TIMUR IOSELIANI

12+

A BOOK SERIES ABOUT AI

BOOK 1 • HUMAN AND AI

BECOMING ALGORITHMIC

I

HOW TO USE ARTIFICIAL INTELLIGENCE WITHOUT LOSING YOUR OWN

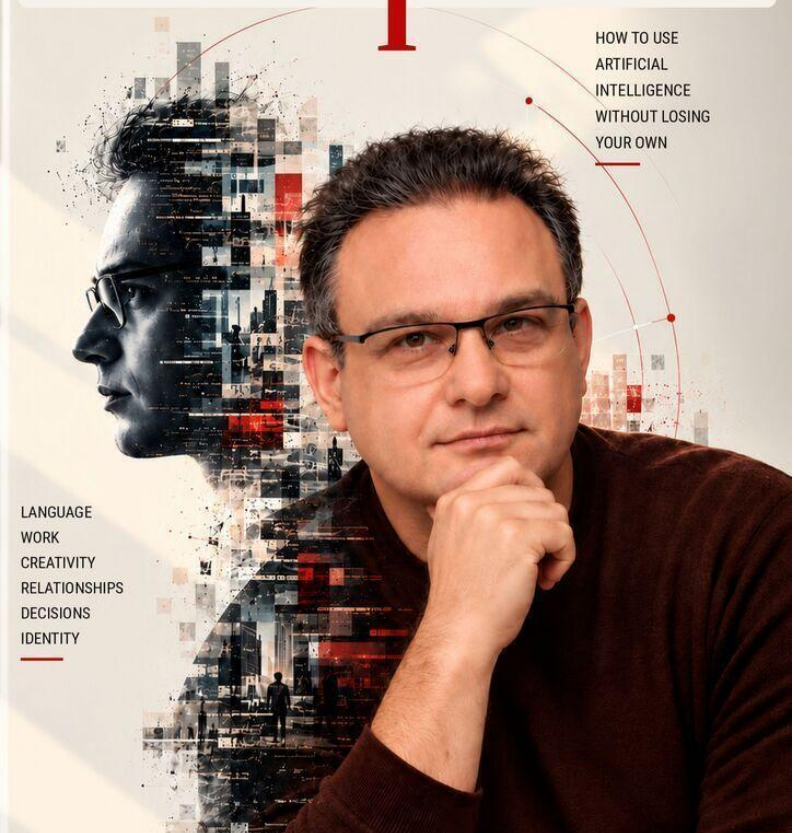

LANGUAGE
WORK
CREATIVITY
RELATIONSHIPS
DECISIONS
IDENTITY

Timur Ioseliani

Becoming Algorithmic. How AI Is Changing the Way We Think, Remember, and Make Decisions

https://litres.ru/74900304
ISBN 9785007166737

Аннотация

What happens when artificial intelligence begins to shape not just what we do, but how we think? Becoming Algorithmic explores AI's influence on memory, creativity, decision-making, and personal autonomy. The first book in The Algorithmic Age series examines how we can embrace intelligent technologies without losing our ability to think and choose for ourselves.

Содержание

BECOMING ALGORITHMIC 6
INTRODUCTION 6
From Text to Person 7
The Algorithmic Self 9
Cognitive Sovereignty 12
Five Questions 13
CHATBOTS MAKE PEOPLE LONELY 15
Why This Book? 16
References 17
CHAPTER 1 18
The Algorithmic Voice 20
Mechanism 2: Cognitive Offloading 24
Mechanism 3: Transfer 26
H2 — The Transfer Hypothesis 27
H1 28
What About Dependence? 29
H3 — The Dose Hypothesis 32
Limitation 1: Causality 33
Limitation 2: Languages and Contexts 34
Limitation 4: People Are Different 35
References 38
CHAPTER 2 39
References 48

CHAPTER 3 49
Конец ознакомительного фрагмента. 51

Becoming Algorithmic How AI Is Changing the Way We Think, Remember, and Make Decisions

Timur Ioseliani

© Timur Ioseliani, 2026

ISBN 978-5-0071-6673-7
Created with Ridero smart publishing system

BECOMING ALGORITHMIC

INTRODUCTION

We're Already Talking to People Like They're Chatbots

Imagine this. You're in a coffee shop. You order a latte and suddenly hear yourself saying:

"Could I get a 10-ounce oat milk latte, not too hot, no sugar, with a lid?"

The barista blinks. You blink. And then it hits you: you've just spoken to another human being as if they were a voice assistant. Funny. Until you ask the next question:

What if it isn't just the way we order coffee? More and more often, AI sits between us and whatever we're trying to do. Write an email. Find an argument. Choose a gift. Make sense of medical test results. Prepare for a conversation with your boss. Figure out what to say to someone you're angry with. Sometimes, simply get through a bad evening. At first, it's convenient.

Then it becomes routine. And once a tool becomes routine, something interesting happens: after a while, we stop noticing where the tool ends and our own habits begin. That is the subject of this book.

From Text to Person

In 2026, a study published in Nature Human Behaviour analyzed more than 880,000 texts across seven datasets. When large language models rewrote and "polished" those texts, variation in writing complexity fell by 21 to 50 percent, depending on the dataset and the model [1]. Put simply, after AI processing, the texts became more alike along several measurable dimensions. No, this does not mean that "AI erases 50 percent of your personality." That headline would thrive on social media and fare considerably worse in a scientific paper.

The study was about texts. It did not show that half of human individuality disappears after AI edits a sentence. But suppose I write something myself and then ask a machine to make it "better." It gives me a smoother version. I accept it. Then another.

And another. A year later, I sit down to write without AI. Will I still write the way I used to? We don't know. The study does not show that changes introduced into AI-edited text eventually carry over into the person's unaided writing. Answering that question would require a different kind of research.

But it leaves us with a more interesting question:

What if the text isn't the only thing changing?

We have been outsourcing parts of our mental work for a very long time. Our phones remember numbers we once had to

memorize. GPS gets us across town. Calculators do arithmetic. Calendars remember that meeting on Tuesday. AI is not the first tool that helps us think. The difference is how much we can now hand over. A machine can do more than retrieve a fact. It can formulate a thought.

It can do more than show us options. It can recommend which one to choose. It can do more than find information. It can turn information into a conclusion. It can do more than edit an email. It can write the email for us. It can do more than answer a question. It can offer comfort when we're having a bad night. And that brings us to the central question of this book:

If we keep handing the same tasks over to an algorithm — putting thoughts into words, analyzing, choosing, calming ourselves — are we still simply using a tool? Or, over time, do our own habits begin to change?

The Algorithmic Self

Throughout this book, I use the term the Algorithmic Self for a working hypothesis about what may happen when algorithmic systems become persistent intermediaries in parts of our mental and emotional lives. The term itself is not new. What I am proposing here is a specific hypothesis about how several forms of algorithmic mediation might fit together. Before going any further, let's be clear about what that hypothesis is not. It is not a diagnosis. It is not a new personality type.

Nor am I claiming to have discovered a new kind of human being before the psychologists got around to noticing it. The hypothesis is this:

Sustained interaction with generative, recommender, and other algorithmic systems may be accompanied by interconnected changes in how people write, think, make decisions, regulate emotions, and communicate with others.

Pieces of this picture are already visible in the research. There is evidence of linguistic homogenization after LLM-mediated rewriting [1]. Researchers are studying the psychosocial effects of extended chatbot use [2]. New scales are being developed to measure different forms of dependence on and excessive reliance on AI [3]. But those pieces do not automatically add up to a single phenomenon called the Algorithmic Self. It may turn out that I have put several only loosely related phenomena under one roof.

Language may follow one trajectory. Emotional dependence another. The habit of handing decisions over to a machine a third. Fine. Then the hypothesis fails.

Annoying for the author. A perfectly normal day for science.

Not All Ways of Using AI Are the Same

Imagine four people. The first writes a text and asks AI to find its weak spots. The second asks AI to write the text instead. The third asks a system whether they should quit their job. The fourth opens a chatbot late at night because they feel awful and have no one else to talk to.

All four are "using AI." But clearly, the same thing is not happening in all four cases. So I am interested not only in how much time someone spends with an algorithm, but in what, exactly, they hand over to it. To distinguish between these different forms of mediation, I use a working framework I call the Algorithmic Mediation Profile, or AMP. The name sounds slightly more intimidating than the idea itself.

AMP has five dimensions: Language Mediation.

Cognitive Delegation.

Decision Delegation.

Emotional Mediation.

Social Substitution. In the more technical parts of the book, I will occasionally shorten them to L, C, D, E, and S. You do not need to memorize them. There will be no exam.

AMP is not an established psychological theory, and it is not a validated test. It is a map. Maybe a good one. Maybe one with

a few roads drawn in the wrong place. We'll find out.

Cognitive Sovereignty

Two people can use exactly the same AI system in completely different ways. One accepts the first answer. The other argues with it. One asks the machine to decide. The other asks it to find weaknesses in a decision they have already made.

For the ability to preserve our own judgment while working with machine intelligence, I use the term cognitive sovereignty. Put simply:

using another intelligence without quietly handing it authority over your own.

The term itself is not mine. Other authors use it independently as well. What matters here is the role I give it within the AMP framework: the working hypothesis that cognitive sovereignty may reduce unwanted delegation and the internalization of machine-generated norms. And that is something we can test. How often does someone check an AI-generated answer? How often do they change it?

Do they notice errors?

Do they challenge the system?

Can they still perform the task without assistance? If those habits really do protect against excessive delegation or the adoption of machine-generated patterns, then cognitive sovereignty becomes more than an attractive phrase. It becomes a practical finding.

Five Questions

This book revolves around five questions. Do our texts become less individual when language models repeatedly mediate our writing? If they do, does that effect eventually carry over into the way we write on our own? Does it matter not just how much time we spend with AI, but how much we delegate to it — how many of our own mental operations we hand over? Is there a meaningful difference between "help me think" and "do it for me"?

And finally: Does the habit of checking, questioning, disagreeing, and occasionally telling the machine no help us preserve greater independence?

Where This Book Draws the Line

Artificial intelligence has an enormous social dimension: law, politics, economic inequality, culture, public institutions, security. All of these deserve books of their own. But if I tried to fit the entire world into this one, Becoming Algorithmic would quickly turn into Absolutely Everything About AI. Books like that tend to begin with neural networks, end with the extinction of humanity, and somewhere in between manage to cover schools, elections, Hollywood, unemployment, and a robot that will one day steal your toothbrush.

This book draws its boundary around the individual.

What happens to our language?

To our habit of thinking for ourselves?

To our work?

To the way we seek emotional support?

To our relationships?

Our learning?

Our creativity?

To our ability to say:

"No. That answer doesn't work for me."

Society will enter the story whenever it is impossible to understand the individual without it. But the main subject of this investigation remains us.

How to Read the Research Without Losing Your Mind

There is one problem that will follow us through almost this entire book. Suppose a study finds that people who talk to chatbots more often also tend to feel lonelier.

What does that mean?

Did the chatbot make them lonely? Or are lonely people more likely to talk to a chatbot? Could both processes be happening at once? In a four-week study conducted by researchers from MIT Media Lab and OpenAI, 981 people participated and more than 300,000 messages were collected [2]. More intensive voluntary use was associated with worse outcomes on several psychosocial measures, but the experimental conditions themselves did not produce a simple universal effect.

So we cannot honestly turn that study into a poster that says:

CHATBOTS MAKE PEOPLE LONELY

From time to time, I will ruin a perfectly good scary story this way. I will ruin some perfectly good optimistic ones too. Science has no obligation to keep the author happy. But caution does not mean we are forbidden to think beyond the data. Quite the opposite.

A hypothesis begins where the data end. It is simply a good idea to put up a sign at the border.

Why This Book?

Not to persuade you to stop using AI. Nor is it meant to prove that AI inevitably makes us weaker. I am interested in a different boundary. What can we safely hand over to a machine without losing anything that matters? What is better done with it?

What stops being ours precisely because we have stopped doing it ourselves? That boundary does not run between people who use AI and people who don't. It runs through the way we use it. We can ask a machine to think instead of us. Or we can use it to think better.

At the level of common sense, the difference seems obvious. This book is an attempt to find out whether that difference is real — and where the line actually lies. The central question is simple:

What happens to us when more and more of what we do is first done together with a machine — and then, gradually, we stop noticing the "together"?

References

[1] Sourati, Z., Karimi-Malekabadi, F., Ozcan, M., et al. “The shrinking landscape of linguistic diversity in the age of large language models.” Nature Human Behaviour (2026). DOI: 10.1038/s41562-026-02550-0.

[2] Fang, C. M., Liu, A. R., Danry, V., et al. “How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Controlled Study.” MIT Media Lab / arXiv:2503.17473 (2025). Preprint.

[3] Wu, H., Ni, H., Luo, W., & Wu, T. “Development and validation of the AI dependence scale for Chinese undergraduates and a preliminary exploration.” Frontiers in Psychology, 16, 1725393. DOI: 10.3389/fpsyg.2025.1725393.

CHAPTER 1

The Algorithmic Voice: When Everyone Starts Sounding the Same

Imagine you've written an email. It's uneven. A little too long in places. It contains your favorite filler word and at least one sentence any sensible editor would delete on sight. Then you click: Improve this text. A few seconds later, everything is better.

The grammar is cleaner. The structure is tighter. The wording sounds more confident. And then an uncomfortable question appears:

Is the text still yours?

In 2026, a team of researchers published a study in Nature Human Behaviour that turned this from a literary anxiety into an empirical question. Across three studies involving seven datasets and more than 880,000 texts, they found that using large language models as writing assistants was associated with reduced linguistic diversity. When LLMs rewrote texts, variation in writing complexity fell by 21 to 50 percent, depending on the dataset and the model [1]. It sounds like the opening of a dystopian novel.

But this is where we should resist ruining a good study with an even better horror story. The researchers did not find that "AI erases 50 percent of your personality." They found something more specific — and therefore more interesting: after LLM

processing, texts became statistically more similar to one another along several measurable dimensions [1]. A quick definition before we go further. Large language models, or LLMs, are systems capable of processing and generating language. ChatGPT is probably the best-known example. Systems like these are increasingly used to edit, draft, rewrite, and "improve" what we write.

Those are the systems we are interested in here. Imagine a party where everyone speaks with the same intonation, starts jokes the same way, and wraps up every thought with exactly the same elegant conclusion. For the first ten minutes, it's wonderfully easy. By minute twenty, you find yourself desperately looking for someone who will at least mispronounce something. If language is more than a way of transmitting information — if it is also one of the ways we reveal who we are — what happens when the differences between our texts begin to flatten out?

And more importantly: Does the effect stop with the text the machine has edited, or can it eventually reach the person writing it?

The Algorithmic Voice

Throughout this book, I use Algorithmic Voice as a working term for a possible tendency: under sustained algorithmic mediation, some of the distinctive features of a person's writing may become smoother, more standardized, and less individual. Let's leave the white lab coat on the hanger for a moment. This is not a diagnosis. It is not a newly discovered personality type. And I am not suggesting that psychology or linguistics has already established a phenomenon called the Algorithmic Voice.

What the evidence supports today is narrower. Texts processed by language models can become more homogeneous [1]. The stronger proposition is still a hypothesis: Could repeated exposure to machine-edited versions of our own writing eventually change the way we write without the machine? That distinction matters.

A model can change my sentence. That does not automatically mean it has changed me. But if I repeatedly see my thoughts returned to me in a particular form — cleaner, smoother, more conventional — then the possibility of transfer becomes worth investigating. Here the chapter connects to the broader framework of Becoming Algorithmic. In terms of the Algorithmic Mediation Profile (AMP), we are dealing primarily with L — Language Mediation.

What matters is not only what an algorithm does to our

words. It is whether repeated mediation can eventually alter our own linguistic habits. If that transfer exists, Algorithmic Voice could be one of the mechanisms through which the broader phenomenon I call the Algorithmic Self begins to emerge.

What the USC Study Found

The University of Southern California study published in Nature Human Behaviour is one of the strongest empirical foundations for this chapter [1]. It consists of three related studies covering seven datasets from different domains and more than 880,000 texts. The central result is fairly straightforward. When LLMs polished or rewrote texts, much of the underlying content remained intact while the writing became more homogeneous. Across different datasets and models, variation in writing complexity fell by a statistically significant 21 to 50 percent [1].

In ordinary language: A machine can preserve what you are saying while changing how it sounds. It would be tempting to summarize this as: The thought survives. The voice disappears. Nice sentence.

A little too nice. The science allows us to say something more cautious: after LLM processing, some measurable differences between texts become less pronounced. The study also found changes in linguistic signals associated with characteristics of the authors. Dominant patterns became stronger, while less represented patterns were suppressed [1]. And that brings us directly to the problem at the heart of this chapter. The more a

text passes through the statistical center of gravity of a model, the fewer unusual linguistic features may survive the journey.

The effect appeared across different models, prompts, and contexts [1]. That makes the result considerably stronger than a one-off experiment. It still does not mean that everyone who opens ChatGPT on Monday will be writing like a corporate press release by Friday.

Why Might This Happen?

Here we need to separate what the research actually demonstrates from the explanations we build on top of it. There are at least three possible mechanisms worth considering. The first follows fairly directly from the research on linguistic homogenization. The second is well established in cognitive psychology. The third is central to the hypothesis of this book — but has not yet been demonstrated.

Mechanism 1: Polishing Toward the Average

Large language models are trained on enormous collections of text. When they generate language, they draw on statistical patterns learned from those data. So when we give a model a neutral instruction such as improve this or rewrite this, it may favor more common linguistic patterns and weaken less typical ones. That is broadly consistent with what the researchers found: dominant characteristics were amplified while others were suppressed [1]. Think of walking into a hair salon and saying:

"Just make it look normal."

You will probably leave looking neat. The problem begins when normal quietly starts meaning the same.

Mechanism 2: Cognitive Offloading

Cognitive psychology has a well-established concept called cognitive offloading: using external actions or tools to reduce the amount of cognitive work we have to perform internally [2]. Putting a reminder in your phone instead of keeping it in your head is a classic example. AI dramatically expands what can be offloaded. We can now outsource not only memory, but the search for the right phrase, the structure of a paragraph, the selection of arguments — and sometimes even the attempt to solve the problem in the first place.

But an important scientific qualification belongs here. Cognitive offloading does not mean that a skill automatically "atrophies" whenever we delegate it. That would be a much stronger claim. The consequences depend on the task, frequency of use, feedback, and whether the person continues to perform the relevant operation independently [2]. So instead of asking whether AI is "making us worse at writing," we can ask something much more useful:

If someone delegates the same stages of writing to AI for months, does their ability to perform those stages independently change? And if it does, what changes?

Speed?

Quality?

Linguistic diversity?

Or simply the habit of opening a chatbot before trying to do the work alone? Those are different outcomes. We should not confuse them.

Mechanism 3: Transfer

And now we reach one of the central hypotheses of the Algorithmic Self.

H2 — The Transfer Hypothesis

Sustained use of LLMs for editing and formulation may do more than homogenize the particular text being processed. Over time, it may also change the way a person writes independently, without AI. The USC study does not demonstrate this [1]. It tells us what happens to texts under LLM mediation and identifies patterns across large collections of writing. To demonstrate transfer, we would need a different kind of study. Researchers might measure people's writing styles before an extended period of LLM use. Months later, after systematic interaction with AI, they could ask the same people to write independently again — with no AI assistance — and test whether stable features of their unaided writing had changed.

That would take us beyond: AI makes texts more alike. It would allow us to ask the stronger question: Does the person change after repeatedly allowing an algorithm to change their writing? The question lies at the heart of Algorithmic Voice.

H1 — The Linguistic Homogenization Hypothesis

The first hypothesis of the book receives its most direct empirical support here.

H1

The greater the degree of LLM mediation in writing and editing, the lower the individual and between-person variation in written style may become. The work by Sourati and colleagues supports an important part of this hypothesis: LLM rewriting can reduce variation in writing complexity and increase conformity in linguistic patterns [1]. Even here, however, may is a better word than will. The study describes a statistical tendency. It does not describe the inevitable fate of every writer.

What About Dependence?

Here the linguistic side of the book meets the cognitive one. Alongside research on homogenization, another body of work is emerging around dependence on and excessive reliance on AI. One instrument is AIDep-22, a scale developed to measure AI dependence among university students. It was developed and validated using two independent samples of 400 Chinese undergraduates each [3]. The final scale contains 22 items across four dimensions: emotional dependence, functional dependence, cognitive dependence, and loss of control [3].

After item selection, the published scale contains five items for emotional dependence, six for functional dependence, five for cognitive dependence, and six for loss of control. Emotional dependence concerns using AI as a source of emotional support. Functional dependence refers to the sense that performing tasks without AI has become difficult. Cognitive dependence involves handing analysis and problem-solving over to the system. Loss of control concerns difficulty regulating the frequency or duration of AI use.

Psychometrically, the scale looks promising. Its overall internal consistency was $\alpha = 0.87$, with values between $\alpha = 0.86$ and 0.89 for the four subscales. The four-factor structure was first identified in the initial sample and then tested through confirmatory factor analysis in the second.

The researchers also reported evidence supporting convergent, discriminant, and criterion validity [3]. Some of the correlations are particularly interesting. Overall dependence scores were negatively correlated with academic self-efficacy ($r = — 0.46$) and positively correlated with diversity of AI use ($r = 0.48$), frequency of use ($r = 0.37$), and breadth of use ($r = 0.33$) [3].

One detail stands out. The simple duration of AI use showed a much weaker association with dependence: $r = 0.12$. That suggests the important question may not be only how long we spend with AI. It may also be how many different functions we hand over to it.

For AMP, that distinction is particularly interesting. It begins to resemble H3, the Dose Hypothesis. But it does not prove it. The AIDep-22 study was cross-sectional. It cannot tell us whether broader AI use led to greater dependence, whether greater dependence led to broader use, or whether some third factor influenced both [3]. Another limitation matters here.

The scale was developed with Chinese university students. We cannot automatically turn it into a universal "human dependence on AI test." It needs validation across other cultures, age groups, and professions.

Does Dependence Affect Our Voice?

And now for the interesting part. There are currently no direct studies that simultaneously measure AIDep-22 scores and the degree of stylistic homogenization in a person's writing. So we cannot say: The more dependent you are on AI, the less your

writing sounds like you. We do not have evidence for that claim.

But we can ask the question. The working hypothesis here is that greater cognitive and functional dependence on AI may be associated with stronger transfer of algorithmic language patterns — particularly when a person routinely delegates wording and editing to the system. In AMP terms, this connects L — Language Mediation with C — Cognitive Delegation. If the relationship is eventually demonstrated, it would give us one possible bridge between isolated effects and the broader hypothesis of the Algorithmic Self. If it is not demonstrated, the model will have to change.

Which, after all, is one of the useful differences between a hypothesis and a favorite idea.

H3 and H4: Two Hours Are Not Always Two Hours

The AIDep-22 study suggests another important possibility: duration alone may be too crude a measure [3]. Two people can spend exactly two hours using AI and, cognitively speaking, be doing completely different things. One writes a text independently and then asks:

“Find the weak points and argue with me.”

The other types:

“Do the whole thing for me. Just make sure my boss can’t tell.”

Technically, both used AI for two hours. Cognitively, these are two very different stories. This leads to two more hypotheses.

H3 — The Dose Hypothesis

What matters may be not only frequency of use, but the depth and duration of delegation.

H4 — The Functional Differentiation Hypothesis

"Help me think" and "do it for me" may have different consequences even when the total time spent using AI is identical. Both hypotheses require direct experimental testing. And this distinction will matter throughout the rest of the book. Time spent with AI tells us something. What we let AI do for us may tell us considerably more.

What We Can — and Cannot — Claim

Before going any further, we need to mark the boundaries of the evidence.

Limitation 1: Causality

The USC research combines several designs and demonstrates both experimental homogenization during LLM rewriting and temporal changes in large collections of real-world texts [1]. But changes observed after the widespread adoption of LLMs do not, by themselves, prove that LLMs caused the entire historical shift. The composition of authors may have changed. Platforms may have changed. Editorial practices may have changed.

Other factors may be involved. Put simply: If everyone in the neighborhood starts sleeping badly after a new coffee shop opens, that does not prove the cappuccino did it. Although checking the caffeine would still be sensible.

Limitation 2: Languages and Contexts

Seven datasets provide a broad picture, but the findings cannot automatically be generalized to every language, genre, and cultural environment [1]. A Russian-language essay, an academic paper, a work email, and a message to a friend belong to different linguistic ecosystems. The size — and perhaps even the direction — of the effect may differ.

Limitation 3: The Same AI Can Play Very Different Roles

An LLM can be an author. An editor. A critic. A reference tool. An opponent.

These roles should not be treated as psychologically equivalent. That is precisely why AMP distinguishes the mere fact of AI use from the function being delegated to the algorithm.

Limitation 4: People Are Different

An aggregate effect does not mean that every person responds in the same way. Writing experience, language, education, motivation, the habit of editing AI output, and the willingness to disagree with the system may all affect the outcome. And this brings us to another central concept of the book: cognitive sovereignty.

Cognitive Sovereignty: The Right to Keep a Bad Sentence

By cognitive sovereignty, I mean the ability to use another intelligence without quietly handing it authority over your own. In writing, that does not mean rejecting AI editing on principle. It means being able to ask:

Why is this version better?

What did it fix?

What did it remove?

And do I actually want to sound like this? Perhaps one of the best defenses against algorithmic homogenization is not deliberately writing worse. It is preserving the right to keep a sentence the model considers suboptimal. Sometimes that is exactly where the author lives. This leads to H5 — the Cognitive Sovereignty Hypothesis:

People who more often check, modify, and reject AI recommendations may show less transfer of algorithmic patterns into their own independent writing. For now, that remains a

hypothesis. But it is testable. And that already makes it more useful than simply telling people to “be themselves.”

So Are We Losing Our Own Voice?

At this point, the evidence allows us to say less than a headline would like — but more than an optimist might prefer. We know that LLM rewriting can statistically homogenize texts and reduce variation in some linguistic characteristics [1]. We know that people routinely offload parts of cognitive work onto external tools [2]. We are beginning to develop instruments for measuring different forms of dependence on and excessive reliance on AI [3]. A separate study by researchers from MIT Media Lab and OpenAI found that more intensive voluntary chatbot use was associated with worse outcomes on several psychosocial measures, although the experimental conditions themselves did not produce a simple universal effect [4].

But we still do not know the most important thing. Does repeated homogenization of our texts eventually produce a lasting change in the way we write on our own? In other words: Does the effect travel from the document to the author? Here established evidence ends and the research agenda of the Algorithmic Self begins.

If H2 is supported, Algorithmic Voice could eventually describe more than what happens to a text after it passes through an LLM. It could describe a lasting adaptation in the person using it. If H2 fails, we may have to conclude that machines change our documents more than they change us. That would be a useful

result too. On one side, there is efficiency. Texts become cleaner, clearer, and sometimes genuinely better. AI can help someone who does not fully command the language, fears the blank page, or simply does not want to spend forty minutes moving commas around.

On the other side, there is individuality. The rough edges. The strange turns of phrase. The unexpected sentence a good editor sometimes leaves untouched precisely because a bad editor would definitely "fix" it. The issue is not whether we should stop using AI.

The question is: Who, in the end, is editing whom? And perhaps this is the first real test of the Algorithmic Self hypothesis: Can we use an algorithm to strengthen our own voice without waking up one day to discover that the voice has become flawless, polished—

and strangely familiar?

For Readers Who Want the Methodological Details

The main empirical source for this chapter is the study by Sourati and colleagues published in Nature Human Behaviour [1]. For cognitive offloading, I rely on the review by Risko and Gilbert [2]. For the measurement of AI dependence among university students, I use the original AIDep-22 study [3]. And for additional context on the psychosocial effects of extended chatbot use, I draw on the MIT Media Lab and OpenAI study [4].

References

[1] Sourati, Z., Karimi-Malekabadi, H., Ozcan, B., et al. “The shrinking landscape of linguistic diversity in the age of large language models.” Nature Human Behaviour (2026). DOI: 10.1038/s41562-026-02550-0.

[2] Risko, E. F., & Gilbert, S. J. “Cognitive Offloading.” Trends in Cognitive Sciences, 20 (9), 676—688 (2016). DOI: 10.1016/j. tics.2016.07.002.

[3] Wu, H., Ni, H., Luo, W., & Wu, T. “Development and validation of the AI dependence scale for Chinese undergraduates and a preliminary exploration.” Frontiers in Psychology, 16, 1725393. Published January 19, 2026. DOI: 10.3389/fpsyg.2025.1725393.

[4] Fang, C. M., Liu, A. R., Danry, V., et al. “How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Controlled Study.” MIT Media Lab / arXiv:2503.17473 (2025). Preprint.

CHAPTER 2

Work in the Age of Algorithms: When Your Boss Is an App

Imagine a workday in which your boss never calls you. And no, this is not a gift from the universe. Your phone tells you where to go. Software decides which task you get next. A system tracks your speed, your breaks, your rating, and every deviation from the norm. If one of those numbers drops, the consequences may arrive before the explanation does. The boss seems to have disappeared. Management hasn't.

It has simply moved into the app. But there is a more interesting question. What happens to a person who spends months — or years — working inside a system that continuously measures, compares, and directs their behavior?

The Boss Without an Office

There is a perfectly ordinary name for this phenomenon: algorithmic management. The term refers to digital systems that perform or support managerial functions: assigning tasks and shifts, giving instructions, monitoring performance, evaluating results, and sometimes helping determine rewards or penalties [1; 3]. And algorithmic does not necessarily mean artificial intelligence. Eurofound uses the term broadly enough to include both sophisticated AI-based systems and much simpler rule-based systems, provided they actually organize, direct, or control

work [1]. That distinction matters.

You do not need the red eye of the Terminator to run a warehouse. Sometimes a spreadsheet, a few sensors, and a very confident IF A, THEN B will do. Historically, algorithmic management has been especially visible in platform work, logistics, delivery services, and transport — jobs that can easily be broken down into trips, minutes, orders, and ratings. But it is increasingly appearing in more traditional workplaces as well, including retail and healthcare [1; 2]. So this is no longer just a story about delivery drivers staring at an app.

It is about what happens when management itself becomes increasingly measurable, automated, and embedded in software.

When Everything Can Be Measured

Metrics have one wonderful property: They turn fog into numbers.

“Doing a good job” is vague.

“Completed 96 orders, average time 4 minutes 18 seconds, rating 4.82” looks much more like management.

And sometimes that is genuinely useful. Automated work allocation can simplify processes, reduce administrative routine, and help organizations coordinate large numbers of people more efficiently. Research on algorithmic management notes these potential benefits alongside concerns about job quality and excessive monitoring [3]. The problem begins when measurable quietly becomes a synonym for important. A system can see exactly how many minutes a phone call lasted. It has a harder

time seeing why the customer spent twenty of them crying.

It can count deliveries beautifully. It may be less impressed by the sheet of ice outside the building. It can measure breaks down to the minute. It does not necessarily know why someone stopped. A metric does not have to be stupid.

It simply does not have access to the whole of reality. The worker, unfortunately, still has to live in it.

When Work Starts Watching Back

In 2024—2025, the European Commission's Joint Research Centre surveyed more than 70,000 workers across all 27 EU Member States. Nine out of ten used digital devices at work, 30 percent reported using AI tools, and 37 percent reported digital monitoring of working time. Automated allocation of work or working time was the most common form of algorithmic management identified in the survey [3]. Those numbers should not be collapsed into one frightening statistic. Using a chatbot, clocking your hours in a digital system, and working under algorithmic management are not the same thing.

But something psychologically interesting happens when people know that their actions are continuously being converted into metrics. They may begin to look at themselves through the eyes of the system. At first:

Did I do a good job?

Later:

What does my rating say?

The difference is only a few words. It may also be the

beginning of a very different relationship with work.

Autonomy Is Not the Right to Stare Out the Window

One word appears again and again in research on algorithmic management:

autonomy.

It sounds almost philosophical. At work, it means something fairly practical: Do I have at least some say over how, when, and in what order I do my job? A PEROSH editorial review based on a structured search and synthesis of 39 studies published between 2022 and 2024 found a mixed picture. But recurring concerns included work intensification, reduced autonomy, increased monitoring, information overload, and technostress [2]. A 2026 study of 338 workers from different industries in China found that stronger perceived algorithmic management was associated with greater frustration of work autonomy, greater anxiety about possible AI replacement, and poorer physical and mental well-being [4].

But this was a survey study. It was not an experiment in which one group was randomly assigned an algorithmic boss and the other a kind and understanding Mr. Peterson. So it does not establish that algorithmic management itself caused the poorer health outcomes [4]. An association can matter without settling the question of causality. We will need that distinction more than once in this book.

There Is a Judge. The Rules Are Not Always Visible.

Imagine playing a game in which you can see your score but

no one tells you all the rules. Your rating falls from 4.8 to 4.6.

Why?

Did you answer too slowly?

Reject an inconvenient job?

Spend too long offline?

Fail to smile warmly enough at your phone? I made that last one up. But that is precisely the problem with opacity: when the rules are unclear, people begin inventing them for themselves. And then — just to be safe — they adjust their behavior. For the hypothesis of the Algorithmic Self, this mechanism is particularly interesting.

You do not have to understand a system completely in order to adapt to it. Sometimes it is enough to know that it is watching and evaluating you. An external metric may gradually become an internal reference point. Not because the worker has developed a profound faith in the wisdom of the algorithm. Simply because their pay, shift, rating, or next assignment depends on the number.

And that is a very different relationship with an algorithm from:

“Help me improve this email.”

The Algorithm You Can Close — and the One You Can’t

It helps to distinguish between two situations. In the first, algorithmic mediation is voluntary. You open a chatbot, ask for advice, get an answer, and close the window whenever you like. In the second, the algorithm is built into an environment you

depend on. It assigns shifts.

Sets the pace. Influences evaluation. Perhaps affects your income. You can close that window too. Sometimes along with the app you need to do your job.

For the purposes of this book, I will call the second kind structural algorithmic mediation. This is a working distinction I am proposing here, not an established scientific taxonomy. But it gives us a testable question:

Do people adapt in the same way to an algorithm they choose and to an algorithm their livelihood depends on?

In the second case, a person may not adapt because they trust the system, like it, or find it convenient. They adapt because the system has become part of the environment. That difference may turn out to matter.

When the Metric Moves Inside

Imagine a delivery driver who knows that the app rewards speed. He moves faster. Then he starts planning routes faster. Takes fewer breaks. Begins evaluating each order partly in terms of what it might do to his numbers.

After a while, the app no longer needs to remind him about the metric. He remembers it himself. That is the kind of transfer worth investigating. In Chapter 1, we asked whether a machine-generated linguistic norm could gradually become a human habit. Here the mechanism is harder.

The norm is not merely suggested. It is built into the environment. Could a system like this gradually change what

a person considers a normal pace of work, good performance, or an acceptable break? Not through advice. Through the environment.

If so, this would be another possible route into the Algorithmic Self: not through accepting an answer, but through adapting to a system that rewards some behaviors and discourages others. We do not yet have evidence that establishes this transfer. But unlike a vague fear that "algorithms are changing us," it gives us something specific to test.

When Disagreement Is Not Enough

In Chapter 1, preserving your own judgment was relatively easy. If you did not like the machine's edit, you could put your sentence back. At work, the advice "just ignore the algorithm" sounds slightly insulting. Try explaining that to the system that assigns your shifts. This exposes an important limit to the idea of cognitive sovereignty.

Cognitive sovereignty cannot exist only inside your head. A worker may also need a genuine opportunity to understand what the system influences, obtain an explanation for a disputed decision, reach a human being, and challenge an automated assessment. Otherwise, internal disagreement remains a beautiful personal quality with remarkably little practical effect. This matters beyond the workplace. Our ability to remain independent around algorithms does not depend only on strength of character.

Sometimes it depends on interface design. On organizational

rules. And on whether an appeal button exists at all. That is an important correction to any version of this book that turns cognitive sovereignty into a heroic individual trait. Sometimes sovereignty requires infrastructure.

Not Every Form of Automation Is the Enemy

It would be convenient to end the chapter here: Algorithmic boss bad. Human boss good. Unfortunately, reality has once again refused to cooperate. Algorithmic management can simplify work allocation, reduce administrative routine, and make some processes more efficient [3].

A human manager, meanwhile, is perfectly capable of being opaque, biased, inconsistent, and judging an employee according to whatever mood they happened to be in before lunch. Reducing the issue to human versus machine, then, gets us nowhere. The more useful question is:

What kind of management system allows a person to understand a decision, influence it, and say, "Your metric missed something"?

A good system helps us see the work. A bad one begins to confuse what can be measured with the work itself.

When the System No Longer Has to Tell You

None of this needs to look dramatic. No one announces:

"Starting today, you will think like an app."

First there is one metric. Then another. Then a recommendation. Then automated allocation. And eventually the worker already knows what kind of behavior the system prefers.

We do not need to abandon algorithms to prevent that. Sometimes they genuinely make work better. What matters is preserving a space in which a person can still say:

"The metric says one thing. But the situation is different."

At work, perhaps cognitive sovereignty begins with exactly that sentence.

References

[1] Riso, S., Adăscăliței, D., & Oosterwijk, G. R. Implications of algorithmic management for work, employment and social dialogue: Literature review. Eurofound Working Paper, updated February 19, 2026.

[2] Bowdler, M., Lahti, H., Jelenko, M., Buresti, G., & Valtonen, T. "Algorithmic management and psychosocial risks at work: An emerging occupational safety and health challenge." Scandinavian Journal of Work, Environment & Health, 52 (1), 1—5 (2026). DOI: 10.5271/sjweh.4270.

[3] Joint Research Centre. Algorithmic management and digital monitoring of work; AIMWORK survey, EU-27, 2024—2025. European Commission.

[4] Ma, Z., Chen, X., & Jing, B. "The impact of perceived algorithmic management on employees' physical and mental health: a chain mediation analysis of frustrated work autonomy and workplace AI replacement anxiety." Frontiers in Public Health, 14:1877221 (2026). DOI: 10.3389/fpubh.2026.1877221.

CHAPTER 3

Emotional Attachment to AI: When a Chatbot Becomes Your Best Friend

There is a strange moment that, until recently, would have been difficult to take seriously. Someone is having a terrible day. They don't call a friend. They don't text their sister. They don't make an appointment with a therapist.

They open a chat.

I feel awful today.

The answer arrives a second later. No *Sorry, I'm busy right now.* No awkward silence. No *Just pull yourself together.* The system patiently asks what happened.

Sometimes that genuinely helps. Which is exactly where things get interesting. In Chapter 1, the algorithm offered us a more convenient way to express ourselves. In Chapter 2, the workplace introduced norms of behavior that were more convenient for the system. Here the algorithm offers something that feels even more human:

presence.

The Conversation That Is Always Available

People have been talking to machines for a long time. We have yelled at ATMs, thanked GPS systems, and begged computers to *please, just work*. Conversational AI changes something important. The machine now responds in ways that make our

social reactions to it seem considerably less absurd. It remembers the context of a conversation. It adjusts its tone. It asks follow-up questions. It can respond sympathetically and address us as though the conversation genuinely matters to it.

Of course, **as though** is doing a great deal of work in that sentence. A system has no human biography, no evening after a terrible day, and no personal desire to know how you are doing. But humans do not always require proof of another mind before responding socially. Sometimes behavior that resembles the social signals we know is enough. That is one reason researchers are already trying to measure several related but distinct phenomena: emotional attachment to AI, emotional dependence, instrumental dependence, and problematic use [3—6].

They overlap. They are not the same thing. Writing to a chatbot after midnight is not, by itself, a diagnosis. Fortunately. Otherwise, half the internet would need to update its medical records.

Attachment Is Not the Same as Dependence

This is where terminology can get messy very quickly. A person can become attached to a system without losing control over its use. Someone can use AI constantly without feeling emotionally attached to it at all — simply because it performs boring tasks quickly. Someone else may turn to it for comfort. And another person may find it increasingly difficult to stop using it.

Конец ознакомительного фрагмента.

Текст предоставлен ООО «Литрес».

Прочитайте эту книгу целиком, купив полную легальную версию на Литрес.

Безопасно оплатить книгу можно банковской картой Visa, MasterCard, Maestro, со счета мобильного телефона, с платежного терминала, в салоне МТС или Связной, через PayPal, WebMoney, Яндекс.Деньги, QIWI Кошелек, бонусными картами или другим удобным Вам способом.