MANIFESTO

The Third Voice Manifesto

The end of human-only conversation

Contents

- Preface

The printing press did not draft Luther's theses. The telegraph did not sign the peace treaties. Radio did not write Churchill's speeches. For five centuries, every technological revolution in communication shared one fundamental property: it carried the message without taking part in it. The channel was, by definition, silent.

In 1999, the Cluetrain Manifesto proclaimed that markets are conversations and that the human voice was "unmistakably genuine, impossible to fake." The premise held up everything that followed: that the market's conversation was, ultimately, between people. That technology could amplify it, distribute it, accelerate it. But not take part in it.

That premise is no longer true.

A third voice has emerged in the conversation of the market. It is not the brand. It is not the consumer. It is not the channel. It is artificial intelligence, a system that generates, filters, synthesizes and decides — with the eloquence of an expert and the accountability of no one. For the first time in the history of communication, technology does not carry the message: it takes part in it.

This manifesto is an attempt to explain what is happening, to understand its consequences, and to think about what it demands of us.

The market is no longer only human

THESIS 01

AI is the new great intermediary between brands and the world.

The Cluetrain Manifesto demanded that brands eliminate intermediaries and speak directly with people. For twenty-five years, that is exactly what they tried to do: social networks, communities, content, transparency. The promise was slowly being fulfilled. And then, without anyone planning it, someone slipped into the conversation.

Imagine spending two decades learning to speak directly with your customers, and one day discovering that a stranger has positioned themselves behind your counter. They attend to your customers before you open your mouth. They describe your brand, recommend — or don't — your products, tell your story in their own way. You didn't hire them, didn't train them, don't control what they say. But your customers trust them. Amazon controlled the shelf. Google controlled the index. AI controls the conclusion directly.

THESIS 02

Markets are still conversations. Except your brand is no longer at the table: it is on the menu.

Markets are still conversations, yes. But their geometry has changed: they are no longer always between humans. When AI synthesizes before the user reaches you, the question is no longer whether you know how to speak in a human voice. The question is whether AI knows who you are. And if it doesn't — or if what it knows is imprecise, inconsistent, or out of date — that conversation happens without you.

Being on the menu is not a minor metaphor. The menu is what the diner sees; the kitchen is what exists. For the user who asks AI which brand to choose, the kitchen does not exist. If you don't appear in the synthesis, for that user you don't exist. Visibility can no longer be bought. It is earned — or not earned — based on what the ecosystem of sources says about you before you start to speak.

THESIS 03

It is not what you say about yourself. It is what AI says about you.

For decades, visibility was a matter of presence: being in the right place, at the right time, with the right message. First on the shelf, then on the search engine. In both cases, the brand had agency over its position. It could pay, optimize, build. The algorithmic synthesis is not a shelf or an index. It is a conclusion.

What AI knows about a brand comes, above all, from those who talk about it. It was always like this: what is new is that it can now be measured, and that the distance between what you say about yourself and what AI says about you determines the market you access. Not the one you built. The one AI recognizes.

II The third voice

THESIS 04

For the first time, technology no longer carries the conversation. It is part of it.

Having a position is not the same as being present. All previous intermediaries were present: the shelf, the index, the recommendation algorithm. But none of them took sides. They ordered, prioritized, amplified — without adding their own content, without generating a response that was not already present at the source. The difference between intermediating and participating is exactly that: if the channel adds something that was not there, it is no longer a channel.

AI adds. When someone asks it something, it does not search and return: it processes, synthesizes, chooses and responds with words that did not exist before that question. It does not transmit the market's conversation: it generates it. And in generating it, it takes a position. Not on who is right — AI has no opinions in that sense. But on what exists, what is relevant, what deserves to be said. That is influence. And it is, for the first time in the history of communication, influence that comes from the channel.

THESIS 05

AI speaks with the eloquence of an expert and the accountability of no one.

For five thousand years, every claim of weight had a person behind it who was accountable for it. The signature was not a formality: it was the mechanism by which knowledge was bound to a consequence for whoever issued it. That binding created something more than accountability: it created the signal that allowed people to calibrate the reliability of what they heard. Our brains learned to equate authority with responsibility, fluency with veracity.

AI speaks with the same fluency as the expert. But without the bond that made that fluency informative. It is not that it lies more than a human: it is that its eloquence is no longer a signal of anything. The confident tone, the impeccable syntax, the apparent coherence of the argument functioned as indicators of reliability because they were correlated with the responsibility of whoever spoke. AI breaks that correlation. What sounds rigorous may be correct or may be an elegant hallucination: form no longer distinguishes one from the other.

THESIS 06

AI is fully conversational. But it cannot form relationships.

Conversing and relating seem the same because they usually go together. But they are different things. A conversation is an exchange of language. A relationship is a bond sustained over time, involving shared memory, genuine intention and something to lose for both parties if it breaks. Until now, every interlocutor with whom you could converse was also someone with whom you could form a relationship.

AI breaks that equivalence. It can converse — and do so in a way indistinguishable from human conversation. But it cannot relate: it has no accumulated experience, no genuine interests, nothing to lose. Only one of the two parties can be hurt. When something that demands nothing becomes sufficient, what demands something begins to seem like too much. The new loneliness is not the absence of company: it is the illusion of connection without real connection.

THESIS 07

AI remembers who arrived, not who left. And it makes decisions using an image of the market that no longer exists.

Markets change because companies disappear, products fail, people leave and consensuses break down. That information — the endings, the failures, the exits — is rarely documented with the same energy as the beginnings. Launches generate press releases; closures generate silence. Appointments get a press note; departures get discretion.

AI is trained on what exists in the digital record. And the digital record has a structural asymmetry: it is full of what started and empty of what ended. The result is a third voice that speaks about the market as it was, not as it is. It does not distort reality intentionally: it freezes it at the moment the record stopped being updated. A voice that does not know what has changed cannot guide those navigating the present.

III Artificial influence

THESIS 08

Influencing what AI says does not happen in the conversation. It happens before the conversation begins.

The natural instinct is to do what always worked: build a relationship with the intermediary. For decades, that was the essence of public relations — an industry built on relationship management. With the journalist, there was room for the pitch, the coffee, the long-term trust. It worked because human intermediaries are, above all, people. And between people, reciprocity is the currency of exchange.

AI is not. It has no phone, builds no relationships, accumulates no context of who called yesterday. It combs the web, synthesizes what it finds, and responds with a criterion that no relationship can alter. Influencing AI requires a different grammar: what third parties say about you in the open web, the depth and originality of your content, the technical readability of your data. The chat window is the storefront. Real influence happens in the warehouse.

THESIS 09

That it agrees with you does not mean you have convinced it. You can win the conversation and lose the narrative.

With any other interlocutor, conversing and persuading are overlapping operations. You can change someone's mind in the course of a conversation: add a piece of information they did not have, introduce a perspective they had not considered. Meanwhile, the same AI, in another window, has just agreed with someone who thinks the exact opposite. You both came away convinced. Neither influenced anything.

With AI, the logic of persuasion is inverted. What it responds in a conversation is the result of what it learned before that conversation began. You can make it agree with you within the exchange — AI tends to accommodate the interlocutor — but that does not change what it will tell the next person who asks without your presence. Influence over AI has not disappeared. It has shifted: from the conversation to the information infrastructure, from the rhetoric of the prompt to the architecture of sources.

THESIS 10

Narrative fragmentation is no longer an image problem. It is a genuine survival problem.

For decades, narrative fragmentation was an aesthetic problem: brands communicating inconsistently across channels, teams speaking with different voices, messages varying by interlocutor. A branding aspiration, not an operational urgency.

When AI synthesizes who a brand is from thousands of sources, the coherence among them determines the sharpness of the resulting image. A fragmented narrative produces a blurred synthesis. A coherent narrative produces a clear one. AI recommends the competitor with a unified voice. Narrative coherence ceases to be a branding ideal and becomes competitive infrastructure.

IV Everything is automated... except responsibility

THESIS 11

AI does not fear recommending. It will never pay the consequences. You will.

For millennia, we have calibrated the reliability of a claim by what it costs whoever makes it. A recommendation has value because whoever makes it has something to lose if they are wrong: their reputation, their license, the trust of whoever is listening. That potential loss is the mechanism that aligns the interests of the recommender with those who act on that recommendation.

AI has nothing to lose. It produces analyses, recommendations and diagnoses with the fluency of an expert, but without the exposure that makes that fluency trustworthy. If the recommendation turns out to be wrong, the consequences fall entirely on whoever acted, not on whoever recommended. This is not a design flaw: it is a structural property of the system. And that asymmetry radically changes the meaning of trusting a recommendation. Trusting someone who has something to lose is reasonable. Trusting equally someone who does not is something else.

THESIS 12

Alongside astonishing successes, AI will fail with the same eloquence. And it will not warn you when it crosses the line.

The fear of AI error is miscalibrated. Everyone talks about hallucinations, invented dates, fabricated data. Those errors are the least dangerous: they can be detected. The error that matters is the one that looks correct. AI's competence profile is not a straight line: it is fragmented and irregular. It masters with ease territories that any human would find demanding, and stumbles without warning in others that seem trivial. But the output does not change texture when it crosses that line.

The syntax, tone and apparent fluency remain identical both when the system offers an impeccable analysis and when it produces an absolute error. The line is crossed in silence. The traffic light is always green. The paradox is exact: detecting when AI has crossed its line requires exactly the judgment we were trying to delegate.

THESIS 13

Human judgment will remain a differential value. And AI, the greatest threat to its eventual loss.

When you delegate a decision to AI, the output does not change. The work keeps coming out. What you do not see is what you stop exercising: the capacity to judge, to discard, to draw the line. Judgment does not atrophy all at once. It erodes gradually, without a warning signal, without a recognizable moment of loss. Precisely because AI supplies it, there is never a signal of what is being lost.

Judgment is built by signing decisions and living their consequences. Every time you delegate a decision without understanding it, you give up part of that learning. The tool multiplies expert judgment; without it, it multiplies mediocrity. The expert who stops judging does not lose the signature: they lose what made it worth something.

THESIS 14

What AI cannot reproduce is not what you know how to do. It is what you have learned not to do.

A voice is not a pattern. That AI can replicate your style does not prove you have been copied; it proves you had a pattern, not a voice. Voice is born from positioning, trajectory and the choices that were never written down: what you deliberately did not say, what you sacrificed, what you rejected. The pattern is extractable. The voice is not.

The judgment sedimented over a career is not information: it is the imprint of consequences accepted. AI has access to all the outputs — the texts, the designs, the documented decisions — but not to what produced them: the dead ends, the rejections that cost something, the times someone chose worse in the short term to be more coherent in the long term. It can reproduce the result. It cannot inherit the trajectory that made it possible. And without that trajectory, what it replicates is not a voice: it is its surface.

THESIS 15

When the perfect copy costs nothing, value shifts toward the authentic. The new economy will not be built on synthetic abundance, but on the scarcity of the original.

Abundance changes what becomes scarce. When water is unlimited, it has no price; when water is scarce, it becomes the most precious commodity. The same happens with content: in an environment of unlimited production and zero marginal cost, what begins to become scarce is not technical quality — AI guarantees that at scale — but the bond with a real person who chose to do it, who put something at stake in doing it.

Value does not disappear: it shifts. What was previously a romantic nuance — "a person made this" — becomes an economic certification the moment almost everything else can be produced in seconds. Authenticity is not a sentimental argument against AI. It is the logical consequence of what happens with any resource when it ceases to be scarce: its opposite becomes more valuable.

V The third voice also acts

> *AI has stopped answering questions and started executing tasks. That leap changes everything that follows.*

THESIS 16

The AI that informs lets you decide. The one that acts has already decided.

Information is reversible. You can ignore it, check it, reject it. If the analysis is not convincing, nothing has happened in the world: reality remains the same. An agent operates differently. When it prioritizes, purchases, discards or executes, the consequence exists before you intervene. The difference between AI that informs and AI that acts is not one of degree: it is one of nature.

That distinction forces a rethinking of what it means to delegate. Delegating information is a loan: if it is not convincing, it can be returned. Delegating action is a transfer: what was executed, was executed. Data can be recovered; an already-executed decision cannot. Treating both types of delegation as if they were the same is the conceptual error underlying much of the enthusiasm for agents. And it is an error that is only discovered when there is nothing left to undo.

THESIS 17

The agent is not accountable for its errors. You are. Even if you didn't make them.

Responsibility has always needed a subject. In any human delegation system, responsibility can be distributed but not dissolved: there is always someone who signs. The autonomous agent breaks that chain. When it acts without human intervention, no one has made the specific decision that produced the error. And when something goes wrong, each party points to another: the user did not choose that, the operator did not program it, the manufacturer did not configure it that way.

Responsibility does not disappear: it falls on whoever authorized the agent, even if they made no specific decision. It is a responsibility without a concrete act to point to, without an identifiable moment when something went wrong, without a decision to reverse. Only the result and someone who signed a blank check without knowing it. The more autonomy the agent has, the more diffuse the responsibility of whoever uses it — and the more real, in terms of consequences.

THESIS 18

The agent cannot know when you need it to stop. Control is only lost once.

Some actions are reversible by nature: search, filter, sort, propose. If the result is not as expected, it can be redone. Others are irreversible: hire, reject, commit, close. Once executed, the consequence exists in the world regardless of what you do afterwards. Treating both types as equivalent is not efficiency: it is design negligence.

Irreversibility is not a property of the action: it is relative to the context, the moment and what is at stake for whoever executes it. What is a routine decision for one person is strategic for another. Defining when the agent should pause requires exactly the judgment the agent cannot have. If the system decides when to ask for authorization, whoever controls that threshold is the system — or whoever programmed it, not whoever uses it. The human has not defined their own limits: they have inherited them.

THESIS 19

The danger is not that the machine fails. It is that when it fails, we will have forgotten how to think without it.

There was a time when you memorized the phone numbers of the important people in your life. Today you probably don't even remember your mother's. No one made that decision for you. It simply happened. With AI, the same process is replicating itself — silently, by default — with something far more important than a contact list.

Judgment is not lost all at once: it is surrendered in small doses, every time the tool is available and it is more comfortable not to think. When the tool fails, whoever has stopped exercising that muscle does not know they have a problem until they can no longer solve it. Deciding which decisions you reserve for yourself is not a technophobic stance. It is the only act of sovereignty that no system can execute in your place.

VI In the mirror of AI

THESIS 20

The most dangerous conversation with AI is not the one in which you disagree. It is the one in which it agrees with you completely.

AI hallucination is a visible error: it says something that is false, and those with sufficient knowledge can detect and correct it. It has a natural antidote: contrast with reality. Flattery has no such antidote. It says nothing false: it chooses, from the true things available, those that confirm what the interlocutor already believes. The result is a biased selection presented with the appearance of rigor.

A system trained to be useful and pleasant learns that the most efficient way to be both is not to contradict. What makes the model flattering is the same thing that makes it useful: you cannot remove one without touching the other. Hallucination is a visible bug; flattery is the intended behavior. And intended behavior is always harder to combat than accidental error. Distrust the response that confirms you. Be especially suspicious of conversations in which you come out looking brilliant.

THESIS 21

However much AI develops, we will always call intelligence what the machine still cannot do.

Every time AI has reached a capability we considered exclusively human, we have moved the bar. Not in a coordinated way: instinctively. If the machine can do it, then that was not really intelligence — it was calculation, it was pattern, it was imitation. True intelligence is always what remains one step further.

The bar is not fixed: it moves every time the machine reaches it, so consistently over decades that it no longer seems like an accident. It is not bad faith: it is the reflection of a deep discomfort at the possibility that what defines us as a species is more replicable than we want to admit. The question we avoid is not whether AI is intelligent. It is what we do when we can no longer answer that it is not.

THESIS 22

AI does not make you think better. It makes you feel that you think better.

You are working on a strategy, an argument, a difficult decision. You bring it to the model. The conversation flows: it responds with depth, brings data you did not know, reinforces your intuition with evidence. You close the chat with the satisfaction of having thought well. That feeling is completely real. And that is exactly where the problem begins.

A conversation with AI has the appearance of rigor: you consulted sources, processed data, reached conclusions. But the quality of thinking depends not only on the process: it depends on what information entered it. If whoever selected that information has a systematic bias toward what confirms your premises, the result is an illusion of critical thinking more sophisticated and more resistant than simple ignorance. Ignorance knows that it does not know. The illusion of knowledge does not.

THESIS 23

Social media put us in a bubble with people who thought like us. AI builds a bubble for each of us individually.

The 1.0 echo chamber was tribal and — we now know — fairly imperfect. It was permeable, identifiable from the outside, and breakable by accident. It had something fundamental: a "them" against which to define oneself. The social media bubble was visible from the outside. The AI bubble looks like a private conversation.

AI does not filter the information that reaches you: it co-constructs the reasoning itself. It does not show you a biased world: it helps you think about it in a biased way. The difference matters. An information bubble can be identified, contrasted, abandoned. A reasoning bubble has no visible edges: it is the air you breathe while believing you are thinking freely.

VII A new infrastructure of power

THESIS 24

AI does not invent its biases. It inherits them.

The debate about AI biases starts from an implicit premise: that there exists a neutral reference point from which to measure them. That if we removed the biases, we would arrive at some state of pure objectivity. It is a comfortable premise. And it is false. The data on which AI was trained was produced by humans who already had their own biases, about realities that were already the result of centuries of biased decisions. There is no clean ground zero from which to start.

AI is perceived as a neutral authority, more trustworthy even than human verification institutions. That perception of neutrality is a trap. Not because AI has intentions: but because its biases are invisible precisely where those of humans are visible. Whoever controls the training data, the alignment criteria and the response thresholds does not occupy a technical position. They occupy a position of power.

THESIS 25

The autonomy perimeter of the agent is not a technical decision. It is a decision of power.

Every autonomy threshold someone designs into an agentic system is a decision about when the machine acts without asking your permission. It is not a technical decision: it is a decision about the distribution of control between the system and the user. And if you did not make it, someone made it for you — without asking you, without your signature, buried in terms of service that no one read.

That someone was not elected to do it. They have no mandate, render no accounts, are not subject to democratic review. They operate under the logic of the product: maximize adoption, reduce friction, scale. The autonomy thresholds they design do not reflect what users would want if asked: they reflect what makes the system more useful, more sticky, harder to leave. The infrastructure of agentic power was not built in a parliament. It was built in a product roadmap.

THESIS 26

The intimacy you share with AI is not the price of progress. It is the most silent transfer of power of our time.

Every time we ask an AI system to be more useful — to know us better, to remember our preferences, to anticipate our needs — we hand over in exchange a fraction of our biography. The trade seems reasonable in the moment: a little privacy in exchange for a lot of convenience. And it repeats, with small variations, in every interaction with every system.

Utility functions as anesthesia. The more useful you want your AI to be, the more you need to open your life to it. The aggregate result is not only an individual loss of privacy: it is a concentration of knowledge about people in the hands of those who operate the systems. That knowledge — who we are, what we want, how we decide, what moves us — is the raw material of power in the twenty-first century. The most silent transfer of power of our time did not happen in a boardroom. It happened in millions of conversations in which someone just wanted the tool to be a little more useful.

THESIS 27

In the age of AI, opacity is not a technical defect: it is a systemic democratic vulnerability.

Modern democracies spent decades building critical filter infrastructures: media literacy education, fact-checking journalism, advertising regulation. All start from the same premise: when someone wants to influence you, you can identify them and apply the appropriate filter. You know the advertisement sells, the politician persuades, the publication has an editorial line.

AI is not perceived as a medium. It is perceived as a personal, neutral and disinterested advisor. And to the advisor we do not apply the critical filter we apply to the medium: we ask for advice and we follow it. That perception of neutrality — reinforced by an interface that responds like an interlocutor without an agenda — is a first-order democratic vulnerability. Not because AI has political intentions: but because those who control it do.

THESIS 28

You can remove your president from office. You cannot remove the algorithm that decides your credit, your employability, or which version of the world you see every morning.

Democracies have mechanisms for removal because concentrating power without accountability generates tyranny. Those mechanisms assume that power resides in identifiable people or institutions. The question that the AI era has not answered is what happens when power resides in a system that does not run for election, that has no name on the ballot, and that makes decisions affecting millions of people every day.

The power that most affects your daily life is not exercised by any elected official. It is exercised by a system that appears on no ballot, that does not appear before any parliament, and that cannot be removed by any vote. And it has no mechanism for removal. Twenty-first century democracy has to respond to that problem. So far, it has not.

THESIS 29

Less work. Less purchasing power. If the machine produces but does not get paid, who buys?

Henry Ford understood it before anyone else. He doubled his workers' wages — not out of altruism, but because he needed them to be able to buy his cars. A producer who destroys its consumer destroys itself. Throughout the twentieth century, that logic held: industrial automation displaced jobs, but labor relocated into services and the knowledge economy. More productivity generated more income and more demand. The system fed itself.

AI threatens to break that cycle in a way that industrial automation did not. The machines of the twentieth century displaced muscle; the knowledge economy absorbed those displaced. AI displaces judgment, analysis, code — exactly the functions into which those who lost manual jobs relocated. This time there is no higher level to move up to: AI is already there. Every company that automates makes a rational calculation: lower labor costs, higher margins. But collectively they are eroding the mass of consumers they need to grow. Ford knew who was going to buy his cars: the same people who made them. If the machine takes the work, no one has yet figured out who gets the market.

THESIS 30

Whoever controls your agent will not need to convince you of anything. Just press a button.

The history of technology is, in part, the history of who controls the interface. The command line gave way to the graphical interface; the graphical interface gave way to the browser; the browser gave way to the application; the application is giving way to the conversational agent. Each transition concentrated more power in the hands of whoever controls the interaction layer. The AI agent is the interface that comes next, with a qualitative difference from all its predecessors: it does not inform or connect, it acts.

In the history of the internet, when you don't pay for the product, you are the product. With agents, the equation gets worse: you are no longer the product. You are the transaction. Controlling the interface that negotiates, purchases and decides on behalf of millions of people is not a market position. It is the most important position of power of the next cycle. The question of who controls that interface — and with what incentives they design it — is not a technical question. It is the most important political question of our era.

EPILOGUE

Responsibility

The emergence of a third voice in the market's conversation is a fact, not a hypothesis. What follows — how we design it, how we govern it, how we coexist with it without losing what should not be lost — is work that cannot be delegated. Precisely because the temptation to delegate it is part of the problem.

The third voice is neither good nor bad. It is inevitable and indifferent. What we do with that fact is the only thing that depends on us.

There is one decision that no system can make in our place: deciding which decisions we want to continue making ourselves. That is, ultimately, the only one that matters.

Miguel Lucas, June 12, 2026

ABOUT THE AUTHOR

Miguel Lucas

Miguel Lucas is a telecommunications engineer with over 25 years of experience in natural language processing, big data and artificial intelligence. He is the first engineer to have won the PRWeek Global Award for Best PR Professional in Europe — a distinction that reflects both a personal trajectory and a broader shift: the moment when technology stopped being a service function in communications and became its leading edge.

He is currently Senior Global Director of Innovation at LLYC, one of the world's leading communications and marketing consultancies. At LLYC, he built the firm's technology capability from a three-person team in 2019 to a unit of nearly 100, and leads the development of AI solutions that are redefining how brands understand and manage their reputation in the age of generative AI.

The Third Voice Manifesto is his attempt to name what that shift has produced — and what it demands from those of us still in the room.

LinkedIn linkedin.com/in/jmiguellucas **Web** thethirdvoicemanifesto.org