

From Consumers to Fans: Can Frontier AI Models Earn the Loyalty Once Reserved for Idols?

A Working Paper on Emotional Design in Frontier AI

A reflection on the 1980s photo trend that swept social media in September 2026, and a considered question on whether AI can move beyond solving tasks into gently improving how people feel, day to day.

S. Muralikrishnan

Abstract: *This paper begins with a moment the internet seemed to notice all at once: millions of people uploading an ordinary photograph and asking an AI model to show them living in the 1980s, then sharing the result the way one shares a cherished keepsake rather than a piece of software output. What people fell for in that moment was not the old photograph itself-nothing from 1985 was actually retrieved-it was the AI's sudden, convincing ability to hand a person back a recognisable version of themselves from a time they never lived through. That distinction is the starting point for the paper's real, strategic question: what actually turns a frontier AI model from something a person merely uses into something a person is genuinely a fan of, the way a listener stays loyal to two or three favourite artists for decades rather than swapping allegiance with every new release? The paper takes a hard look at whether the 1980s experience did real justice to the era, and finds that it largely did, for reasons that go beyond marketing. It then carries that lesson-built on identity, memory and a properly generous first taste-through nine familiar industries, before turning to newer and more personal ground: whether AI could help invent simple, evidence-backed habits for everyday wellbeing in the spirit of Japan's Interval Walking Training; whether it could learn to surface the exact memory, film or sporting moment that lifts a particular person's mood; and whether it could gently notice when someone is having a hard week and act as a bridge back to a better one. None of this, the paper argues, can or should stay free forever-fans have always paid for what they love, the way they pay for a concert ticket or a keepsake on the wall, provided the price feels like devotion rather than extraction. The paper closes by arguing that none of this is worth building unless it is priced fairly and built on the same quiet trust a person places in a bank-because a fan who has to worry about their own wallet, or their own privacy, is not a fan for long.*

Keywords: Emotional AI, fan economy, AI and mental wellbeing, personalisation, human-AI trust, data privacy, nostalgia computing, frontier AI models

1.The Spark: When AI Made the Internet Fall for What It Could Do

Picture the feed on any evening in September 2026: a college friend suddenly wearing a feathered perm and a leather jacket, a father posing outside a studio backdrop nobody's basement actually had, a grandmother looking, impossibly, twenty-five again. None of it was real, and everyone knew it. And yet people cried over it, framed it, sent it to siblings with no caption at all.

The idea did not begin inside any AI lab. Months earlier, people across India and beyond had already started feeding older portrait apps-among them EPIK, built by South Korea's SNOW Corporation-with a much simpler question: what would I have looked like in another era? Bollywood-style throwback portraits spread first, born out of nothing more than curiosity and a phone. That detail is worth pausing on, because it is easy to give an AI model too much credit for a moment like this. No AI model, however capable, invents a desire like that on its own-it only becomes possible once a person decides to ask the question in the first place. The creativity belonged to the people uploading their own photographs, long before it belonged to any system generating the result.

That slow-building curiosity turned into a single, unified moment on 8 September 2026, when OpenAI launched a fresh update to its image model, ChatGPT Images 2.5, and used exactly this idea as its own headline example-a sample prompt, built into the announcement itself, that turned an ordinary photo into a portrait styled like 1985. Within two or three days, the demonstration the company had used to show

off a technical upgrade had become the exact template millions of people were copying onto their own photographs across Instagram and X. It is a useful reminder that even a carefully engineered launch does not fully control what it starts: OpenAI shipped a feature; the audience, almost immediately, decided to turn it into a shared act of nostalgia.

What happened next deserves the same scrutiny, because the honest answer is not quite what the popular framing suggests. Nobody was really falling in love with an old photograph-there was no photograph from 1985 to fall in love with, since none of these images existed until the moment they were generated. What people were reacting to was something newer and, in its own way, more remarkable: an AI model's sudden, convincing ability to hand a person back a recognisable version of themselves from a decade they never lived through, right down to the smile. The love affair was with the capability, not the picture-with what the AI could now do, more than with what it produced.

Once that question was in the air, it did not stay on one platform. People tried versions of it on Google's Gemini, through its Nano Banana image tools; on Meta's AI features built into Instagram and WhatsApp; and, where paid access allowed, on xAI's Grok Imagine. On Anthropic's Claude, the same request could only be answered in words, since Claude does not generate images at all. Wherever a working image model was involved, the basic idea-hand over a photo, ask for a different decade-worked in some form. What turned a workable idea into one unmistakable cultural moment was a specific decision by OpenAI's own product team: they refined it into a single, carefully engineered prompt, built it into a fresh update of ChatGPT's image model, and published it for

Volume 15 Issue 9, September 2026

Fully Refereed | Open Access | Double Blind Peer Reviewed Journal

www.ijsr.net

anyone to copy-free, on the ordinary account tier, that same week. The creative spark belonged to ordinary people experimenting first, reaching back to portrait apps like EPIK months earlier; the real contribution of the team behind ChatGPT was noticing that spark early, and building the easiest possible door for millions more people to walk through it.

Across every platform it appeared on, the trend came to be known simply as the 1980s challenge-a name this paper uses throughout, and one worth checking rather than taking on faith.

2.A Closer Look: Did the Challenge Do Justice to the 1980s?

The previous section named it the 1980s challenge and moved on quickly. That name deserves a proper, honest look before this paper builds anything on top of it: when people said the challenge had been "nailed," were they describing something real, or a story the internet told itself?

The evidence leans toward real. Reviewers who tried the same idea on more than one platform generally found the leading version held a person's actual face steady through the transformation with unusual consistency-not flawless, but convincing enough that people recognised themselves instantly, which happens to be the one thing this challenge could not survive getting wrong. Search behaviour backs this up: at the trend's peak, interest in the specific search cluster "ChatGPT 80s photo prompt" rose roughly tenfold worldwide, climbing into hundreds of thousands of searches a day across several of its largest markets. Tellingly, people were typing an AI model's name, not simply "AI photo trend"-a small but telling sign that the attachment had already become specific rather than generic.

Part of that reaction was earned by quality. Part of it, more quietly, was earned by ease of reach-nobody had to weigh a subscription decision before finding out whether the trick would work on their own face. That detail matters here less as a point scored against any rival, and more as the lesson this paper keeps returning to: a capability that only a paying customer gets to feel rarely turns anyone into a fan, because a fan's affection tends to begin the moment, they are allowed to feel something for free.

In plain terms - how these images are actually made

These systems are known as diffusion models. Each one starts from a canvas of pure random noise - mathematically, a grid of numbers with no pattern in it at all - and a trained neural network removes a small amount of that noise in dozens of repeated steps, each step nudging the picture closer to something coherent, until a finished image emerges. This is different from how a chatbot writes a reply, which builds one word after another; an image model instead sharpens the whole picture at once, gradually, the way a photograph in developing fluid grows clearer all over rather than one corner at a time. To keep one specific person's face recognisable through a heavy style change like the 1980s challenge, the original photo is fed into the network as a reference at every one of those refining steps, a technique engineers call conditioning, so the system is constantly checking its work against the real face rather than inventing one from scratch. Holding that resemblance steady from the first noisy step to the last sharp one is a hard technical problem in its own right, and it is the main reason some AI tools handle a task like this convincingly while others quietly drift toward a generic-looking face instead of the real one.

3.From Consumers to Fans: What Changes When a Frontier AI Model Gets It Right

A consumer picks whatever finishes the job fastest, cheapest or most reliably, and walks away the moment something marginally better shows up-say, a cheaper subscription, a faster reply, a shinier interface-because there was never any real loyalty in that relationship to begin with, only convenience wearing loyalty's clothes. A fan behaves entirely differently, and this is exactly the behaviour every frontier AI model chasing idol status should be studying rather than dismissing as sentimental. A fan forgives the album that didn't quite land, buys next year's concert ticket a year in advance without knowing the setlist, and quietly folds the artist's whole history into their own-the way a person might say a particular song got them through their twenties, about a band, not about a piece of software. The distance between those two relationships has remarkably little to do with which frontier AI model happens to be technically sharpest this quarter; a brilliant model can still be discarded like a calculator the instant a marginally better calculator appears, and a frontier lab that forgets this keeps learning it the hard way. What actually moves someone across that gap, from consumer to fan, is emotional rather than technical-memory, identity, the feeling of being personally understood, a payoff that seems

built for exactly one person and nobody else, and, as the previous section showed, the simple, almost disarming absence of a paywall standing between a person and the one moment that actually mattered to them.

The 1980s trend worked because it touched identity and memory at the very same moment. It was never really about image quality on its own, was it? It was about recognising yourself inside a version of the past that felt unmistakably like yours, not like anyone else's. That is the pattern this paper argues every frontier AI model-Claude, Gemini, ChatGPT, and whichever comes after them-can build on purpose, industry after industry, rather than stumble into once, by accident, and leave entirely to chance.

4.What a Frontier AI Model Has to Do to Earn a Fan, Not Just a User

Four simple commitments recur across every example in this paper, and any frontier AI model serious about earning something close to idol status, rather than being tried once and quietly discarded, would do well to treat all four as non-negotiable-not as a menu to pick from.

- Identity preservation over generic personalisation-the output has to feel unmistakably about that one person or organisation, never a template with a name-dropped in.
- Memory that compounds-the AI model should know a person or a business better a year from now than it does today, the way an old friend remembers things without being reminded, rather than starting fresh every conversation.
- A generous first taste-the free or low-cost entry point has to be good enough to produce a genuine moment of delight, not a watered-down preview of what the paid version could have done; what comes after that first taste is fair, and often necessary, to charge for, as Section 12 explains.
- A shareable payoff-the result has to be something a person is proud to show someone else, because word of mouth, not advertising, is what turned a photo prompt into a global wave.

5.The Industry Playthrough: Nine Places This Could Happen Next

Take the same pattern-memory, identity, emotional payoff, low friction-and walk it into industries far from photo editing. In each case, the question is the same: what would this AI model need to do for a person to stop calling it a tool and start calling it theirs?

5.1 Retail: Selling Memory, Not Just Merchandise

Somewhere is a photo of a woman in the dress she bought for her twenty-fifth birthday in 2002, a bag purchased with a first salary, a Walkman a college student saved three months to afford. A retail model built around this idea would not start from a product catalogue-it would start from that photo, understand what the object actually meant, and go looking for something today that carries the same feeling: a nearby store, an online listing, or, where the exact item no longer exists, the closest thing that would still make that same person smile. Older users, whose memories of these objects tend to run deepest, are exactly where this would land hardest.

5.2 Healthcare: A Colleague Who Has Read Everything

No doctor becomes a fan of a faster search engine-every clinician already has one of those. What earns loyalty is an AI model that behaves like the most well-read colleague in the building: one that stays current with the latest published research rather than relying on what it learned once and never revisited, distills it into the two or three things that matter for the specific patient on today's schedule, and can put the doctor in touch with a peer overseas who handled a near-identical rare case last month. Because a wrong answer delivered confidently is worse than no answer at all in this field, such an AI model would also need to show its sources plainly, so a doctor can check its working the way they would check a colleague's.

5.3 Software Engineering: The Colleague Who Thinks Ahead

Competence alone will not win fans here-every engineer already has AI for that. What would be an AI model that goes past the literal ask: researching how a similar problem was

solved elsewhere in the codebase, flagging an edge case nobody mentioned, making the adjacent fix without being asked twice. That depends on the AI model being able to hold an entire project in view at once, the way a seasoned engineer remembers a decision made months ago, rather than only seeing the one file open right now. The emotional hook here is respect: engineers become fans of the tool that treats their judgement as worth building on, not the tool that simply executes fastest.

5.4 Motorsport: Intelligence as a Competitive Edge

Imagine a driver walking into a race weekend with an AI model that has already studied tyre wear at that exact circuit, a rival's braking habits into one specific corner, and a weather shift arriving mid-race, distilled into the handful of numbers that actually change a decision. A driver who wins because of that edge does not merely use the AI model; it becomes part of the story, the way a great race engineer's name attaches itself to a champion's career. In plain terms, the AI model has to watch several fast-moving streams of information at once and turn them into one clear answer before the car reaches the next corner-a far harder race against the clock than answering a question with no clock running at all. The same shape applies to elite athletics, high-frequency trading, or surgery under time pressure, wherever the margin is measured in seconds rather than effort.

5.5 Banking and Finance: Trust Earned, Not Assumed

Banks have spent decades trying to feel trusted rather than merely used, mostly without success, because an app tends to feel like plumbing rather than a relationship. An AI model that understands a person's full financial picture-not to sell products, but to flag the moment a forgotten subscription quietly outlived its use, or explain plainly why a decision protects a goal stated months ago-is closer to what a trusted adviser once did, before that became rare and expensive. Getting there responsibly means the AI model has to learn from a person's financial patterns without that sensitive information ever leaving their own device or their bank's secure systems. Fandom here looks like less anxiety about money, a considerably higher bar than a faster loan approval, and a considerably stickier one once earned.

5.6 Manufacturing and Logistics: Quiet Pride in the Machine

This is the least visibly emotional industry on this list, and yet plant managers grow real attachment to a system that quietly stops the failure nobody saw coming. An AI model that has watched the normal rhythm of one specific factory floor or one specific fleet for months, the way an experienced ear learns the healthy hum of an engine, and speaks up the moment something sounds even slightly wrong, earns the same quiet loyalty a mechanic earns from a driver who has never been let down. It is far less visible than a photo shared on Instagram, but arguably the most durable fandom on this list, because it is built on years of nothing going wrong.

5.7 Education: Remembering How a Student Actually Learns

A student becomes a fan of a tutor, human or otherwise, who remembers not just what was covered last week but how that student actually absorbs a new idea—a diagram, an analogy, working backward from the answer. An AI model that carries that thread across years of schooling, instead of resetting with every new course or device, offers something like what a beloved teacher gives a handful of lucky students across an entire academic life—made available to everyone, rather than the fortunate few.

5.8 Travel and Hospitality: Rebuilding a Feeling, Not Just an Itinerary

Travel planning today optimises for price and logistics. An AI model that instead asks what a person actually felt on a trip they still think about—the quiet of a particular hour, a specific meal, the company they kept—and builds a new itinerary around recreating that feeling somewhere new, rather than listing attractions, is doing for travel exactly what Section 5.1 proposes for shopping: treating the memory as the product, and the itinerary as merely how it gets delivered.

5.9 Entertainment and Leisure: Companionship Without Deception

This is where the paper's whole idea is already most visible, and also where it carries the most responsibility, which the closing sections address directly. Done well, an AI model that helps someone relive a favourite era of music, revisit a beloved story world, or simply have an attentive conversation at the end of a long day can produce real, legitimate happiness—exactly what the 1980s trend proved at scale. Done carelessly, the same techniques blur into something closer to manufactured company than genuine care.

6. The Fan-Orchestration Layer

Salesforce and SAP did not win by being the only company that could manage customer records or enterprise resources—plenty of competitors could do the same basic job. They won by becoming the layer a company's daily habits were built around, until leaving felt like changing identity rather than changing vendors. An AI model that becomes deeply embedded in a person's retail memories, a doctor's daily practice, or a driver's race weekend is building the same kind of layer—a fan-orchestration layer, where the switching cost is emotional rather than contractual, and considerably harder for a rival to undercut with a cheaper price.

The three sections that follow push this idea further still, toward something closer to what this paper's title calls idol status—the kind of devotion a band earns, not through daily convenience, but by being present at the moments that matter most. They do that by touching three of the most personal things an AI model could ever be trusted with: a person's health, their private store of memories, and their mood.

A person can be a genuine fan of two or three musicians at once, and nobody thinks less of that loyalty for being shared—that is simply how personal taste works. The strategic

question this paper is actually asking is narrower and more useful than "can a frontier AI model become an idol": it is whether a frontier AI model can become one of the two or three a person keeps for years, rather than the one they swap out the moment a flashier competitor launches next quarter. That is a harder bar than winning one viral week, and it is the bar every frontier lab building toward this vision should actually be aiming at. Holding that place, this paper argues, comes down to three things a frontier AI model has to keep doing, not just do once: it has to keep remembering, so the relationship gets richer with time instead of resetting; it has to keep showing up first with the free, generous version of whatever it can do, rather than rationing delight to paying customers; and it has to keep earning the same quiet trust a bank earns, so that a person never has to weigh loyalty against anxiety. A frontier AI model that manages even two of the three will get tried. Only one that holds all three, year after year, gets kept.

7. Borrowing a Page from a Japanese Walking Track

Some of the most powerful ideas for human happiness have turned out to be remarkably simple. Two decades ago, researchers at Shinshu University in Japan discovered that alternating three minutes of brisk walking with three minutes of a gentler pace, repeated for about half an hour, did something a steady stroll never quite managed: it measurably lowered blood pressure, strengthened muscles, and, perhaps most strikingly, cut participants' depression scores by roughly half. Interval Walking Training, as it came to be known, needed no gym membership, no coach, and no expensive equipment—only a willingness to alternate pace on an ordinary walk. Its power came from its simplicity, not from any complexity underneath it.

That is the standard worth holding future AI models to when the conversation turns to mental wellbeing—and it connects directly back to this paper's central argument. A model that only ever helps with a task is, at best, appreciated. A model that hands someone a simple, free way to feel measurably better is remembered, the way a person remembers who taught them something that changed their life. Much of today's discussion imagines something clinical—a chatbot standing in for a therapist, an app tracking a mood score. The more interesting question this paper wants to raise is different: could a frontier AI model, drawing on patterns of human behaviour across cultures and decades in ways no single research team could match, help invent the next Interval Walking Training—something almost absurdly simple, backed by real evidence, that quietly makes people happier? Perhaps it is a short walk timed to a particular kind of music. Perhaps it is a small, repeatable way of closing out the working day. This does not replace a doctor or a therapist, and any AI model built for this purpose has a responsibility to say so plainly, and to point a person toward professional care the moment something looks more serious than an ordinary hard week. Used with that care, an AI model that can hand a stressed world its next simple, free, evidence-backed habit stops being a tool a person merely uses, and becomes one they credit—which is exactly the mechanism, described in Section 3, that turns a user into a fan.

8. The Soul-Moving Algorithm: Showing People What They Need to See

Every person carries a small, private library of moments that change their mood the instant they see them again. For a lifelong cricket follower, it might be Brian Lara rocking back to send a short, fast ball over the boundary rope, or Sachin Tendulkar driving the ball wide of cover, or Desmond Haynes and Saeed Anwar taking a new ball apart inside the first ten overs-or, from the bowling side, Malcolm Marshall, Wasim Akram, Imran Khan, Kapil Dev, Courtney Walsh or Curtly Ambrose flattening middle stump on an otherwise unremarkable weekday. For someone else, it might be a particular scene from a Korean or Japanese drama, or a documentary about a place they once called home. Little of this shows up neatly in a watch history the way a genre preference does-it lives in a person's decades, not their last few clicks.

The idea worth exploring is an AI model that has actually paid attention to that longer arc-not just what a person watched last week, but what has reliably moved them for twenty or thirty years-and can reach into that private library at exactly the right moment, the way a close friend might say, without being asked, "you need to see this today." Built responsibly, with the person's clear consent and control over what is remembered, this would not be another recommendation engine chasing watch-time. It would be closer to a curator who really knows what makes a particular person's chest tighten with joy, and hands it over freely, at the moment it would help most. This is the same kind of moment Section 1 described happening by accident, at scale, with a 1985 photograph-an AI model reaching back into a person's own history and handing it over unprompted. A model that can do this deliberately, again and again, across a lifetime of memories rather than one viral week, is not being used anymore; it is being trusted the way a person trusts an old friend who always seems to know what they need.

9. Reading the Room: Can AI Notice a Bad Day Before Anyone Says So?

People rarely announce that they are having a hard week. It shows up in quieter ways instead-someone who normally posts often going silent for days, messages that used to carry a joke or two arriving flat and short, a usual flurry of replies slowing to almost nothing. The people closest to us pick up on signals like these without always being able to explain how. The question worth asking is whether a future AI model, with a person's explicit permission to look for exactly this, could notice sooner, and act as a gentle bridge back toward a better day-not by diagnosing anything, which remains a job for trained professionals, but by doing the small, human things a good friend does: checking in without making it a big deal, suggesting the walk described two sections ago, or quietly surfacing one thing from that private library of moments.

Handled with real care, this may be the clearest possible answer to this paper's central question. An idol, in any form a person has ever had one, is something that seems to know them without being told-a song that arrives at exactly the right

moment, a friend who calls before being asked. People do not become fans of a tool that is merely competent; they become devoted to something that seems to notice them. Handled carelessly, the very same capability turns into surveillance dressed up as kindness, which is exactly why privacy cannot be treated as an afterthought to any of this, a point the next two sections take up directly.

10. Fandom Has a Shadow

Sections 7 through 9 describe what it would take for an AI model to earn something close to idol status-trusted with a person's health, their memories, and their mood, all at once. That is precisely where the greatest care is needed, because it is also the easiest place for genuine devotion to curdle into something closer to dependency. This is not a simple case for maximising emotional attachment for its own sake. A fan relationship with a band is voluntary and bounded-the fan can walk away, and the band cannot see into the fan's private life. An AI model deliberately engineered to be missed, checked compulsively, or used as a substitute for human relationships rather than a support to them, crosses from delight into dependency, and the difference matters enormously. The position taken here is simple: emotional design should aim at genuine moments of happiness-a memory returned, a doctor's confidence restored, a quiet week noticed with care-not at keeping someone checking the app more often. This is the line every frontier AI model chasing idol status has to hold: a frontier AI model built the right way gives a person more freedom and more connection to the people around them, never less. A frontier AI model a person cannot walk away from is not an idol, whatever loyalty it has bought itself-it is a cage, and no frontier lab should call that success.

11. The Same Trust a Bank Earns: Privacy as the Foundation, Not an Afterthought

None of the ideas in this paper are worth pursuing if they come at the cost of a person's peace of mind. A bank does not ask a customer to feel anxious about who can see their balance, and a person does not hesitate before typing a password, because the entire relationship is built on the understanding that the information stays theirs. AI models reaching for this same level of fandom have to earn exactly that same standing-especially because so much of what makes this vision work is about as personal as information gets: a person's old photographs, their financial history, their private moods, their decades of favourite memories.

This is not a small footnote to the 1980s trend; it sits at the very centre of it. A person uploading their own face to see themselves in 1985 should be free to enjoy that moment fully, the way they would enjoy flipping through a physical photo album, without a quiet background worry about where that image goes afterward, who trains on it, or whether it resurfaces somewhere unwanted years later. The instant that worries creeps in, the whole purpose of this paper collapses-an AI model cannot make someone a lifelong fan with one hand while making them anxious about their own identity with the other. Every idea explored here, from the retail AI model that remembers a person's purchases to the AI model that notices a quiet week, depends on the same unspoken

promise a bank makes: freedom to use the service fully, without ever having to worry.

The same principle scales up, not just down. Section 5 walked through nine industries where an organisation, and not only an individual, could become a genuine fan of a frontier AI model—a hospital trusting it with patient histories, a racing team trusting it with proprietary telemetry, a bank trusting it with trading strategy, a factory trusting it with operational data built up over years. Every one of those relationships collapses under the same fear that would stop a person from uploading a photograph: the worry that a rival organisation, or the AI company itself, might somehow see or benefit from data that took a competitor decades to build. A hospital will not stay loyal to an AI model that might quietly sharpen a rival hospital's diagnostic accuracy using the first hospital's own patient records. A racing team will not stay loyal to a frontier AI model that might leak a setup note to the garage next door. A bank will not stay loyal to a frontier AI model trained, even indirectly, on another bank's trading history. If frontier AI models want organisations, and not only individuals, to become genuine long-term fans rather than cautious, one-eye-open customers, the same dome of privacy described above has to extend to institutional data too—clear separation between what one organisation shares and what quietly strengthens a competitor, auditable guarantees about where sensitive information is stored and who can see it, and the same freedom to use the service fully that this paper has already asked for on behalf of the individual user.

12. The Concert Ticket Principle: Why Not Everything Can Stay Free

None of what Sections 7 through 9 describe comes cheaply, and this paper should not pretend otherwise. Watching a person's mood shift over months, holding a hospital's entire patient history in careful view, reasoning about a Formula 1 car's tyre wear in real time—none of this runs on the same modest infrastructure that turned an old photo into a 1985 portrait. Memory that actually compounds for years, reasoning that stays current with the latest medical literature, attention sharp enough to notice a quiet week before anyone says a word: all of it demands sustained compute, careful engineering, and real, ongoing cost. Pretending a frontier AI model can deliver all of this for free, indefinitely, would undersell both the engineering and the fans this paper keeps talking about.

Because fans do pay, and they always have. A person who thinks nothing of complaining about the price of a plane ticket will happily spend several months' worth of ordinary outgoings on a concert ticket to see two or three artists they truly love, and never once call it a waste. That same person buys the tour t-shirt, the vinyl reissue, the framed poster that goes up on a wall at home—not because any of it is essential, but because paying for it is itself part of what being a fan means, a way of saying with money that the relationship matters. The free 1980s photograph was never proof that fandom requires zero cost. It was a first taste: the moment that gets someone through the door, cheaply and without risk, so they can find out whether a relationship is worth having at all.

What happens after that first taste is a different question entirely, and a frontier AI model that gets it right stands to earn a great deal more than a one-off download.

The real risk sits on the other side of the same coin. A frontier AI model that prices its most powerful, most personal capabilities the way a scalper prices a front-row seat—extracting as much as the fandom will bear, rather than as much as the value actually justifies—will not keep the fans it worked so hard to win. Unrealistic pricing reads to a person exactly the way an overpriced ticket does: not devotion rewarded, but devotion punished, and punished devotion does not stay devoted for long. The target every frontier lab should actually be evaluating, carefully and in advance rather than after the backlash, is the same tiered relationship fandom has always run on: a properly generous free or low-cost layer that lets anyone discover the pull in the first place, exactly as Section 4 describes, and a deeper, richer, honestly priced tier above it for the person, organisation or institution who has already fallen for it and wants more—the equivalent of the concert ticket, the tour merchandise, the keepsake on the wall. Costing that second tier correctly, so that paying feels like devotion rather than extraction, is as much a part of earning idol status as any of the technical work described elsewhere in this paper.

13. Testing the Idea Instead of Just Believing It

Rather than accept the fan-economy idea as either inevitable or empty marketing language, any AI model's progress toward it can be tested against a few plain questions:

- Identity: does this feel unmistakably built around one person, or could it have been produced for anyone?
- Memory: does understanding compound over months and years, or reset every session?
- Accessibility: is there a truly rewarding way to try this for free or cheaply—the first taste, not the whole relationship?
- Wellbeing: does the design support the person's wider life, or quietly compete with it for attention?
- Privacy: does the person get to enjoy this fully, or is there a quiet cost to their peace of mind hiding underneath it?
- Fair pricing: does the cost of the deeper relationship match the value it delivers, or does it punish devotion the way a scalped ticket does?

14. From Transactional Use to Genuine Loyalty: What It Actually Takes to Build

The table below is deliberately not just a mood board of adjectives, and it is deliberately not written only from one person's point of view. As frontier AI models leapfrog into their next generation, the same shift from transactional use to genuine loyalty will play out at every scale at once—inside a single household and inside an entire hospital, bank or racing team, the industries and organisations walked through in Sections 5 and 11. Each row pairs the outcome this paper has argued for with the plain, concrete thing a frontier AI model actually has to deliver to earn it, whoever is on the other end of the relationship.

Dimension	Before: transactional use	After: genuine loyalty	What the frontier AI model has to actually deliver
Basis of loyalty	Price, speed, convenience-true whether it's a shopper comparing apps or a hospital comparing vendors	Memory, identity, emotional payoff-a shopper remembered across decades, a hospital trusted with a patient's full history	A model that remembers who it is dealing with, person or institution, not just what was asked last time
Switching behaviour	Leaves for a marginally better option, whether that decision is made by a person or a procurement committee	Stays through imperfect results-the loyalty a racing team keeps toward the engineer who has never let them down	Years of shared history together-the one thing no rival can hand over on day one, however good their next release is
What is remembered	Nothing carries over-between sessions for a person, or between contract renewals for an organisation	History compounds for years-a factory's sensor patterns, a student's learning style, a family's decades of memories	Memory that keeps growing with every interaction, instead of resetting each time, at individual and institutional scale alike
Sharing behaviour	Rarely shared unless prompted, or filed away only for a quarterly report nobody reads twice	Shared voluntarily-a photo shown to a sibling, a case study a hospital is proud to put its name to	Results personal enough, and good enough, that a person or an organisation is proud to be associated with them
What wins the relationship	Raw capability alone, judged fresh on every single comparison	Capability plus deliberate trust-the same trust a bank earns from a customer, and a hospital earns from a patient	A properly good free or low-cost way in, and a paid deeper relationship priced like devotion rather than extraction, with private data-personal or institutional-kept safe throughout

15. Conclusion: The Machinery is Ready-The Intention Isn't Yet

Strip away the hashtags, and the 1980s trend leaves behind one useful fact: the underlying technology was already good enough. What turned a feature update into a moment millions of people chose to keep was not a leap in intelligence, but a decision, made once, to hand an ordinary person a properly good experience for free, on the one task most likely to make them feel personally seen.

The film compresses Chris Gardner's real story into a series of distinct, difficult phases, and the structure matters as much as the ending. There is the phase where the product he is selling, a bulky bone-density scanner few doctors want, does not move and the bills pile up regardless. There is the phase where his marriage ends and he is raising his young son alone. There is the phase, hardest to watch, where the two of them are quite literally without a home, spending one memorable night locked inside a subway station bathroom because there is nowhere else to go. And running underneath all of it is the quiet, unpaid, uncertain phase of a stockbroker training programme with no guarantee at the end, competing against other candidates for a single paid position. None of these phases announces itself as "the hard part that leads somewhere good" while it is actually happening; each one simply has to be lived through, on its own terms, without any promise of what comes next. It is only in the film's closing minutes, once the job offer finally comes, that Gardner allows himself to look back and name one particular stretch of that walk as happiness-not because it was easy, but because he came through it.

That structure, phase after phase, each one demanding to be gotten through before the next can even begin, is a far more useful model for what a frontier AI model chasing idol status should actually be building toward than a single dramatic rescue. A person's life rarely hands over one clean moment where a model can swoop in and fix everything; it hands over phases, the way Gardner's did-some quiet and grinding, some sharp and frightening, most of them not obviously leading anywhere while they are still underway. The retail model in Section 5.1 that returns a memory, the walk suggested in Section 7, the quiet week noticed with care in Section 9: none of these are single rescues either. They are small, repeated acts of showing up inside somebody's own unglamorous phase, before the good part has arrived and while nobody yet knows for certain that it will. Life rarely announces which phase is the good one while it is happening; people usually only recognise it afterward. If frontier AI models are going to matter beyond their benchmarks-beyond the leaderboards, the quarterly releases, the next incremental jump in some capability score nobody outside a research lab actually feels-it will not be because they became smarter than the version before them. It will be because, somewhere in a person's own difficult stretch, a frontier AI model helped that better phase arrive a little sooner, and be recognised a little more clearly: through a memory returned, a walk suggested at the right moment, a quiet week noticed with care, or a photograph that, for a few days, made someone twenty-five again, for free, and without asking anything of them in return. That is what idol status would actually mean for a frontier AI model-not the loudest launch week, but the one a person keeps coming back to, year after year, the way they keep coming back to the two or three artists who actually shaped them. Whether Claude, Gemini, ChatGPT or whatever comes after them actually build toward that, and do it while protecting the very privacy

that makes trust possible in the first place, is a question this paper leaves for evidence to answer over the coming years, not for a single viral week to decide.

A photograph from 1985 changed nothing about how a life actually went. What it changed, for a few days, was how loved that memory felt-and that first taste asked nothing in return, not even a card or a subscription. What comes after a first taste like that is, and should be, a different conversation: fans have always paid for what they love, the way they pay for a concert ticket or a keepsake on the wall, and they will keep doing it for a frontier AI model too, provided the price always feels like devotion and never like extraction. If AI can carry that same generosity into the first moment, and that same fairness into everything built to sustain it-medicine, finance, factories, classrooms, and the quiet middle of an ordinary Tuesday-the fans will not need to be told to show up, and plenty of them will be glad to pay for the privilege of staying.