THE CORE MESSAGE: THE TRAJECTORY OF AI IN EDUCATION IS NOT INEVITABLE; IT'S A MATTER OF STRATEGIC LEADERSHIP, CURRICULUM DESIGN, AND ETHICAL PRACTICE.

BEYOND THE AUTOMATION ABYSS

LEARNING AGENCY VS. AI AGENTS IN EDUCATIONAL FUTURES

ABSTRACT

As education enters the era of agentic AI where tools act autonomously on behalf of students, education leaders must grapple with both unprecedented opportunities and systemic risks. Drawing on two decades of teaching experience and empirical research, I have framed a critical challenge: the automation abyss. It is a growing divide between learners who retain agency and those who rely entirely on automated systems. Using techniques used in strategic foresight and future studies, I outline plausible futures for education, the cognitive implications to learners of AI overuse, and curricular implications for educators.

The thought experiment, tied to real data, reveals that while AI adoption in education is accelerating rapidly, institutions can shape outcomes through strategic curriculum design and risk management. By implementing human checkpoints, redesigning learning to preserve inquiry, and establishing ethical frameworks for agentic AI, educational leaders can navigate this transition successfully. The core message: the trajectory of AI in education is not inevitable; it's a matter of strategic leadership, curriculum design, and ethical practice.

ABOUT THE AUTHOR

ERIC HAWKINSON

Eric is a learning futurist, tinkering with and designing technologies that may better inform the future of teaching and learning. Eric's projects have included augmented tourism rallies, AR community art exhibitions, mixed reality escape rooms, and other experiments in immersive technology.

Eric Hawkinson
LEARNING FUTURIST

erichawkinson.com

Roles

Professor – Kyoto University of Foreign Studies
Research Coordinator – MAVR Research Group
Founder – Together Learning
Developer – Reality Labo
Community Leader – Team Teachers
Chair – World Learning Labs

CORE VALUES

- Open Knowledge - Free and open access to information is a foundation to a productive modern life, connected to ideas of the open web and platform agnosticism.
- Privacy by Design - Business models are increasing moving toward supporting revenue by collecting, curating, and trading behavioral surplus through technology. These models should be tempered with safety, ethics, and privacy concerns and designed as such.
- Digital Literacy for All - An informed public about the use of technology is key for a responsible and engaged digital society.

PASSION PROJCTS

Together Learning
A community of technology minded learners. Exploring human potential... together.

Reality Labo
Augmented reality enhanced learning environments and mixed reality rapid prototyping.

Corefol.io
Showcase the Core of your Skills with AI Assisted Student Portfolios

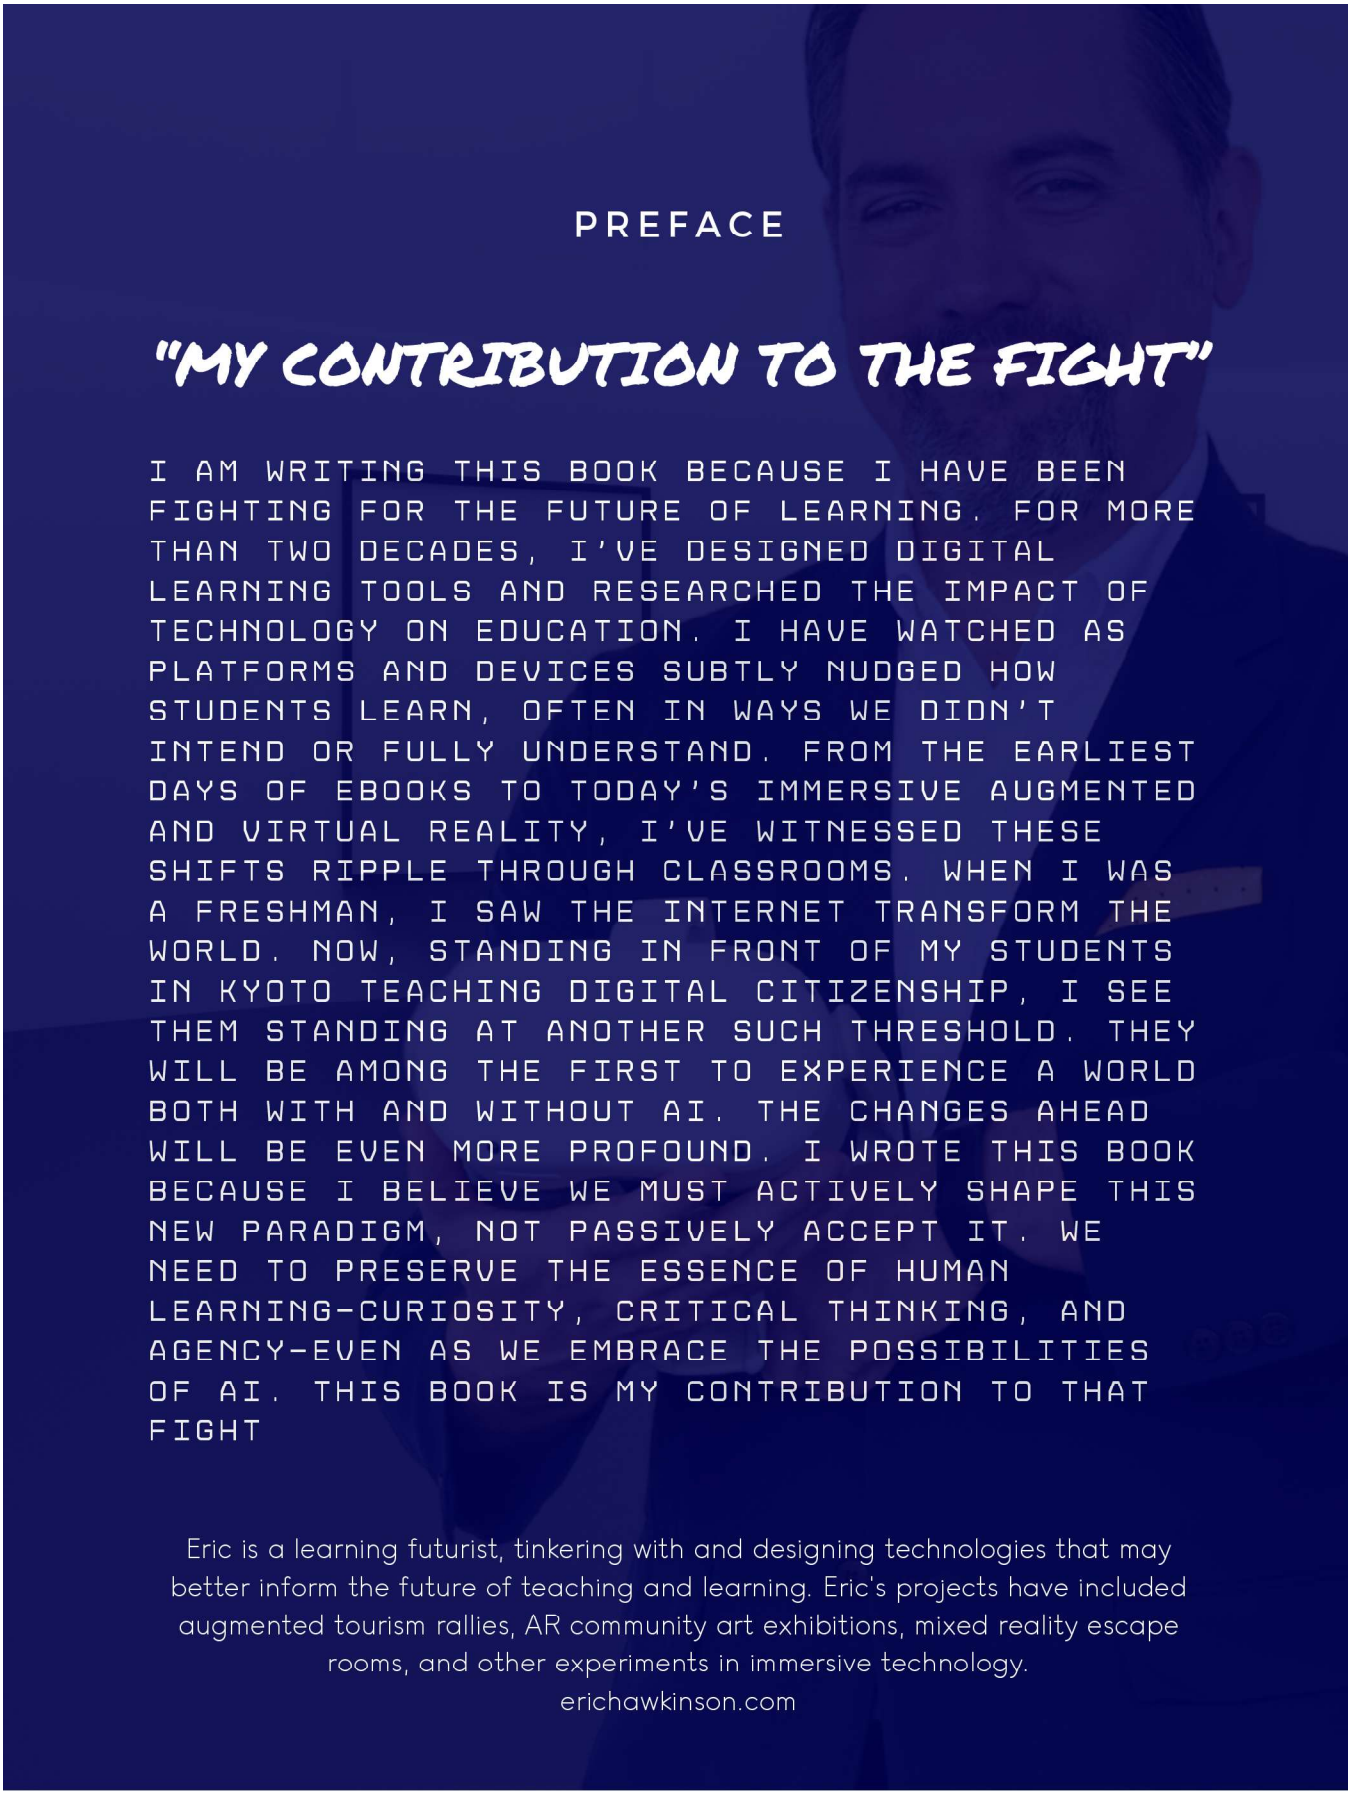

PREFACE

"MY CONTRIBUTION TO THE FIGHT"

I AM WRITING THIS BOOK BECAUSE I HAVE BEEN FIGHTING FOR THE FUTURE OF LEARNING. FOR MORE THAN TWO DECADES, I'VE DESIGNED DIGITAL LEARNING TOOLS AND RESEARCHED THE IMPACT OF TECHNOLOGY ON EDUCATION. I HAVE WATCHED AS PLATFORMS AND DEVICES SUBTLY NUDGED HOW STUDENTS LEARN, OFTEN IN WAYS WE DIDN'T INTEND OR FULLY UNDERSTAND. FROM THE EARLIEST DAYS OF EBOOKS TO TODAY'S IMMERSIVE AUGMENTED AND VIRTUAL REALITY, I'VE WITNESSED THESE SHIFTS RIPPLE THROUGH CLASSROOMS. WHEN I WAS A FRESHMAN, I SAW THE INTERNET TRANSFORM THE WORLD. NOW, STANDING IN FRONT OF MY STUDENTS IN KYOTO TEACHING DIGITAL CITIZENSHIP, I SEE THEM STANDING AT ANOTHER SUCH THRESHOLD. THEY WILL BE AMONG THE FIRST TO EXPERIENCE A WORLD BOTH WITH AND WITHOUT AI. THE CHANGES AHEAD WILL BE EVEN MORE PROFOUND. I WROTE THIS BOOK BECAUSE I BELIEVE WE MUST ACTIVELY SHAPE THIS NEW PARADIGM, NOT PASSIVELY ACCEPT IT. WE NEED TO PRESERVE THE ESSENCE OF HUMAN LEARNING–CURIOSITY, CRITICAL THINKING, AND AGENCY–EVEN AS WE EMBRACE THE POSSIBILITIES OF AI. THIS BOOK IS MY CONTRIBUTION TO THAT FIGHT

Eric is a learning futurist, tinkering with and designing technologies that may better inform the future of teaching and learning. Eric's projects have included augmented tourism rallies, AR community art exhibitions, mixed reality escape rooms, and other experiments in immersive technology.
erichawkinson.com

SCAN

MORE RESERACH ONLINE

Innovating Research Publication

Together Research champions the spirit of co-creation without the constraints of traditional disciplinary boundaries. By deliberately eschewing a fixed list of research areas, we foster a uniquely flexible environment that encourages scholars from varied fields to collaborate and innovate together.

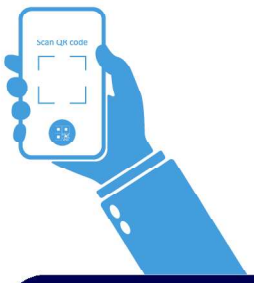

togetherlearning.com/research

BEYOND THE AUTOMATION ABYSS

LEARNING AGENCY VS. AI AGENTS IN EDUCATIONAL FUTURES

Chapters

togetherlearning.com

— CHAPTER 1 —

Agentic AI
Agency for who?

The pencil gave way to the keyboard, for the most part. The library surrendered to the search engine, for the most part. Moving forward, will the proverbial thinking cap get replaced by a thinking app, for the most part? The quiet moment before a student tackles a problem likely now begins with a prompt, not a pause. A student faced with complexity reaches for ChatGPT, Claude, Gemini, or another large language model rather than struggling through fundamentals. Zooming out, learning institutions are starting to deploy these AI systems to automate feedback, assessment, and create new course materials. So now an unsettling scenario is possible, in which a teacher uses AI to create coursework, the student uses AI to complete that coursework, and the university uses AI to accept, score, and give formative feedback to students. Soon, these systems will anticipate and take actions on our behalf, possibly exacerbating this trend from using tools to being tools. Agentic AI is what all the big tech firms are touting, seen as the next iteration of the Internet, where swarms of AI entities intermediate human interaction with the world's technology infrastructure. At this level of automation, what should educators teach or why should students learn? These questions hit at the heart of this concept of the 'Automation Abyss', a state of teaching and learning that pits big tech's

vision of AI as an intermediary between everyone and everything, and educators and students who are already struggling to know when and if they can assert self-determination, or when to take the convenience of AI assistance.

Over two decades of teaching with technology, I have observed this pattern intensifying, marking a shift from using technology as a supportive tool to employing it as a substitute for essential learning processes. How much thinking gets outsourced, especially as it pertains to how we teach and learn, is what this book explores. On June 1, 2009, Air France Flight 447 plunged into the Atlantic Ocean. All 228 people aboard died. The cause? A perfect storm of automation and human skill decay. The pilots, so accustomed to automated systems, had lost the fundamental ability to fly their aircraft when those systems failed. Today, in classrooms worldwide, we are creating an educational version of this same tragedy. Students reach for ChatGPT, Claude, or Gemini instead of wrestling with problems themselves. They're losing not just the ability to write or research, but something more fundamental: the capacity to think through difficulty, to discover the unexpected, to grow through struggle. This is the Automation Abyss… and we are already falling into it.

This concept follows the well-documented phenomenon known as the paradox of automation: the more automated systems become, the more skilled human operators must be to manage them when they fail. Yet the more these systems are used, the fewer opportunities humans have to develop and maintain those essential skills. The very automation designed to make flying safer (autopilot) had created a generation of pilots unprepared for moments when human skill was most needed. The Federal Aviation Administration classified this as a "loss of control" event, a category that has stubbornly persisted even as other types of aviation accidents have declined. Yet automation need not lead inevitably to dependency and disaster. Thirty-nine years earlier, Apollo 13's Commander Jim Lovell faced an equally desperate

situation when guidance computers failed after an explosion crippled their spacecraft. Unlike the Air France pilots, Lovell had been trained with the assumption that automation would fail. He performed complex orbital mechanics calculations manually, using nothing more sophisticated than a slide rule and verified by ground control with pencils and paper. Those manual calculations, executed under unimaginable pressure, brought three men safely home. The difference between tragedy and triumph lay not in the technology, but in how humans were prepared to function when that technology inevitably fell short.

Recognizing these parallels between aviation automation and educational technology, I developed what I call the "Budding Botanist Paradox" to illustrate how this same dynamic operates in a learning environment. This came mostly from my work integrating augmented and virtual reality into various educational contexts. The paradox looks at how technology is nudging the fundamental process of inquiry at all levels. Traditional botany studies require students to observe carefully, ask systematic questions about plant characteristics, and develop analytical skills necessary to identify species through direct engagement with nature. When augmented reality systems can instantly identify plants with computer vision and instantly provide comprehensive information, students may achieve task completion without developing the fundamental inquiry skills that true understanding requires. The more automated the learning process becomes, enhanced with immersive technology, the more learners become dependent on these automated systems for basic learning, more susceptible to algorithmic influence, and more vulnerable when the systems fail or provide inadequate guidance. The symptoms of this dependency are already emerging across multiple domains, even into the simple skills of social interaction. Roughly one in three American adults now feels lonely every week, with young adults reporting the highest rates, social media is playing a big part in this, but AI chatbots

are perhaps becoming an even bigger problem. As traditional social bonds fray, many turn to AI companions for emotional support, spending hours chatting with AI personas. When I discuss AI companionship with students, many have not considered that these digital relationships might be making them lonelier in the long term. They have automated away not just intellectual struggle, but human connection itself.

Meanwhile, institutional responses reveal troubling misunderstandings of the challenge we face. Universities deploy AI detection systems with false positive rates too high for practical use, wrongly flagging innocent students while missing sophisticated AI use. They create elaborate "traffic light" policies categorizing assignments by permitted AI use levels, then wonder why large numbers of students report directly copying AI-generated content despite clear prohibitions. These institutions suffer from what researchers call the "enforcement illusion," borrowing the language of structural systems like actual traffic lights while lacking any meaningful enforcement infrastructure. The confounding reality is that we cannot simply ban these tools. Without mastery of AI systems, students may find themselves at a terrible disadvantage in an automated world. Yet the more we integrate these technologies into learning, the more we risk creating the very dependency we should be preventing. We face the same bind that confronts the aviation industry: the tools that make us more capable also make us more vulnerable. Yet as I write this introduction, the big tech companies are in a fervor to quickly disseminate tools that represent something far more profound than students using AI to complete assignments. These new systems allow AI agents to bypass traditional websites entirely, automatically querying databases, scheduling appointments, and coordinating complex tasks without human intervention. What we are witnessing is not merely the adoption of new educational tools, but the convergence of two simultaneous revolutions that amplify each other in ways we are only beginning to

understand. This agentic AI is the next wave of this paradigm shift in all aspects of modern life, and it calls into question the very meaning of education, what to study, and maybe even why to learn.

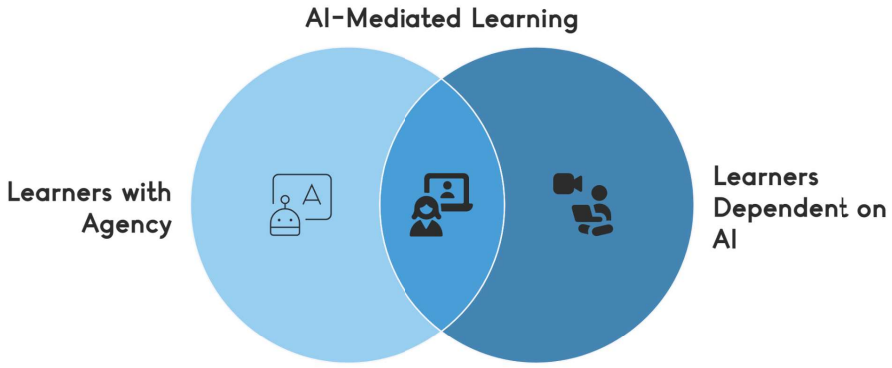

The traditional digital divide of access to technology is evolving into something more subtle and more dangerous: a divide between those who maintain agency over their relationship with technology and those who become dependent on AI-mediated experiences. This shift, accelerated by powerful language models and automated learning tools, raises crucial questions about the future of human learning and development.

I have termed this growing chasm the "Automation Abyss", a fundamental transformation in how humans learn and develop skills that threatens to separate learners into two distinct populations: those who direct AI and those who are directed by it.

The Receipts

In academia, we call this 'showing your work,' but I prefer to think of it as showing the receipts.

The Automation Abyss isn't just theoretical as controlled experiments reveal the exact trade-offs happening in student minds. When efficiency comes at the cost of memory and genuine learning, we can measure the cognitive price tag.

Reference numbers: 3, 4, 10

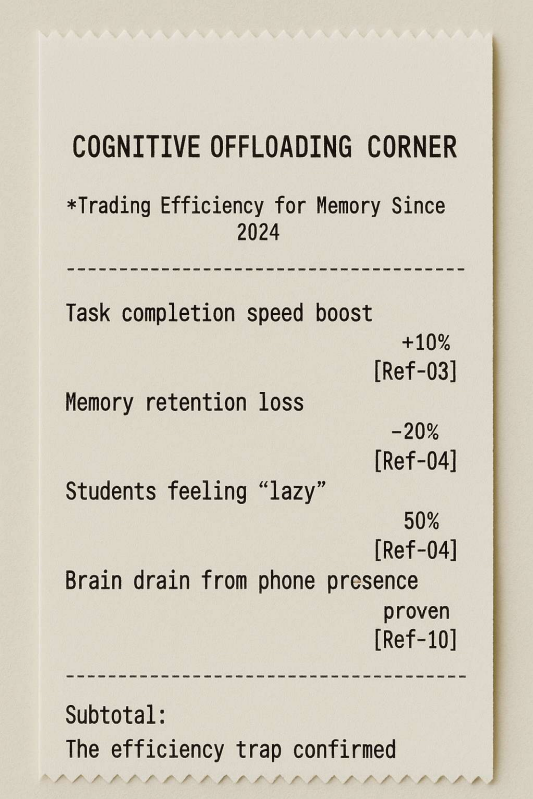

The Receipts *are evidence-based snapshots that appear throughout this book, presenting key statistics and findings from recent research. Think of them as quick-reference cards that distill complex academic studies into digestible facts. Each "receipt" is backed by peer-reviewed sources listed at the end of the book. This receipt captures the fundamental trade-off at the heart of the Automation Abyss: we're gaining speed but losing substance. The data reveals a troubling pattern where cognitive offloading delivers a modest 10% boost in task completion speed, but at the cost of a devastating 20% drop in memory retention. Half of students now report feeling "lazy" when relying on AI tools, while research confirms that even the mere presence of smartphones reduces our cognitive capacity. This is the efficiency trap in action, we're becoming faster but fundamentally less capable, trading our mental muscles for digital convenience. The Automation Abyss isn't just a theoretical concern; it's a measurable phenomenon already reshaping how students think, learn, and retain information.*

Defining the Automation Abyss: A New Digital Divide for the Age of Agentic AI

The Automation Abyss represents a fundamental shift in educational and societal inequality. It's one that moves beyond the traditional definition of the digital divide, but a more insidious separation between those who maintain agency over their learning and those who become dependent on AI-mediated experiences. Where the 1990s digital divide separated those with internet access from those without, creating disparities in information access and opportunity, the Automation Abyss divides humanity into two distinct populations: those who direct AI systems and those who are directed by them.

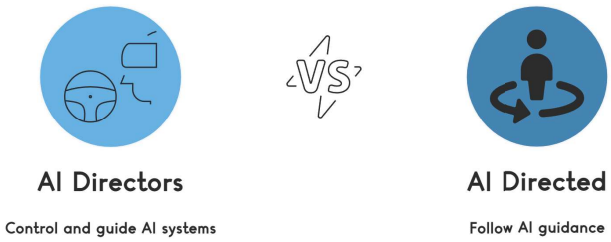

This isn't merely about academic dishonesty vs. the access to AI tools. It represents a fundamental change in how students conceptualize the boundary between their own intellectual work and machine-generated content. When paired with some findings of a significant negative correlation between AI tool usage and critical thinking abilities, we see the Automation Abyss taking shape: increased efficiency paired with decreased capability. The Automation Abyss manifests through what I term the "Budding Botanist Paradox." Traditionally, a botanist develops expertise through observation, questioning, and deduction, fundamental inquiry processes that build cognitive capacity. When AI systems provide instant plant identification and comprehensive information with the aid of augmented reality enhanced tools, these essential intellectual processes atrophy. The paradox deepens as

students may report feeling more confident in their learning while simultaneously losing the capacity for independent thought. AI users often admit feeling "lazy," and many recognize their growing dependency, yet usage continues to accelerate. This divide operates across three interconnected levels. At the cognitive level, students offload thinking processes to AI at unprecedented rates. At the behavioral level, AI substitutes for intellectual effort, with students now using AI for brainstorming and ideation. At the systemic level, entire information architectures are being reconfigured around AI agents rather than human users, as the platform economics incentivize big tech to race for AI systems to be an intermediary in every aspect of our lives.

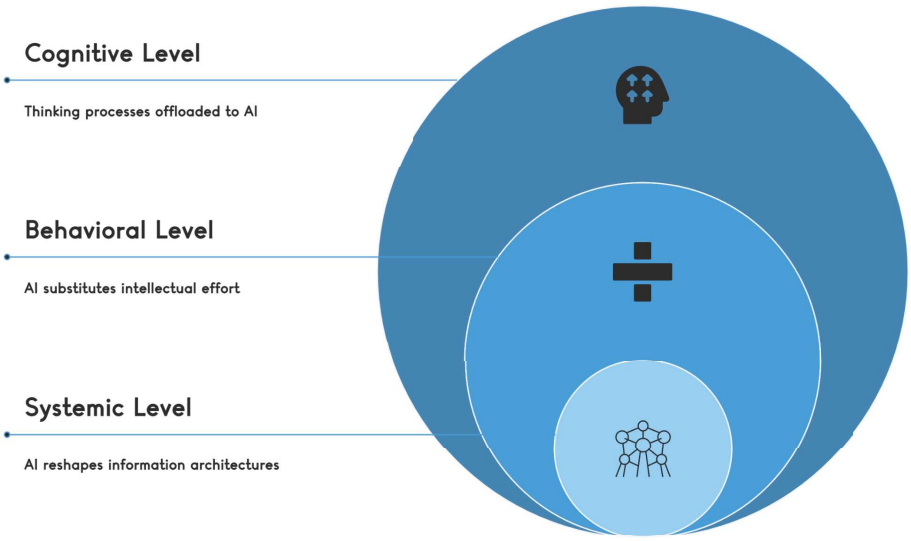

The Automation Abyss isn't inevitable... It's a choice. Unlike the digital divide, which required infrastructure investment to bridge, escaping the Automation Abyss demands intentional design of learning experiences that preserve human agency, critical thinking, and the productive struggle essential for intellectual development. The stakes couldn't be higher: we're determining whether future generations will be architects of intelligent systems or merely their subjects.


I'm telling you, these AI tools are great!
You don't even have to be good at drawing or writing!
But I want to be good at those things....
She doesn't get it
He doesn't get it

Two Paths Diverging
in the Automation Abyss

This comic captures the essence of the automation abyss in a single classroom moment. The boy's enthusiasm for AI tools that eliminate the need for foundational skills represents one path forward—immediate productivity and impressive outputs without developing underlying competencies. The girl's concern about wanting to be good at drawing and writing represents another path, understanding that skill development itself has value beyond the final product.

Neither student is wrong, but they embody the new digital divide this book explores. The original Digital Divide was about access to computers for learning. Today's automation abyss reveals a different pattern: while AI tools are widely accessible, the crucial advantage comes from having humans-in-the-loop during education. Students with access to skilled mentors learn to maintain agency while partnering with AI effectively, they develop the ability to direct rather than depend on AI systems.

Meanwhile, students without such human guidance may gravitate toward AI replacement rather than AI collaboration, potentially limiting their long-term skill development. The automation abyss doesn't fragment learning communities, it deepens existing inequalities through different relationships with technology and human capability.

The million dollar question is

How do we design educational experiences that harness AI's benefits while preserving students' motivation to develop their own skills and agency?