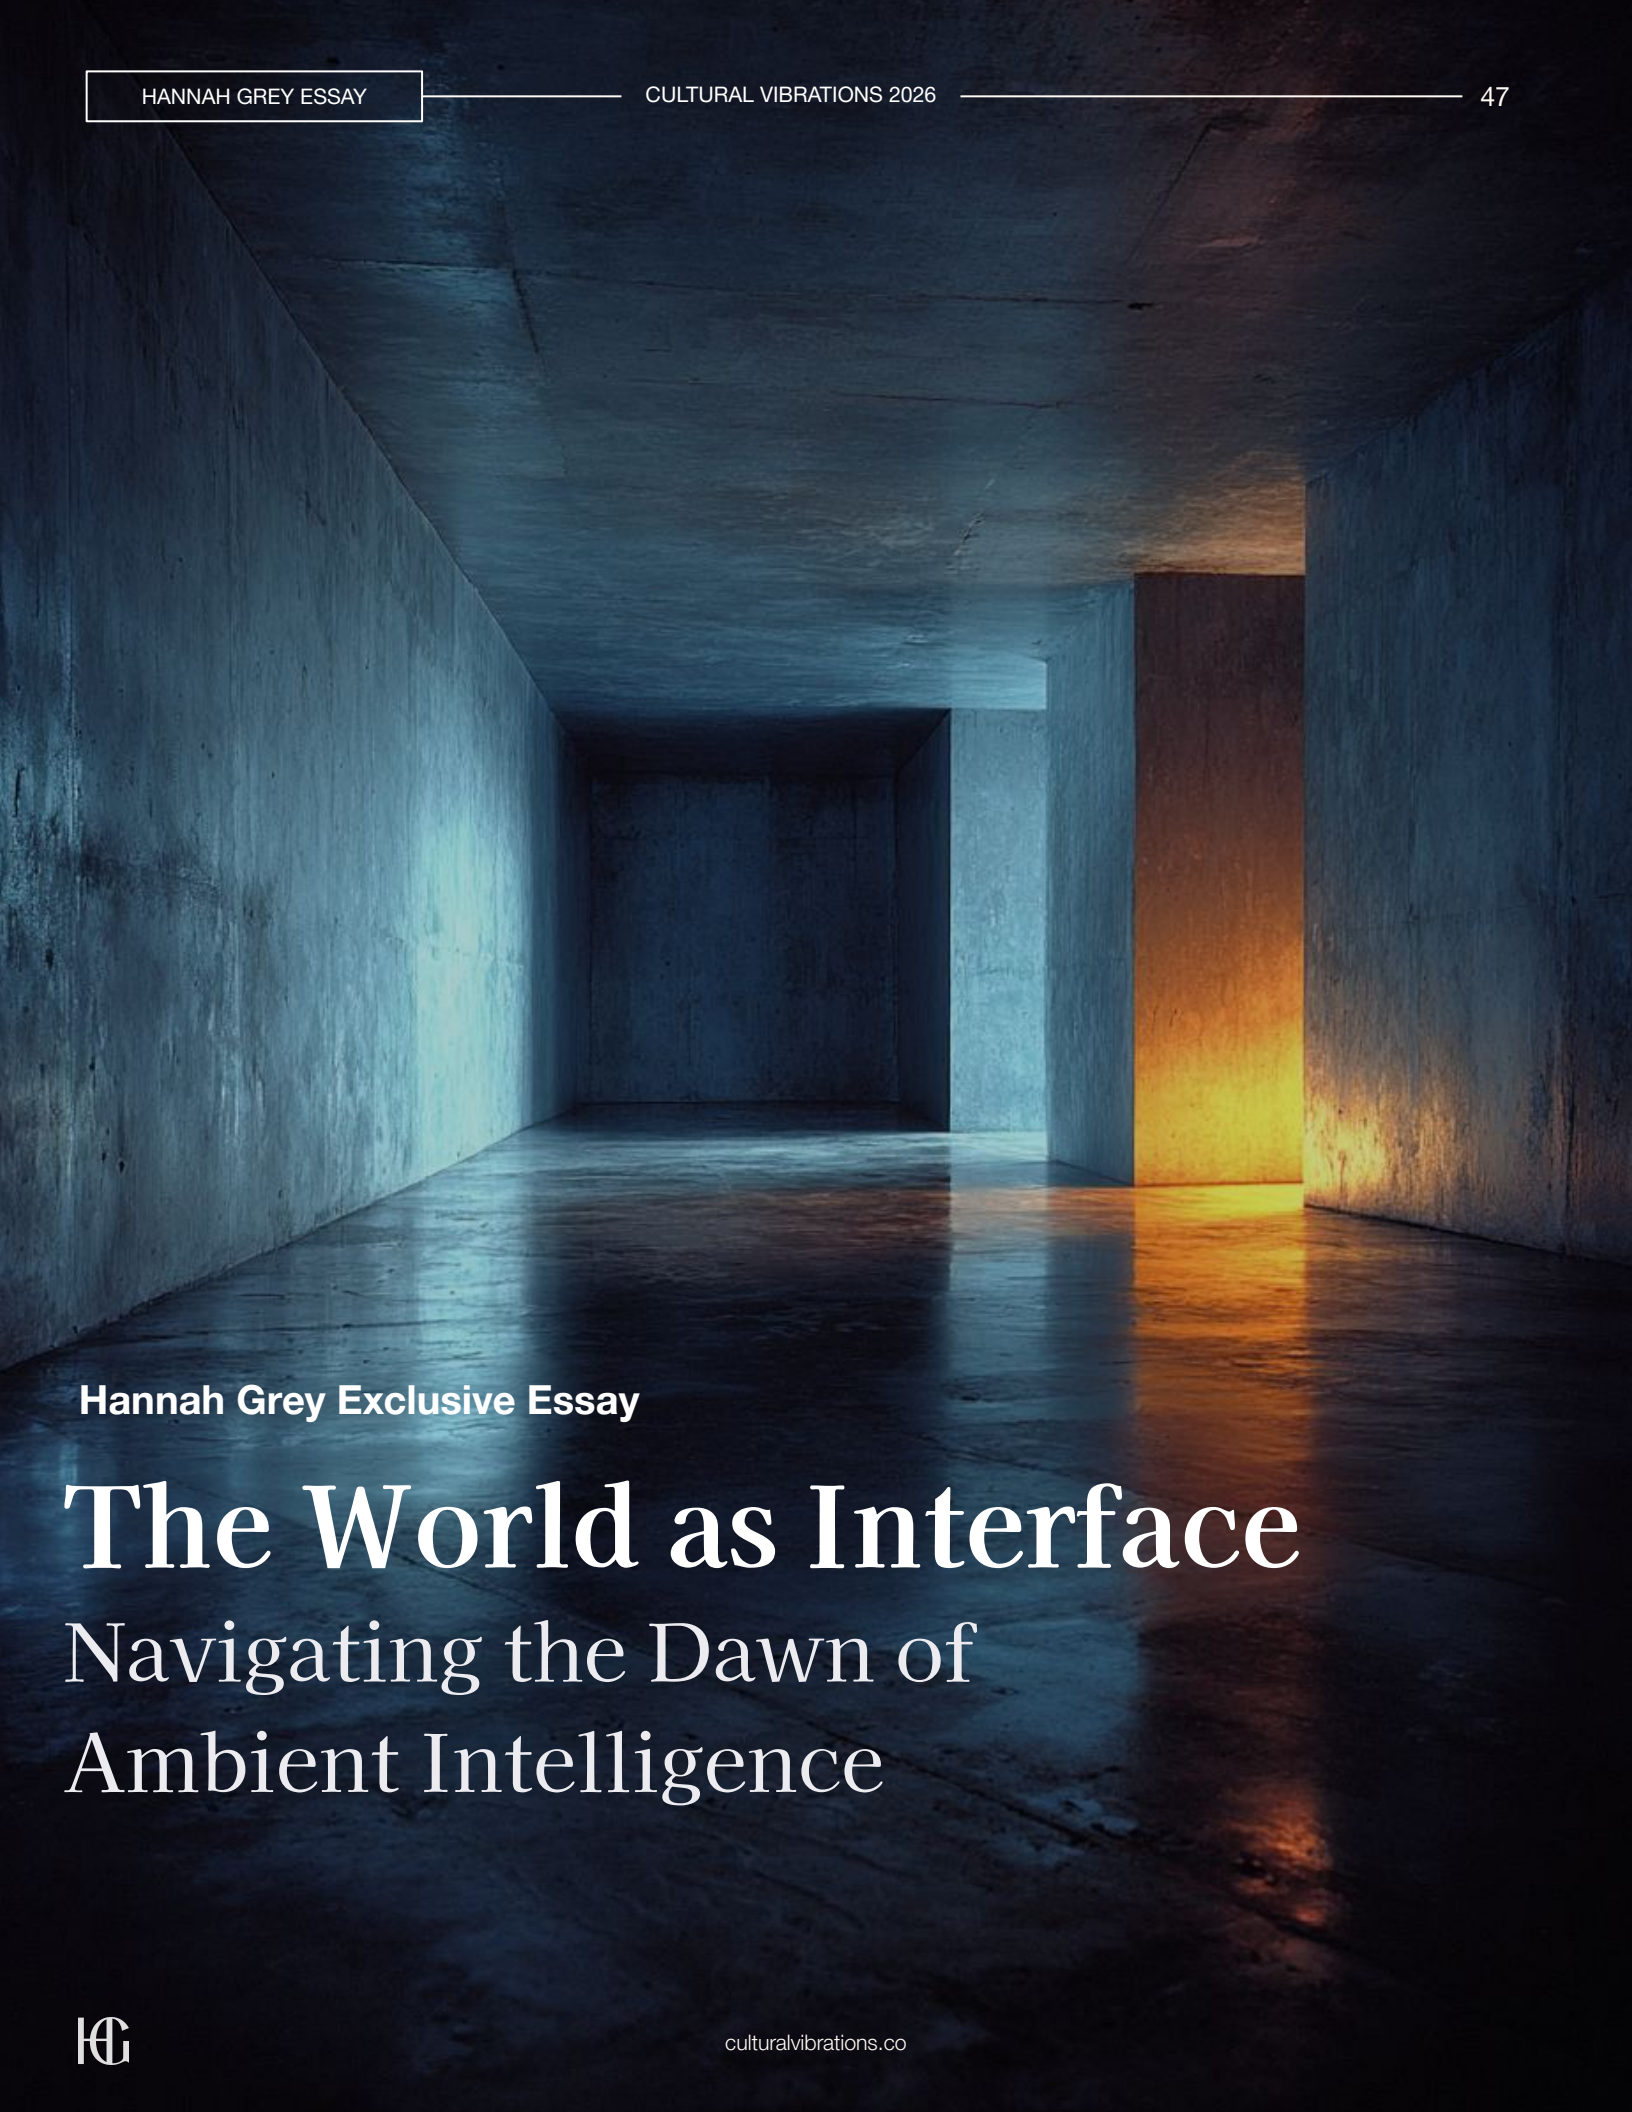

Hannah Grey Exclusive Essay

The World as Interface

Navigating the Dawn of Ambient Intelligence

The defining technological interaction of our time - the conversation with a screen - is beginning to dissolve. We are moving into a new paradigm where intelligence is not something we command, but something that surrounds us. This is the era of **Ambient Intelligence (AmI)**, a state where our environments become sensitive and responsive, anticipating our needs and shaping our behavior without demanding direct attention. This is shifting us from activing commanding technology to passively interacting, reshaping our experiences and social structures.

The allure of AmI is a world of frictionless comfort and intuitive care. It promises a home that adjusts lighting and sound to our emotional state, or a clinical setting where an ambient AI scribe transcribes a doctor's conversation, freeing them to focus entirely on the patient.

This vision of technology as a supportive, invisible partner is powerful.

Yet, this seamlessness creates a new cultural tension. The constant presence of a sensing, learning intelligence raises anxieties about autonomy and privacy, giving rise to counter-trends like "AI veganism," where individuals actively reject algorithmic assistance. A "Humanity Premium" is now being placed on verifiably human-made content and experiences as a direct response to synthetic, invisible systems.

The market is already transforming in positioning. Business models are pivoting from selling discrete devices to offering integrated, outcome-based services, not a smart speaker, but a "restorative environment." We see this in the design of "Third Nature" spaces that blend natural settings with computationally augmented technology. However, as companies build these frictionless worlds, they must contend with a societal pushback demanding transparency. Without clear "proof of human" and ethical guardrails, the magic of AmI can quickly breed deep-seated suspicion.

As this intelligence becomes more integrated, its consequences ripple outward, reshaping assumptions about work and community. When AI agents can automate complex cognitive and creative workflows, we face a potential "Post-Work Identity Crisis," forcing a re-evaluation of human value. The very concepts of presence and community are redefined when intelligence is diffused and our environments anticipate our needs, potentially eroding the space for serendipity and discovery. We will need to develop a new "AI etiquette" and new social contracts for a world where our co-worker might be an algorithm and our home is a constantly learning entity.

Ultimately, the journey into the ambient era is not merely a technological challenge; it is a cultural and philosophical one. The goal is not simply to build a smarter environment, but a more human-centered one. The most promising opportunities in this landscape will be those that understand how to integrate this pervasive intelligence in a way that augments, rather than diminishes, our agency, connection, and sense of self.

> “In a world where everybody is focused on their phones and computers, I think AI can be a bit dangerous. It disconnects people further from each other. So while it might be useful for some things, we need human-to-human interaction.”
>
> *— Lynda, Suburban OK*

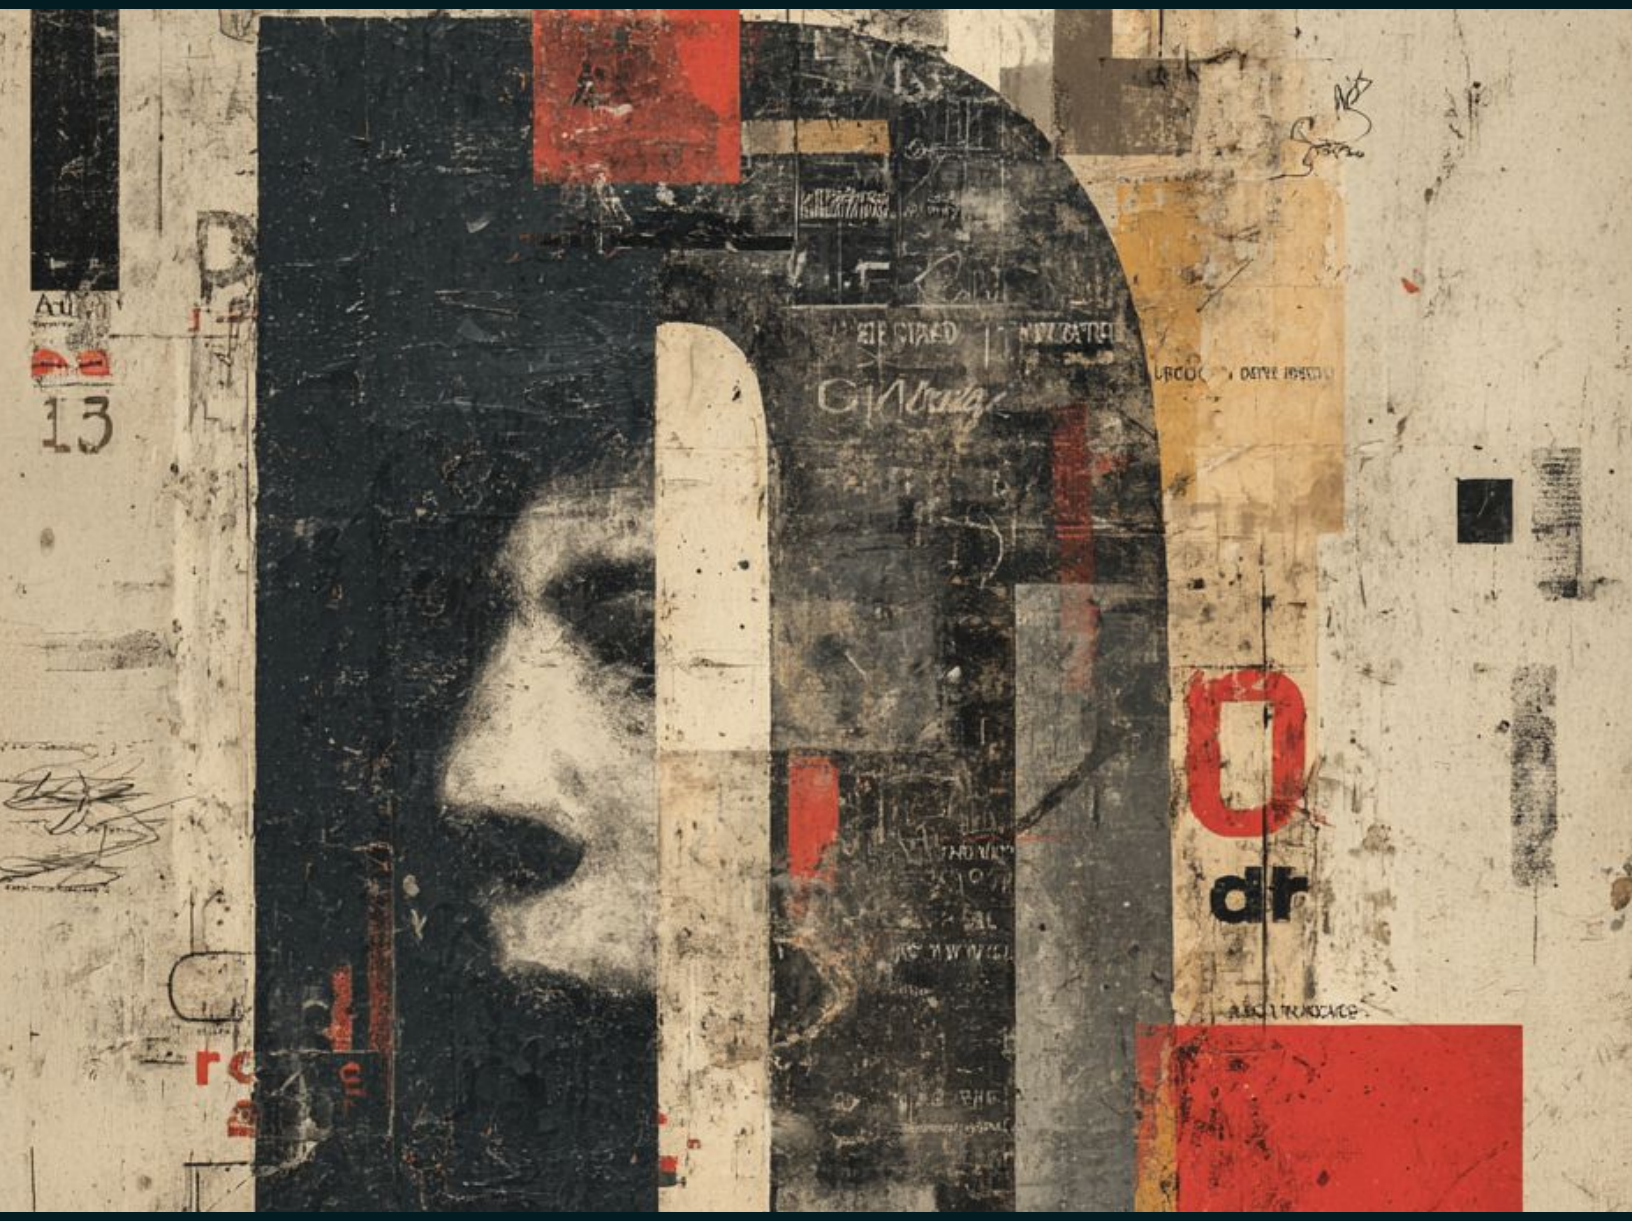

Make Language *Dangerous* Again:

How to Reanimate Creativity by Reintroducing Drama

Ben Jenkins, Co-Founder of Okay Human

Make Language Dangerous Again

How to Reanimate Creativity by Reintroducing Drama

*The following is an excerpt from **Becoming Human.** Read the full article on Substack.*

When Giorgio Moroder harnessed the Moog synthesizer to fuse funky futurism with Donna Summer's ethereal voice in 'I Feel Love' disco didn't die and musicians didn't panic, handing over the reins to synthetic music. It grabbed the tools by the electric horns and cranked out Heart of Glass, Gary Numan's Car, and Don't you Want Me. It made *more* soul and *more* funk as man and machine went to work together.

50 years later our stomach for this duet of tech and talent seems to be weaker. In the age of antagonism we play it safe - viewing every opposite with suspicion. So the latest technology is quickly tagged as an existential threat to creativity and all human agency. Yes, something big definitely is happening. But it's not just about LLMs displacing McKinsey grads, Brad Pitt and Tom Cruise fighting convincingly (despite the plagiarism), or muggles dictating websites with 'vibe coding'. It's about how we non-technical types have fallen in behind AI, meekly awaiting further instructions. It's how we've forgotten how to be boldly human. And how a larger language muddle must be averted before we start Gofundme's for our coder mates.

AI Panic

Last month the internet nearly broke after Matt Schumer posted an essay on X entitled 'Something Big is Happening.' In it he 'magnanimously' claimed that he was compelled to share the reality of AI's imminent lockdown moment - asserting that we're at the technological equivalent of February 2020 and still unaware of COVID's catastrophic implications.

COVID is an interesting analogy, but not for the reasons he states. Rather because this moment is prompting a reawakening of our humanity - just as the pandemic reminded us of the compassionate parts of us that had been sleepwalking. I spend a lot of time thinking (worrying) about this topic, ably assisted by a close friend, who works with engineers and raises the specter of humanity's demise on a daily basis.

But what about the humans that don't live in or around code? Matt Schumer is a computer engineer who spends his life coding. And it is coders who dominate the doom spreading while linguistics experts lurk in the shadows. When you live in a world of commands, UX, and GitHub deployments, language is narrow and structured - unlike the language most humans use, which is fluid, embodied, and chaotic. So a coder might overstate the abilities of AI. While I too feel the AI revolution palpably, including the existential plight of my coding brethren, I do suspect their day job may skew perspectives on the demise of non-coding humans. Therefore I'd like to widen the debate by exploring a fuller definition of 'language' because, while the engineer class has been speaking, the linguists, storytellers, and creatives silently and politely listened. It's time for them to re-enter the chat and make it messy again.

A Large Language Misunderstanding

The creative community hasn't merely listened to tech's words, it has swallowed its language and rituals hook, line and sinker. Humans have a funny habit of dressing up in the culture of that which they worship. There's even a term for it.

The 'linguistic founder effect' is where a new platform or culture births new linguistic patterns which are adopted by those wanting to share in its success.

This happened when influencers copied speech patterns of early reality stars like Paris Hilton, or further back when movie stars all cultivated transAtlantic accents like Katherine Hepburn and Elizabeth Taylor. The irony of LLMs' naming is that they only take into account the smallest part of language - syntax. Linguists know that human language is way vaster than the small category of written words in an LLM. In fact LLMs don't strictly model language at all; they are textual sequence prediction engines. Human language is embodied, social, emotional, even moral. It deploys context, tone, touch, sight, and feeling. Already this separates humans from the world of AI, but you can see why AI may not look too different to Matt and his colleagues - who work entirely with text-based code.

Language has always evolved as new ideas and technologies reshape the world. But in the last few years tech-speak has saturated our lives. As it did so it replaced raw, passionate humanity and agency with sterile, mechanical terms that leave you cold.

Language births culture

Worse than merely dehumanizing us, tech jargon often conceals meaning altogether. Did you really just inspire your team by calling it BX instead of brand strategy? And doesn't putting your beautiful, inspirational insights into a data stack cheapen it? Good language forges deeper connections and unifies communities around ideas, values - even ancestral beliefs. Bad language segregates, divides and diminishes. But all language has a remarkable knack of seeping into culture, so you'd better ensure it's feeding and not poisoning it. Today's 'technocracy' imposes a worldview suspicious of intuition and emotion, favoring clean, frictionless order. This binary worldview informs our metaphors and in turn how we think. If the metaphors of our time struggle with friction, abstract thinking, chaos, or drama then our ability to materialize the complex or philosophical will suffer too.

So language is not an inert byproduct of a culture, but an active, living participant in forging it.

And aspiring to tech's ends does not require the wearing of its clothes. Humans have better ones if only we can remember where we put them. While tech's success is powered by logic, creativity demands boldly reanimating what nature gave us, and accessing this is a job for language that viscerally moves us.

The return of the chaotic human

So let's reintegrate language that was designed for the full range of human expression. This may be tough at first as we have conditioned ourselves to subordinate our human skills to the big expensive machines. Operating these machines favored rational, safe thinkers, while the right-brain, lateral and messy 'feelers' were 'streamlined' out of the kitchen. But, AI's rapid growth brings good news for the nutjobs and the fire starters: the new machines don't need coherence or delicate handling anymore! We can go back to cultivating our chaotic, unhinged and emotional selves.

As the investor and business columnist Eric Markowitz says:

"AI has arrived at precisely the moment when we have already hollowed out so much of what makes work meaningful. I believe that the companies that survive the next era won't be the ones that move the fastest. They will be the ones that moved with purpose. The ones that kept their people. The ones that chose meaning over margin, long-term resilience over short-term extraction, humanity over efficiency."

I want to leave you with three practical steps you can take to help clean up our polluted linguistic water system.

1 - REINJECT THE LANGUAGE OF OUR RIGHT BRAINS.
For many years, we've curbed our language to fit neatly into the systems and tech of the 'machine'. Today's machines are not only comfortable with messiness but work better when you challenge them. We've coined a term for this: 'Chaos Prompting'. You don't politely stand back for AI to create, but instead assign AI the role of the 'firestarter' at the brainstorm - maintaining your own role as arbiter of taste and beauty. This ensures that you retain your agency as well as that of those around you.

2 - DESIGN PROCESSES THAT WARM THE NEURONS
Human brains must be warmed up to retrieve memories and make creative leaps that only humans can. This means your job is to design for this neuron-warming. Change your state, through play, trickery, humor, even vulnerability and emotional openness. It's also about expanding your modes of expression. We communicate through tone, gesture, imagery, silence, metaphor, facial expressions. If we limit ourselves to dry text, we limit our true human nature. When we widen the definitions of language we expand the neural pathways by which we arrive at new ideas.

3 - NORMALIZE A CULTURE OF CREATIVITY IN EVERY DEPARTMENT.

Vibe coding is about to create a whole new generation of long tail, offbeat, zany makers - just as YouTube did with video in the mid 00s. The era of inaccessible creativity is behind us, yet we've been conditioned to see any form of creation as a pursuit for others - not us. BBH etched this idea onto its wall for two decades - 'Ideas that come from anywhere'. Our job now is to condition the 'human' back into our employees, providing permission for them to tap their intuition, their instinct for play and their imagination. I found John Hegarty's latest version of this in his Rules for A Creative Life:

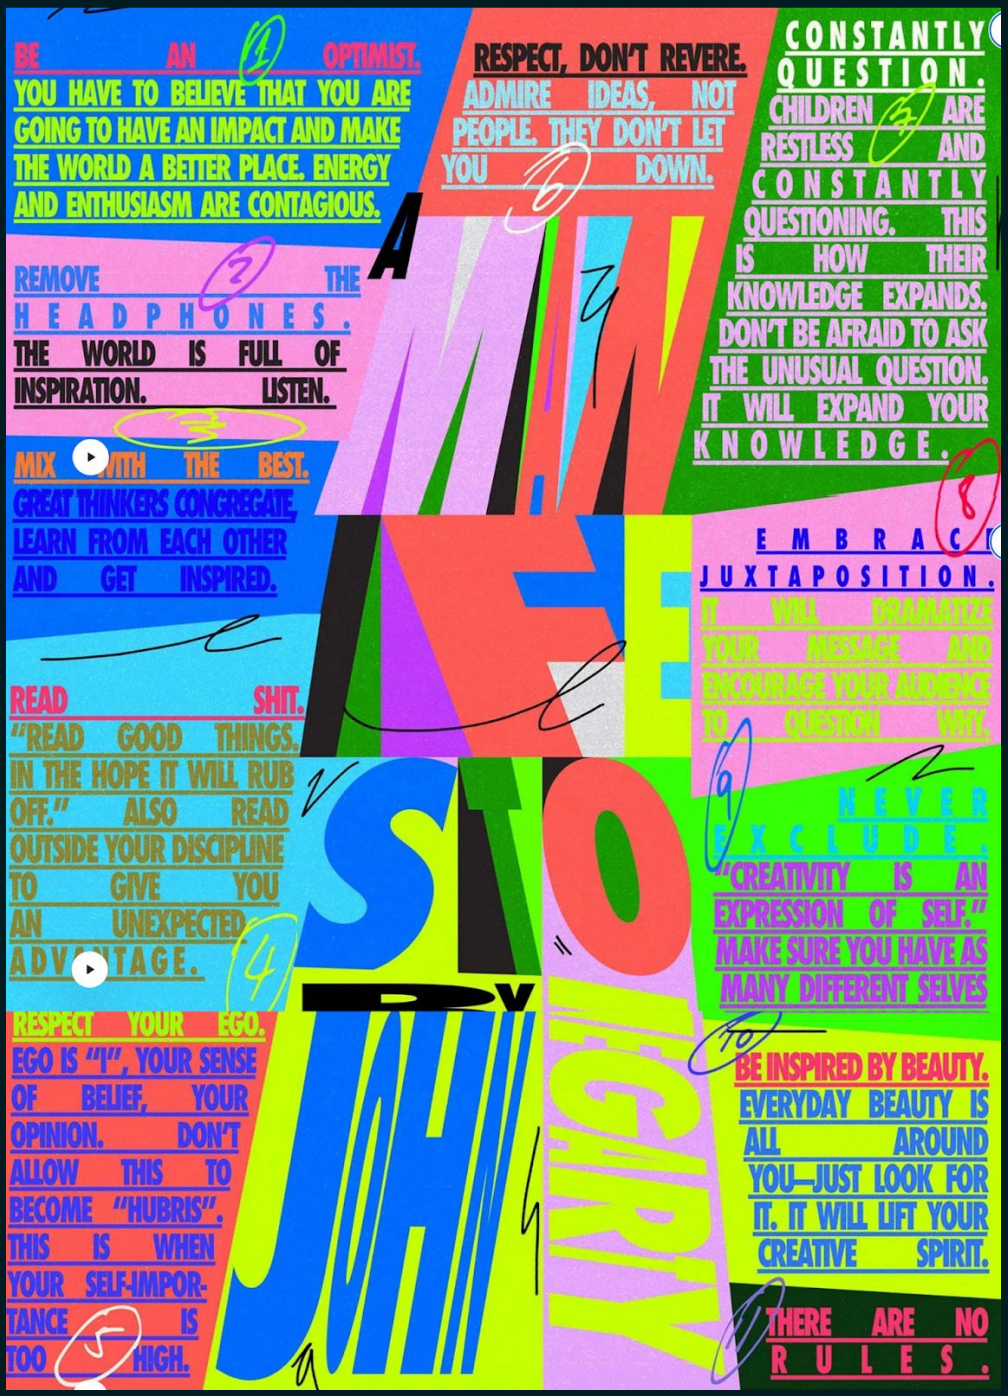

Let's reclaim these remaining human instincts from the domain of children, tinkerers and artists. Language is the first step towards this. Choose inspiring, emotionally rousing words and ditch the cold and the mechanical. If vibe coding can manifest every dream, vision and idea that we can produce linguistically, then evocative language becomes a huge differentiator in a world of prompting. The nihilism and anxiety we feel today isn't simply because we're alienated from our outputs, but from any hint of understanding of the tool that makes the product. We *became* the tools. Regaining agency means getting 'the prompted' (AI) to prompt us right back.

Thoughts of Optimism

Friends From The Hannah Grey Network

Thoughts of Optimism

Ding Zhou

VP, Ads Engineering, Coupang

"Advertising technologies have long used ML to match shoppers with relevant, high-quality, fairly priced products. What energizes me most is bringing Generative AI into how we build, work, and serve our customers. Beyond optimizing ad delivery decisions, ***AI is now giving businesses—especially small and micro merchants—the superpower to tell compelling product stories with high-quality, personalized content, once accessible only to large agencies.*** *This shift is making high-quality, personalized commerce experiences more scalable, creative, and accessible than ever. Leading that AI-native shift at Coupang is exciting."*

Nikki Scott

Head of Industry, TikTok

"There's so much to be excited about as we move towards innovation in AI (which, hot take, I believe will allow us to move offline rather than online) but I am personally most motivated by advances in the creator economy. More than ever, ***we're seeing creators that were born from niche online followings (rather than celebrity elsewhere) launch and fund successful products, start their own communities, and build commercial empires.*** *I, for one, love that power is shifting from brands back to people, who can operate with more candor, speed, and authenticity than product-led organizations, and can't wait to see how AI will allow creators to scale their empires"*

Thoughts of Optimism

Ben Lang
Head of Community, Cursor

> *Been fun meeting founders getting started recently,* ***seems like there's an entire new generation/type of company being built with all the changes in the world.***
>
> *Curious to see where this as all goes!"*

Taj Singh
Head of IT Strategy & Transformation, Coreweave

> *"AI agents operating across structured and unstructured systems are redefining IT from a service function into a coordination layer for digital labor. This unlocks entirely new operating models where* ***workflows are composed, not executed, and scale is driven by orchestration*** *rather than hiring."*

Justin Borgman
Chairman & CEO, Starburst

> *"It's the obvious answer but it is also the genuine answer: AI!* ***At Starburst, we have been helping customers get value out of their data for years, but AI makes it easier than ever before.*** *You don't have to be an expert in SQL or python, you can just ask questions of your data in plain english. Now even your summer intern can be an expert in your business on their first day!"*

Thoughts of Optimism

Jon Levy

Behavioral Scientist & NYTimes Bestselling Author

> ***What excites me most as a leader right now is seeing how much more attention people are giving to connection and meaningful relationships**, even amid all the transformation and possibility AI is creating. I find it encouraging that, no matter how advanced the technology becomes, people keep returning to the human elements as essential. By implementing AI agents, our team has increased productivity by an order of magnitude. We have built custom agents to support nearly every part of our work. Now the real challenge is deciding what we want to do next and how to think even bigger."*

Katie Klumper

Head of Marketing Transformation & Operations, GEICO

> *"The role of a marketing leader is becoming one of a 'Chief Transformation Officer,' with AI as our most powerful tool. The opportunity is to move beyond using AI for simple optimization and instead use it to architect entirely new, dynamic customer experiences in real-time. This elevates the 'art' of brand and storytelling - **it's the human insight and creative vision that will guide the immense 'science' of AI.** We're now empowered not just to interpret customer needs, but to anticipate them, creating unprecedented value and shaping the future of how businesses connect with people."*

Thoughts of Optimism

Kim Kreuzberger

Founder and CEO, Pivot Projects

"I'm incredibly optimistic about the maturation of the creator economy by 2026. We're moving past the era of vanity metrics and into a world of verifiable impact. ***What excites me most is the rise of the 'creator-as-founder' – influencers who are not just channels for brands, but are building authentic, community-driven businesses themselves.*** *The technology and data platforms emerging now will make it possible to scale these authentic relationships in a way that was previously unimaginable, leading to a more vibrant and equitable ecosystem for creators and brands alike."*

Nick Uhas

Founder & CEO, Uhas Media Co.

"I'm very optimistic about the convergence of the creator economy and AI. While a lot of people think AI will replace creators, I see it doing the opposite and supercharging them.

We're no longer stuck with bad stock images, clunky Google searches, or hours of manual transcription. We now have the most powerful creative and thinking tool in human history at our fingertips.

It really feels like we've arrived at a true "media democracy," where ***everyday creative people can produce work that rivals what used to only come from big studios****. I see the AI wave as an artist's dream, not some Terminator-style destruction machine."*

The Road Ahead

In this edition of Cultural Vibrations, we've tracked how abstract AI is becoming a concrete collaborator, remapping everything from corporate structures to economic models. What emerges is not a narrative of replacement, but a **blueprint for human elevation**. It's a guide for designing systems where **our unique abilities are not automated away, but augmented and amplified**.

As inception investors, we believe this shift from command to collaboration represents a generational opportunity. The next market-defining companies will win not by building a better model, but a fundamentally better partnership - one where human participation builds a compounding data moat, continuously improving the experience for every user. They will enable humans to orchestrate complex systems of intelligence and operate with a emotional quotient machines cannot replicate.

The Collective Coefficient is our framework for this. It asks: Where is human judgment most critical? How do we design organizations for fluid, human-AI teaming? What infrastructure is needed to ensure trust, transparency, and control?

Thank you for engaging with this exploration. We hope it prompts you to not think about the tasks AI will perform, but the new workflows it will unlock, and how our own roles will be elevated as a result.

Let's get to work designing a future where our greatest creation is a better collaborator.

— Kate, Jessica & the Hannah Grey VC Team

Ben Jenkins

Ben has spent 25 years studying the inner worlds of humans across the globe for brands including Levi's, Coca-Cola, Johnnie Walker, British Airways & Axe. His first half was spent as a brand strategist in leading ad agencies - including BBH, Droga5, and Ogilvy. Then as the co-founder of Okay Human, a new generation of ethnography that pairs the art of deep qualitative probing with technology designed for intimacy and creativity.

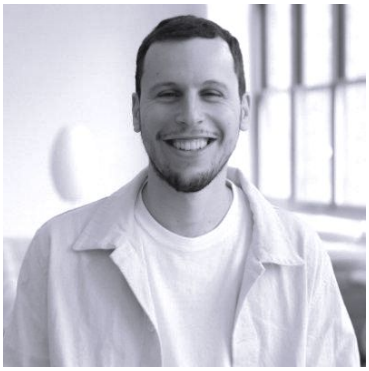

Ben Lang

Ben Lang works on Community at Cursor, previously was an early team member at Notion. He is an active angel investor and runs the Next Play newsletter for folks who are thinking about what's next.

Brian Lange

Brian Lange is the co-founder of Future Commerce, an independent research and media organization focused on commerce, technology, and culture, where he writes about commerce as a cultural and media system. His work examines how markets blur boundaries between business and culture, media and retail, and lived experience and abstraction.

Ding Zhou

Ding dedicates his career to serving consumers and empowering business owners. He is an expert in building large-scale intelligent systems and applied machine learning products, driving innovation across the advertising, e-commerce, and media industries. He held key leadership positions at global technology companies, including VP of Search & Discovery at Snapchat, VP of Ads at Pinterest, and VP of Engineering at DoorDash. Ding currently leads Ads Engineering at Coupang, building AI-driven systems for advertising and driving AI-native transformation.

Jon Levy

Jon Levy is a behavioral scientist and New York Times bestselling author exploring the art and science of human connection. A trusted advisor to leadership at Microsoft, Google, and Samsung, Jon specializes in applying behavioral research to transform C-Suite performance, culture, and the AI-enabled workplace. Over a decade ago, he founded The Influencers Dinner, a private community and behavioral experiment where leaders from diverse fields connect not on status, but on substance.

His bestselling books, *Intelligence You're Invited* and *Team Intelligence*, codify the principles of trust, belonging, and collective genius, offering frameworks for leaders navigating the future of work.

Justin Borgman

Justin Borgman is a subject-matter expert on all things big data & analytics. Prior to founding Starburst, he was Vice President & GM at Teradata (NYSE: TDC), where he was responsible for the company's portfolio of Hadoop products. Justin joined Teradata in 2014 bi the acquisition of his company Hadapt where he was co-founder & CEO. Hadapt created "SQL on Hadoop" turning Hadoop from a file system to an analytic database accessible by any BI tool. He founded Starburst in 2017, seeking to give analysts the freedom to analyze diverse data sets wherever their location, without compromising on performance. Recently, Justin was named one of the most Exceptional Entrepreneurs of 2022 by Goldman Sachs and received the 2022 BIG Awards for Business Entrepreneur Award.

Katie Klumper

Katie Klumper is a visionary leader and a trusted partner for CMOs who want to drive growth, innovation, and impact in their organizations. She is currently the Head of Marketing Transformation and Operations at GEICO, where she leads marketing strategy, investment strategy, transformation, AI and operations.Katie formerly served as CEO and Founder of Black Glass, a CMO consultancy that offers a unique subscription model and a progressive strategic framework for marketing transformation. With over 18 years of experience in marketing, advertising, technology, and finance, she has a deep understanding of the challenges and opportunities facing CMOs in a complex and dynamic world. She has been recognized by Fast Company, AdWeek, Direct Marketing News, She Runs It, and BCIT for her leadership, creativity, innovation, and growth.

Kevin L. Walker

Kevin L. Walker is a foresight strategist who translates cultural signals into strategic decisions. As the former Head of Cultural and Consumer Insights at Creative Artists Agency Brand Consulting, he built the cultural intelligence practice that informed brand investments for T-Mobile, Chase Bank, Modelo, and Adobe. He is the founder of Experience Intelligence Lab and co-producer of Bring Your Own Vinyl a Dallas institution that anticipated the global vinyl resurgence. Today, Kevin advises organizations navigating cultural disruption as a fractional strategic officer, retained intelligence partner, and keynote speaker. He has devoted his consultancy around tracking the cultural signals shaping the next wave of human-centered economic opportunity from community wealth models, to digital sovereignty, to the emerging re-analog economy.Kevin advises organizations on cultural strategy, foresight frameworks, and human-centered market opportunity. He holds an MBA from SMU and a BA from Morehouse College.

Kimberly Kreuzberger

After 14 years working on prestige fashion publications including Lucky, Vogue and InStyle, Kimberly joined goop, an early stage contextual commerce start-up. During her 4.5 year tenure, goop raised over $100M+, experienced hyper growth from 16 to 260+ employees and grew to be valued at over $300M. As Chief Revenue Officer, Kimberly doubled publishing revenue year-over-year by creating a breadth of media offerings including podcasts, wellness summits, retail pop-up activations, licensing collaborations, and e-commerce product integrations. In 2019, Kimberly founded Pivot Projects, a full-service business partner for celebrity talent. Pivot builds impactful storytelling platforms by offering marketing, e-commerce and operating services to new and existing talent led brands. Clients include Chrissy Teigen (Cravings), Julianne Hough (KINRGY), Channing Tatum (Magic Mike + Sparkella), Winnie Harlow (Cay Skin), and Molly Sims (YSE Beauty).

Leopold Pluto Hermansson & Nikita Khomich

Leopold Pluto Hermansson and Nikita Khomich are both Founding Members of the Technical Staff at Tzafon. Focused on research and infrastructure development, both Leopold and Nikita have experience across applied quantitative research, and machine learning, and software development.

Matt Klein

Matt Klein is a cultural theorist, strategist, Webby-winning writer, global speaker and “Top Voice in Market Research.” With a decade of experience in trend forecasting, innovation, and strategic communications, he is a trusted advisor to organizations seeking to make sense of cultural change. Klein partners with brands, the UN, investors, philanthropies, and the press to provide provocative, alternative explanations of the zeitgeist and author future- proofing strategies. His publication ZINE studies our technology, media and unspoken trends defining our future. He’s a source for The NYT, WSJ, VICE & DAZED and speaker at SXSW, MoMA, Fast Company and House of Beautiful Business. He is currently the Head of Global Foresight at Reddit and Resident Futurist with Hannah Grey VC.

Nick Uhas

Nick Uhas is the Founder & CEO of Uhas Media Co., a social media strategy and production company built for modern businesses and entrepreneurs. Leveraging more than 14 years of hands-on experience in digital media, Nick and his team help brands turn social media into a true growth engine - combining research, scripting, filming, editing, posting, and performance analysis into a single, scalable system. As a science-focused host and producer, Nick has created, produced, and hosted content forDreamWorksTV, AwesomenessTV, the Discovery Channel, and Netflix. He has over 7M followers on TikTok, 1M subscribers on YouTube, and 25M subscribers on Snapchat.

Nikki Scott

Nikki Scott is the Head of Industry at TikTok where she leads financial services partnerships. Prior to TikTok, she held leadership roles at influencer marketing agency Social Studies, Foursquare and Evite. Nikki also teaches yoga and megaformer Pilates and lives in Los Angeles with her husband, 2 sons, and 2 dogs.

Okay Human

Okay Human is a human-centered research and design consultancy that helps organizations decode the "why" behind human behavior. Their methodology translates ethnographic research, qualitative interviews, and behavioral analysis into actionable strategy. By building foundational frameworks from deep human insight, they provide teams with the clarity and conviction to design more meaningful products, services, and systems.

Scott Anthony

Scott D. Anthony is a Clinical Professor of Strategy at the Tuck School of Business at Dartmouth College, where he focuses on disruptive change. Scott's new book - his ninth - is *Epic Disruptions: 11 Innovations that Shaped Our Modern World*. Scott previously spent more than 20 years at Innosight, a growth strategy consultancy founded by Harvard Business School Professor Clayton Christensen. Thinkers50 named him the world's fifth most influential thinker in 2025 and named him the world's leading innovation thinker in 2017.

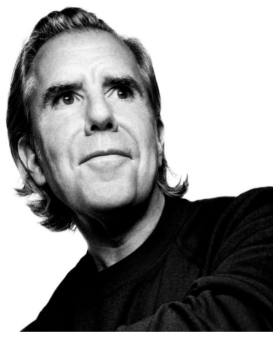

Scott Dadich

Scott Dadich is the creator and executive producer of *Abstract: The Art of Design*, the Emmy-nominated, International Documentary Award-winning Netflix series that profiled some of the world's most innovative creators across two seasons and 14 episodes. The former editor-in-chief of WIRED and co-founder of Godfrey Dadich Partners, he has spent three decades working at the intersection of narrative, technology, and design—advising leadership at Apple, Netflix, TED, Nike, Microsoft, and the Obama Foundation, among others. He is a recipient of the National Design Award for Communication Design from the Cooper Hewitt, Smithsonian Design Museum. His recent film work includes *American Family,* a 2024 film for Vice President Kamala Harris, and *The Gift of Time*, a documentary short commissioned by the chairman and CEO of SEIKO as a meditation on time, craft, and human creativity—a gift to the people of Japan. He is now a senior advisor to several ventures—particularly those operating at the intersection of narrative, AI, and creativity—and serves on the boards of The People's Portfolio, advancing human rights through narrative storytelling, and the Bay Area Host Committee, helping steward the region's ongoing civic momentum following Super Bowl LX and into the 2026 World Cup. As ever, his work lives between disciplines: design and diplomacy, product and prose, systems and stories.

Taj Singh

Taj is a forward-thinking CIO and technology leader focused on helping organizations harness AI, modern platforms, and enterprise infrastructure to drive meaningful business transformation. His work centers on translating emerging technologies such as AI and GenAI into scalable platforms that improve customer and employee experiences while enabling faster innovation across the enterprise. At Coreweave he leads the technology and product organizations responsible for building secure, resilient, and compliant systems across enterprise infrastructure, fintech services, IT and HR platforms, compliance and SOX environments, cybersecurity, as well as supply chain, construction, data center operations, and logistics.

Hannah Grey

Hannah Grey is a thesis-led venture capital advisory services firm investing at inception in high velocity founders building outlier companies.

Kate Beardsley, Founding Partner

Jessica Peltz Zatulove, Founding Partner

Gillian Katz, Investor

Pawan Murthy, GTM Lead

Yujin Chung, Customer Intelligence Lead

Haley Lopez, Chief of Staff

Matt Klein, Futurist in Residence

www.hannahgrey.com
www.culturalvibrations.co

For any questions, feedback, or press inquiries on **Cultural Vibrations 2026**, reach out to us at:

vibrations@hannahgrey.com

Design — Bossa Design Studio