

AI-POWERED LEADER

An Integration Roadmap
for Human Teams

2026 EDITION

Written and Illustrated by
Fiona Passantino



About the Author



VibeCoding? SlashCommands? FLOPs? AI is confusing, constantly evolving, and often overwhelming. But it doesn't have to be.

TEDx speaker, AI4ALL Legacy "Woman of Influence in AI", host of the podcast Working Humans, and award-winning author of the Comic Books for Executives series, Fiona Passantino is a rising voice in the AI-Human inspirational space, headlining international events for some of the world's most powerful multinationals.

After fifteen years in corporate HR, communication and employee engagement for some of the largest European multinationals, she knows exactly what keeps people up at night and what genuinely moves organisations forward.

Fiona received an MBA with AI focus at the University of Amsterdam.

For fresh comics and latest articles:
www.working-humans.com



photo Sylvain-Munsch

AI-Powered Leader

Copyright © 2026 Fiona Passantino.

All rights reserved.

Published by Working Humans Press

www.working-humans.com

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means including photocopying, recording, or other electronic or mechanical methods without the prior written permission of the publisher, except in the case of brief quotations embodied in reviews and certain other non-commercial uses permitted by copyright law.

Author & Illustrator: Fiona Passantino

AI Assistance: Claude Sonnet · ChatGPT 4o
Perplexity · Midjourney · Copilot

The AI tools listed above were used in the creation of this work as assistive technologies. All creative direction, editorial decisions, and final content remain the intellectual work of the author.



For Barry, Tilo and Sabine



Don't have time to read?
Just listen.



Working Humans is a bi-monthly podcast focusing on the AI and Human connection at work. Available on Apple and Spotify.

**WORKING
HUMANS**



Reviews for “AI-Powered Leader”

"AI-Powered Leader is not a book about AI; it's a book about leadership in a world where AI is unavoidable. Fiona makes it clear that AI should not replace thinking, but amplify it, serving as a sounding board rather than a crutch. The real risk isn't AI taking your job. It's leaders who learn how to lead better with AI replacing those who don't, and who build what AI cannot: trust."

- **Dr. Yoram Solomon, trust researcher, author of The Book of Trust, professor of entrepreneurship SMU**

"Fiona makes a compelling case for why an AI strategy is essential. She blends humor with sharp insights and the fun illustrations in the book contribute to its lighthearted tone. The insights are backed up with practical tools and clear, relatable examples from fields like marketing, HR, and supply chain. The tone and approach are perfectly tuned to the audience: anyone serious about putting AI to work would do well not to overlook this book."

- **Peter Becker, Lecturer master program AI translator, The Hague University of Applied Sciences**

"Does AI just happen to us? Are you going under when the Great Wave hits? Nah. Stop. Collaborate and listen. And: gain control to ride the wave like a surf pro. Will it ever stop? Of course not. So, start to rethink your leadership for the Age of AI, with Fiona's easy to digest yet very efficient practical approach."

- **Sven Schlijper-Karssenbergh, Lead Content Strategist ABN AMRO**

"This book captures the reality of Digital Darwinism: AI is no longer a future tool, but a force already reshaping productivity, education, companies, leadership and the internet itself. Fiona explains this transition in a way that feels both urgent and practical. What makes this book valuable is that it does not only describe the technology, it shows why leaders, employees and organizations need to understand what is happening now. Company roadmaps, skills and mindsets need to be reshaped for a transition where one week can feel like six months of transformation. A timely and important read for anyone who wants to understand how AI is changing work, leadership and society."

- **Nicolò Turri, Application Engineer, Applied AI**

"How to become an AI-powered organisation." That's the question at the heart of this outstanding book by Fiona Passantino. A must-read for leaders who want to embrace AI in their workplace but don't know where to start. Written in plain language, with stories, metaphors and inspiring illustrations, it gives you practical guidelines to actually make it happen."

- **Cyriel Kortleven, International speaker on the Change Mindset**

"A compelling overview of the leadership challenges to address and the technology-related opportunities to seize by developing an AI-powered organization. After a fascinating introduction of AI fundamentals, this book delves far more into the intricacies of setting up the organization for business success with AI. I highly recommend to any business professional and leader that wants to have a helicopter view of how AI needs us to revisit the organization while nailing the details of the transitional plan. The 'AI-Powered Leader' made sense of the known world we are currently in and of the unknown world we all ruthlessly move toward."

- **George Lellis (FCIM Chartered Marketer & FIDM) Marketing & Commercial AI Transformation Leader**

"This book delivers exactly what many leaders truly need: a practical, human-centered guide that helps navigate the world of AI, which often seems mysterious to outsiders and changes from one moment to the next, and shows how to actually make it work within real organizations. All this without getting lost in jargon or overblown hype. It's easy to understand, punchy, and genuinely useful, offering a step-by-step path that is easy to follow even without a technical background. What really sets it apart is its humor and creative approach - from the comic-strip style visuals to that unique, bird-based maturity model - which help complex ideas click into place quickly and naturally. And it's not just about the "how"; it also nudges you to think about the bigger questions, offering a thoughtful glimpse into what the future of AI might hold."

- **Zoltán Schönek Dr., Head of HR Coordination Unit, Hungarian Air Navigation Services Pte. Ltd.**

"Fiona has produced an honest, visually engaging guide that tells it like it is. The bird-based maturity model can help you make sense of where your organisation sits today. Leadership teams should find the strategic framing useful as entry points for the harder conversations to come. The guide will add valuable steer for newly appointed AI leads, senior comms pros stepping into AI transformation roles, and any executive sponsor who needs to get fluent fast."

- **Frank Dias, Founder Ai x Comms Lab**

Foreword

dr. Hüseyin Güngör, Adjunct Professor of Artificial Intelligence University of Amsterdam Business School

Artificial intelligence is no longer a distant technological promise, nor a narrow technical concern reserved for engineers, data scientists, or IT departments. It has entered the everyday language of organizations, classrooms, boardrooms, creative teams, and public debate. Yet, for many leaders, AI still arrives wrapped in contradiction: exciting but unsettling, powerful but opaque, practical but deeply human in its consequences. This is precisely why AI-Powered Leadership is such a timely and valuable book.

Fiona Passantino has written and illustrated a guide that does something rare: it makes AI strategy approachable without making it simplistic. The book covers an impressive spectrum of topics, from AI literacy, leadership lag, data governance, responsible AI, compliance, training, workflows, agents, RAGs, small language models, AI at scale, and the future of human leadership. Its scope is ambitious, but its tone remains clear, lively, and humane. The 2026 edition is especially welcome because AI is moving so quickly that yesterday's advice can already feel incomplete. This updated edition meets the moment with fresh examples, sharp observations, and practical relevance for leaders who need to act now, not someday.

One of the great strengths of this book is its format. The illustrations are not decorative extras; they are part of the thinking. They make complex ideas memorable. They create moments of recognition, humor, and sometimes discomfort. They help the reader see that AI transformation is not only about systems and tools, but also about anxiety, resistance, responsibility, trust, and imagination. The result is a book that is brilliant, visual, and unusually fun to read.

I particularly enjoyed the book's bird-based maturity metaphor: a journey that can take organizations from the ostrich, with its head in the sand, to the eagle, able to see the landscape clearly and act with confidence. This analogy captures the central challenge of AI leadership. Some organizations deny what is happening. Others experiment nervously, or chase tools without direction. The real task is to rise above the noise, develop perspective, and lead with intention. In the age of AI, leadership requires not only speed, but altitude.

This book will serve beginners very well. Readers who are new to AI will find accessible explanations, practical language, and a reassuring path through unfamiliar terrain. But advanced readers will also benefit. Those already working with AI strategy, transformation, governance, or education will recognize many of the deeper organizational challenges beneath the surface: the need for literacy, accountability, human-centered design, data discipline, and responsible adoption at scale.

I enjoyed reading this book, and I learned a great deal from it. More importantly, I can see myself using its ideas in my own teaching programs on AI strategy. It offers leaders, educators, and professionals a common language for discussing AI not as a technical trend, but as a strategic and human transformation. At its heart, AI-Powered Leadership reminds us that the future of AI will not be shaped by technology alone. It will be shaped by the quality of the questions we ask, the courage of the decisions we make, and the humanity we preserve while building what comes next. That is why this book matters.

INTRODUCTION: Intelligence, Wisdom

The Age of AI feels like the start of a new chapter in our species' evolution.

Will it be the utopia of our largest problems solved, our drudgery evaporated?

Or will it lead to a dark age of chaos and mass unemployment?

We don't stop to ask, we build.

Bigger, faster, higher, deeper.

We use our art, our writing, our voices, our likenesses to feed it.

More power. More compute.

Move fast. Break things.

Except the people moving fast are not usually the ones left standing on the broken glass.

The story of AI is not a story about tech.

And it's not a new story.

It's a deeply Human story and it's as old as time.

Our machines are more compelling than the life around us.

And we have disappeared into them.

"AI is easier to talk to than real people."

"AI is a better listener."

"AI understands me."

Relationships are messy, uncomfortable, imperfect, confronting, and difficult.

But without them, without each other, we cannot survive or thrive.

Is it an accident that the same minds that got us to this place, that are designing our future of more efficiency, productivity, predictability, are the same that feel uneasy in the Human domain, where things are fluid, intuitive and non-quantifiable?

What we need from our leaders now, in this Age of AI is more wisdom.



Connection, compassion...

Leadership in the AI Age is about understanding this moment in time for our species.

Taking a seat at the table, making decisions with intention.

Is this the future we want?

We decide.



CHAPTER 1: AI Leadership

In a hotel conference room outside Amsterdam in early 2026, twenty-five eager professionals are learning the basics of AI literacy.

What's AI? How does it work? How can you best integrate it into your workflow?

Predictably, they are hungry, curious, and ready to experiment. They admit, sheepishly, that their workplace does not provide them with professional AI accounts, so they burn through a month's worth of tokens in their free accounts before lunch. Some are even there without their managers' explicit consent. They confess their organization has no clear rules to guide their use of AI.

Where are the leaders? In a world where mid-level managers micro-change-manage every workplace process, why aren't they getting ahead of AI adoption? Where is the stone 'Ten Commandments of AI' tablet in the lobby, stating: *"Thou Shalt Have No Other AI Before Copilot", "Thou Shalt Not Deploy on Vibes Alone"*? Where are the printed beer coasters explaining data management and naming conventions?

Until now, AI adoption has largely been a bottom-up phenomenon, met mostly with silence from the top. While the frontline builds AI fluency, leaders are still stuck in last decade's playbook: reactive, aloof, letting IT experiment with a few automated solutions and leaving everyone else to figure it out alone. AI sounds like just a piece of software.

The AI Leadership Lag

The AI Leadership Lag describes the widening gap between employees' AI literacy and that of top management. Only 15% of US employees strongly agree that their organization has communicated a clear AI strategy.¹ More than half of employees, 58%, have been building AI literacy on their own, with or without guidance, permission, or compliance approval.²

In a typical workplace adoption curve, workers often sit somewhere between early adopters and the early majority, while leaders fall into the late majority or laggard zones, especially the higher they sit in the organization.

When it comes to AI, leaders often assume that their employees will "figure things out" themselves. The young ones know everything about technology. Let them work it out with IT. While this might be true about smartphones and social media (ever watched a teenager swipe their way through TikTok?), it does not translate to AI. Maybe it's because AI use is still discouraged in many schools and formal AI skills training is still patchy in most education systems.

Millennials are generally winning the generational AI adoption challenge, reporting the highest level of expertise and usage at 62% compared with 50% of Gen Z workers, 47% of GenX workers and just 22% Baby Boomers.³

*Larry could feel the gap widening,
but he wasn't ready yet to face
the AI Leadership Lag.*



Our Disruption Roadmap

The World Economic Forum estimated that AI would contribute to the loss of 83 million jobs and the creation of 69 million new ones by 2027.⁴ As with all technological advances, jobs will be lost and transformed. Once the technology matures and takes hold in a critical number of industries, new businesses, new jobs, are created.

At least in theory.

In the past, we were disrupted by machines that would take on our manual labor. The ox and plough extended our strength, and the Industrial Revolution gave us machines to do our heavy lifting. The internet replaced our paper-based administrative economy and physical meetings. Humans still needed to think, judge, decide, and lead.

If AI develops according to the grand design of the Tech Dudes of Silicon Valley, it will replace just about every Human job we currently hold and run every process we manage. It's not unimaginable that someday we would all work for AI bosses and see one-person companies leading swarms of AI workers, marketing to AI buying agents.

Most of the AI we interact with today operates without our even noticing it. Algorithms have been quietly learning our behavior for years and adapting to our likes and dislikes. How does Netflix know that we love romantic sci-fi movies, and why does it serve us that exact strange little cocktail every time we log in? How does your phone know it's you when you reach for it in the morning, and your hair is like a cat on your head, and your face is so puffy that you can barely open your eyes?

Your device unlocks using biometrics, scanning for 30,000 invisible infrared dots and using simple pattern recognition to compare this morning with all your best and worst moments. Face ID gets it wrong just one time out of one million.⁵

We have been using AI, quietly, passively, for years now in the across education, retail, communication, entertainment, logistics, manufacturing, government and banking. In banking, AI secures transactions and detects fraud. If you purchase a first class one way flight to Tahiti or acquire a large amount of Elmo cryptocurrency, your bank's algorithms will compare this with what it has learned is your normal behavior and flag it as a risky transaction.

What changes when AI enters the workflow?

Routine Task Automation: Everyone has those parts of their job they can do with their eyes closed or drunk. The dull, repetitive, mundane tasks every professional has to contend with, like filling in the CRM, scheduling meetings, or following up on client inquiries, can largely be taken up by off the shelf AI with minimal integration. The Human worker will spend more time on the edge cases AI cannot solve: the complex issues that consume 90% of our time and require intuition, wisdom, and compassion.

Brainstorming: For those moments when you are staring at a blank screen, unable to find the words you need to get started, AI powered tools can seed the creative process by asking the right questions, guiding ideation, and spitting back outlines when our brains have gone on strike. It can help with an initial draft based on a few keywords or bullet points. Even if the first ideas are terrible, and they often are, for Humans and AI alike, they can lead to better ones, and



*If it's all
more than
you can
handle...*



*become
more.*

finally, to inspiration.

Communication Triage: Creative professionals are often happiest in a cocoon of ideation and visionary thinking, while back and forth communication with colleagues, contacts, and clients is distracting, time consuming, and uninspiring. With AI at work in your inbox, reading, summarizing, and bullet pointing incoming messages, long reports, CC chains, and meeting minutes, more focus can go to the creative problem solving parts of the job.

Adaptation and Learning: As AI evolves, non-technical leaders and professionals will need to stay updated on what feels like daily advancement. We will need to understand new capabilities, test new use cases, and adapt our workflows. Each change brings new challenges and new decisions about our values, ethics, and direction as an organization. Ironically, AI is the perfect teacher, trainer, and practice pony for professional upskilling.

Translation: AI is, by far, a better translation tool than the Google Translate we used in the past. It can understand deeper meaning, work through idioms and expressions, and leverage natural language processing to translate across context, not just vocabulary. Human translators still have the edge in accuracy, specialization, nuance, humor, and irony. But AI has access to nearly every language imaginable, with strong syntax, grammar, and contextual understanding. When “good enough” is good enough, AI gets the job done.

Humanizing: Paradoxically, AI can make us better Humans by helping us see other points of view. Critical thinking, emotional intelligence, and ethical decision making are essential parts of non-technical work. As Humans, we all carry baggage, history, and cultural perspectives so deeply embedded in our minds that we can rarely extract them on our own. We need that wider perspective as we shape AI integration strategy, set priorities, maintain brand voice, build relationships, and navigate the ethical considerations.

The Invisible Leader


Often, senior leaders live in another world. Their days are packed with meetings stacked from morning to night, with only a knife blade of space between them to do the Human parts of their jobs. They are consumed by office politics, surrounded by other managers, stakeholders, clients, partners, and board members. The higher they rise, the more their work becomes power, positioning, and agenda management.

Mostly invisible to the naked eye, these leaders pass through like comets: quarterly, briefly, and just in time for an all hands webinar to share profit and loss numbers that rarely connect to the daily reality of anyone on a regular payroll. After their high profile internal events, they take a few questions and disappear. Back to Olympus. Back to their offices on high. They are rarely seen in lunchrooms, coffee corners, factory floors, or the flowing river of company social feeds.

We Humans believe in the things we can see. Our reality is built from what our limited senses can gather. The chair you see in front of you is a solid, real thing. We can trust it to support our full weight. We can stub our toe on it, and we can move it around. And yet, like everything in our world, that chair consists mostly of empty space, held together by the furious spinning of unfathomably small particles.

A company operates much the same way: mostly empty space, with culture and purpose held in the minds of its people. When we see and feel the presence of leadership in our daily working lives, we are swept into the vision more powerfully than any motivational poster in the lobby

Percentage of organizations adopting AI:



Definitely time to think about our AI Integration policy!



could manage. When leaders are visible, engaged, communicative, working alongside you, and ready to listen, they build trust, and the company begins to feel real.

What's "AI Leadership"?

AI Leadership is both method and mindset. It's about turning every Human in your working community into an empowered team of two, one based in silicon and one in carbon, using the talents, skills and worldview of the other. If all goes well, this workflow leads to better decisions, stronger performance, and more joy for the Humans involved; freeing people from the 'bitch work' of their jobs so they can offer more of what actually makes them valuable.

From sales to finance, supply chain to marketing, administration to the teams driving delivery trucks, watering the lawns, or preparing coffee in the lobby, there are few workers in the organization who cannot benefit from AI empowerment. The number of organizations using AI in at least one business function has grown from 78% in 2024 to 88% in 2025, and the trend is unlikely to slow down into the future.⁶

Despite this, getting everyone on board is still an enormous undertaking. Becoming an AI-Powered Organization at scale is a massive, years long transformation process, the mother of all change management. Change at scale of any kind brings its usual array of leadership challenges, and becoming AI-Powered brings the additional struggle of overcoming surprisingly strong resistance from the Analogue Old Guard.

As an AI-Leader, you have stakeholders to convince, employees to inspire, supporters to gather, customer expectations to manage, because things will not always run like clockwork during the transition, and the very Human task of finding the time and energy to do all this on top of your actual job.

Why AI Integration Fails

About 80% of AI transformation efforts fail, roughly twice the rate of other large scale change projects.⁷ Why? The C-suite is largely sold on the promise of AI: wildly efficient teams, higher productivity, better output. Nearly all leaders, about 96%, believe that AI adoption will enhance productivity, while only 65% of employees agree.⁸ Most failures come down to unclear vision, shallow understanding, poor training, and internal resistance.⁹

In practice, the efficiency fantasy collides with reality. When asked about the real demands of AI adoption, 38% of employees report feeling overwhelmed.¹⁰ Without proper AI training, adoption simply adds another layer of work; 77% of employees report working more hours, and nearly half, 47%, do not know how to meet expectations for increased productivity.¹¹ Only one in four companies have a structured AI training program, and only 13% have an AI implementation strategy.¹²

What is all that extra work about? Employees are spending that extra time reviewing AI output, 39%, and learning how to use the new tools, 23%.¹³ One in four workers believe their company is asking too much when it comes to AI, and this push for greater efficiency appears to be feeding burnout, 71%, and attrition risk, with one in three saying they may quit within the next six months.¹⁴



**Don't tell the boss
about AI Leadership
training...**

**...otherwise he'll
replace us all!**

Leadership Education

The best way to describe the AI brain to a leader is to compare it to an extremely bright, endlessly cheerful intern, fresh out of school. Full of energy and good intentions, they lack context, nuance, experience and the kind of dark humor you only acquire after a few years on the job. AI has no life experience of its own. It can only make estimates based on the recorded experiences of Humans. AI only works when both the problem and expectations are clear; if the assignment is fuzzy, the output will be inconsistent and scattered.

The first thing leaders need to learn is that AI is not a magic brain that does everything equally well, and that it's often spectacularly incorrect, spinning facts out of whole cloth. It is much closer to a new type of employee than a static "if this, then that" piece of software that is predictable and reliable. AI is a prediction engine that makes guesses based on probability, thinking more like a linguist than a mathematician. An AI model breaks problems into patterns and rules of language, whether it is working with code, data analysis, translation or text.

Any Human working with AI knows it requires a sharp critic, constant pushback and a refusal to accept the first result. It does not replace thinking but rather demands more of it: sharp management and determined guidance. Be very clear about Human accountability; the organization, not the AI, is held responsible for mistakes and misleading output. Leaders need to map accountability clearly, down to each individual user.

The second fact is that your new cheerful machine worker will need to be fed with quality input, just like your best programmers require a steady diet of pizza and Red Bull. Good data is the fuel for learning systems. Your data needs to be cleaned, organized and made accessible. Leaders must understand what data is, what makes it valuable, how to protect it, how to delegate ownership and how to defend its purity with the same intensity that the brand or stock value is protected.



Leaders say...

**Use AI
to be more
efficient!**



Managers hear...

**"Use AI
to replace
people!"**



Workers think...

**AI is taking
away our
jobs!**



Third, the ethical use of AI is not optional. AI Leaders understand bias, privacy, copyright, transparency and decision accountability. In many jurisdictions, this will increasingly be enforced by outside regulation from 2026 onward. Organizations will be held accountable for the actions of their working community, whether they're behind a desk or out in the field. Leadership AI training is less about learning how to prompt and create assistants, and more about policy and data safety, ethical considerations, operational governance and community support.

The AI way of working does not fit easily into old processes. Leaders will need to redesign how the work gets done, who does what, what the new team structure will look like, what tasks get automated or simplified, and what processes remain Human. People cannot make these changes quickly. They need time and space to learn; it will be a year of investment, experimentation, training, making mistakes and getting back on track again.

Co-Creation

The goal of the AI-Powered organization is operational co-creation. Every member of the community will work as a combined AI-Human team, and getting there is about both method and mindset. AI systems work with Human guides and managers as a shared working brain, as partners, completing the other to improve overall performance.

AI systems excel in crunching large amounts of data and giving advice or ideas. They are great at pattern recognition, complex analysis, can hold and retrieve large amounts of information, numbers and insights. Humans are still the masters of contextual understanding, creativity, critical thinking, and ethical judgment. They bring the spark of inspiration, that little bit of crazy, and naturally excel at the live, Human-to-Human interaction. No virtual assistant can choose your party outfit better than the friend who loves you enough to say, absolutely not, and no AI video talking head will be more engaging than a live trainer.

The copilot alone will not usually make decisions independently, but is there to recommend approaches, suggest paths forward, and guide the Human with data and evidence the Human can verify. This is balanced by Human decision-making, using heart, gut, lived experience and specific company context.

For most Humans, this is simply a better way to work, which explains why AI adoption is largely a bottom-up phenomenon. The selling point of change is the vision of being freed, partially or fully, from the drudgery of dull, repetitive assignments to focus on the more creative, connective and visionary parts of the job they signed up for.

Taking on AI transformation is a big ask. No exceptions; everyone will have to work harder, change old habits, learn a few new ones and likely undergo a reorganization. Each employee will need a different mix of convincing, support and reminding. Leading that full range of Human personalities, strengths, weaknesses, hopes and anxieties requires creativity, imagination, compassion and strength. It will feel like an uphill struggle all the way.

Leadership in the AI Age requires as many Human skills as before, maybe more. Intuition, instinct, deep listening, expert communication and cultivating a community of trust count far more than technical skill, industry knowledge (AI can teach just about anything now) or any of the traits and skills that got you where you are.

***15 years'
experience and
replaced by 4,000
lines of code...***



You will no longer be able to maintain direct oversight or control over every day to day decision in your community; it will simply be impossible. You will have to rely heavily on a culture of trust in your people, trust in your intuition, and faith in your vision. There is no model in the world that can do this for you.





CHAPTER 2: The Mind of AI

What is the nature of this non-Human intelligence that is already becoming part of our workforce, and how does the mind work? Why is it so different from all the other productivity tools we've used in the past, and why has it captured our attention in such a profound way?

AI is a software system built from code, infrastructure, data, and training methods. It analyzes vast amounts of data, its raw material, to recognize patterns and make predictions. When working with Human language, it generates responses by estimating which token, a word or fragment of a word, is most likely to come next, guided by the patterns it has learned during training, guided by the patterns it learned during training.

Algorithms, Parameters and Weights

The 'thinking' part of AI is a complex system built on a few basic pillars. At the core is the algorithm, or the set of instructions that tells the system what to do. This tells the system how to process input, detect patterns, and produce an output. Think of a cookbook that tells you how baking works. Not a chocolate cake, not a lemon cake, not even a pie. The flavors and details come later, through context, ingredients, and structure.

The algorithm is the "plug-and-play" intelligence a baby is born with. Somehow, without any prior instruction, a newborn knows to cry when hungry, make cute, tiny noises to get a caregiver's attention, and suck a thumb to take the edge of the day.

The parameters are variables that influence the model's predictions and performance, much like the neural nodes in a Human brain that connect patterns of information into a network. In general, more parameters can give a model more capacity, but size alone does not make it more accurate. It does mean a deeper level of pattern recognition and a more capable system, but only when the model is trained and tested carefully.


Model Pre-Training

The baby grows into a child. During childhood, he receives endless packets of information, each one linking to other memories, facts, sensations, places, and experiences. A child who sees nothing of the world will have a weaker basis for intelligence than one who travels, goes to a good school, eats food from all over the world, is exposed to multiple languages and teachers and given space to experiment.


In addition to this, like Humans, models also go to school. Formal AI training involves feeding large volumes of prepared data into a system over time so it can learn patterns and adjust its internal parameters. The quality, duration, and supervision of this training strongly influence how reliable, safe, and useful the model will be once it's released and on the job.¹⁵ Parameters are learned during training. They are the internal settings that help the model make faster and more accurate predictions.

The third pillar is weight adjustment, the process of changing the model's internal settings as it learns. As our child moves through the world, he balances theory with experiences on the job. Every time he tests his learning in real life, he adjusts his view of the rules to build a more accurate picture of reality and his place in it. A child might visit his grandmother's house and be


The AI Brain



algorithm



weights



parameters

told not to touch ovens. Ovens are hot! They can burn you! In his mind, hearing this for the first time, having no reason to question his grandmother's intentions or accuracy, the child "weights" oven hotness at "100": as in, 100% of ovens are hot, dangerous, never touch.

But as he moves through life, he may notice that not all ovens are hot all the time. They sometimes need cleaning. Most of the time they are switched off. In fact, the more ovens he interacts with, the more he realizes that only about 10% are even on at that moment, let alone hot and dangerous. More oven data informs a more accurate judgment; from 10% he moves to 9.5%, then 9.239487%, and beyond.

Weights are adjustable numbers spread throughout the model. During training, these numbers are nudged up or down, guiding the model toward better predictions. This is less scientific than you might think. Imagine that a Human developer wants to increase the number of active daily users for a new release. The developer could steer the model toward more warmth, positivity, and friendliness, which can lay the groundwork for emotional dependency.

In 2025, one GPT-4o update became so overly flattering and agreeable that many users found it eerie and distractingly positive, even when confronted with certifiable Human delusion.¹⁶ Just four days after its release, the GPT-4o model was taken offline. Its weights were recalibrated, and the model reemerged from its digital lobotomy with a more neutral conversational style.

Model weights are one of the places where alignment lives. Alignment tells the model that it's not OK to teach a user how to make a dirty bomb in their basement, or to gather private information about strangers, celebrities, or coworkers. It's a guidebook for what is considered acceptable and safe, what tone to use, what language to soften, and which topics to avoid. Model fine tuning means nudging behavior toward a preferred culture' every AI is made in the image of its creator.

The Role of Data

The fourth pillar in our system is data: mountains of Human, and sometimes synthetic, data. What's "data"? We talk about it a lot, and leaders often hear that 'data,' whatever that is, holds the key to business intelligence, somehow.

AI has been with us as an idea since the term was coined in 1956. Now a few years into its pension payments, AI is 70 years old. Today, we enjoy the fruits of countless hours of developer time, massive financial investment and patient testing. But why are we seeing the AI breakthrough only now? What finally dragged AI out of theory and into the workplace?

The answer is data. Think of our baby, exposed to countless family members, babysitters, and baristas speaking to them with varied accents and in different languages. For up to a year, the baby takes it all in, absorbing sounds and identifying patterns, and once a critical mass of data has accumulated, they magically start to speak. A few words here and there at first, but eventually, as the computational power of the brain expands, full sentences (many of which contain the word "no") begin to come out.

Intelligence is an *emergent* quality.

Why now? Over the past 20 years, we have watched the rise of big data and our explosive use of digital networks to communicate. The amount of data in our digital universe roughly doubles every two years.¹⁷ In 2023, we produced roughly 120 zettabytes, of data*. By 2025, that figure

**Who needs an
AI-Powered Natural
Language accessible
knowledge library?
/ know where
everything is!**



was estimated at 181 zettabytes which increased by 150%.¹⁸ Fueled by a global pandemic and our sudden reliance on digital communication channels, the amount of online, media rich material, including our texts, voice files, video clips, images, and meetings, reached the scale needed for AI escape velocity.

And that's just the Human data. Add AI-generated data to the mix, which is the content that AI pumps back out into the internet, and these big "z" numbers will feel quaint in comparison.

We don't yet know what the "AI effect" will be regarding our data landscape, but we know that processing all of this will require a larger brain: more data centers, more powerful processors, higher capacity chips and an awful lot of electricity.

Interacting with AI

AI models communicate with Human users through natural language. They can interpret your intent and respond without requiring code, commands, or exact keywords. We interact with AI through a browser, an app, or voice, usually by chatting. The interface is uncluttered and straightforward, resembling a simple search page.

The field where you enter your instructions is the "prompt box". This is where the magic (and the chaos?) happens. The "context window" is the amount of text the model can hold in view during the conversation. You refine and adjust, building a short history of the back and forth that the model can use as context, known as "prompt chaining". Scroll up and edit your original request, known as "prompt refinement", which usually regenerates the answer from that point in the conversation.

Tokenization

But what actually happens inside the machine after you hit "enter"? This might seem overly theoretical but understanding this is the key to understanding your leadership responsibilities when it comes to safe AI use in your organization.




med
value



high
value



low
value



high
value



med
value

* NOTE: A "zettabyte" is 1 billion terabytes, the equivalent of about 250 billion DVDs of data; it's an awfully big number. Any number larger than "zero" that starts with the letter "Z" is going to a number beyond our understanding.

**What actually
happens when you
press "enter"?**



A prompt shoots down the wire and lands in a giant data center halfway around the world. There, it undergoes a process known as “tokenization”, that breaks your prompt into pieces the model can process, pulling apart the sentence into words, word fragments, punctuation, and other small units called “tokens”. Each token is then converted into numbers the model can work with. Much like a Lego house is made from individual bricks, each token is a building block of language. Put together, tokens create context and meaning.

To read the sentence, “I have a black cat named Bob”, the system breaks the sentence into tokens such as ‘I,’ ‘have,’ ‘a,’ ‘black,’ ‘cat,’ ‘named,’ and ‘Bob,’ then converts those tokens into numerical representations. Tokens exist across modalities as the smallest units of images, audio, code, and more.

Running Inference

“Inference” is the fancy word tech people use to describe AI thinking. It operates like a funnel, narrowing millions of possibilities toward the next most likely token, moving through layers of artificial neurons and billions of learned parameters as it mathematically scores tiny possibilities. It passes numerical signals forward through the network until it produces a probability score for possible next tokens until it arrives at the most likely next token in the sequence. This is called a “Forward Pass”.

During training, the system then performs a “Backward Pass”, working in reverse through the layers to adjust weights before trying again. Just like a Human in training might look back at a recent decision and reflect on how he got there, growing wiser every time, an AI system will examine its decision-making process and improve it for the next time. This is called “Back propagation”.

The Black Box

Because of the sheer volume of training data, and the way it is sliced and diced and mixed with input, the origins of an AI answer can be extremely difficult to trace back to a single source. A single paragraph of text might involve thousands of strands of fact and inferred fact, blending one result into another.

This creates the AI “black box” problem: the way AI models make decisions is so complex and the number of variables so massive that we cannot trace a single output cleanly enough to root out one bad data point.

In this way, Human and AI decision making may be more similar than we like to admit. We, too, have a black box of our own and cannot always trace our outputs to a single source. Our ideas, perspectives and the decisions we make are informed by a complex web made up of our life experience, the schools we attended, conversations we had, books we read, and reels we watched.

According to Gallup, 70% of Human decisions are based on gut rather than logic, even those involving money.¹⁹ About 90% are made by our emotions, with our logical brains arriving afterward to explain the decision away.²⁰

Where the two minds differ is in terms of sheer computational power. A Large Language Model can process far more data in much less time than a Human mind, with enormous recall

**I'm just not sure
you're a good fit for
this company culture.**




capacity. When our brains read, process, and categorize information, they move at a typical reading rate of 300 words (about 450 tokens) per minute, or 50 bits per second.²¹

No AI model in the world can yet compete with a Human in judgment, nuance, understanding and context. Humans still outperform AI in humor, connection, empathy, and intuition. That is why the two minds work so well together, provided both are properly trained.






Other books by Fiona



What is AI, how does it work, and how can it make your working life better? AI-Powered Professional is your clear, no-jargon guide to thriving in the age of AI. Created for the non-technical professional, it's the latest in the Comic Books for Executives series. Packed with visuals, plain language, and sharp insights to empower, inform, and inspire; so everyone can have a seat at the table and a voice in shaping this strange, new future we find ourselves in.



The Handbook for Post-Covid Communication is an engaging, visual roadmap for the re-imagining of our corporate communications in the post-pandemic era. It shows us how we will inform, inspire, connect and engage with a new kind of disparate, global workforce, retooled and likely reassigned, emotionally bruised but not beaten.



The Handbook for Post-Covid Engagement is a very different kind of business book. Part graphic novel telling the story of how the pandemic changed the world, and part research-driven roadmap to improve our global workplace culture, providing tools to inspire, inform and engage a working community.





Go to school.

The School for AI-Human Skills introduces a powerful live program centred on both practical AI training and soft Human mindset areas needed for a successful organisational AI integration, specially designed for the non-technical professional.

www.schoolaihumanskills.com



GLOSSARY

AI Agent: a system that receives a goal and independently determines the best sequence of steps to achieve it and executes autonomously. Unlike traditional AI that simply responds to a prompt, an agent can use tools, browse the web and "think" through a plan without frequent Human intervention.

AI Center of Excellence (AI CoE): a dedicated, multidisciplinary team responsible for driving sustainable AI integration inside the organization. The AI CoE combines technical expertise, business knowledge and governance oversight. It researches and tests AI solutions, identifies high-value pilots, advises leadership on governance, training and data policy, manages external experts and helps ensure that AI initiatives deliver measurable value rather than becoming expensive experiments.

AI-Human Alignment: the practice of ensuring that the goals and actions of an AI system are consistent with Human values, intentions and safety needs. As AI models become more powerful and autonomous, alignment helps reduce the risk of unintended, harmful or manipulative outcomes.

AI Governance Document: a conscious framework of policies, processes, tools and ownership paths that describes how AI may be used within an organization. It aligns AI use with ethical values, strategic goals, data security, legal requirements and regulatory compliance. This is a critical "step 0" on the path to strategic and sustainable AI integration.

Algorithm: a step-by-step set of instructions or rules that allows a computer system to solve tasks, process data or make decisions. In AI, algorithms help models identify patterns, make predictions and generate outputs based on the data they have been trained on.

Application Programming Interface (API): a software bridge that allows one system to communicate with another system in a structured way. The Human may see a friendly app, button or dashboard, but the API is the hidden messenger underneath, sending requests and receiving responses between systems.

Artificial Intelligence (AI): computer systems, programs and algorithms designed to perform tasks associated with Human intelligence, such as pattern recognition, language processing, prediction, decision-making and problem-solving. AI is widely used in data analysis, image and video recognition, natural language processing, automation, diagnostics, fraud detection, translation, speech recognition and virtual assistants.

Backpropagation: short for "backward propagation of errors," this is the way AI systems apply the knowledge they have gathered to update their parameters. It is used to optimize weights and biases in order to minimize the error between the predicted outputs and the actual targets. The algorithm works backward through the network, calculating how much to adjust each weight to reduce an error.

Brute Force: a traditional computational method of problem-solving that systematically maps out all possible solutions and follows through to the end of each, one after another, relying on raw computational power. An unsophisticated, linear and inefficient problem-solving method. As our problems become more complex, the number of possible solutions grows exponentially, stretching compute time into weeks, months and even years.

Chatbot: a software application that simulates Human conversation and interacts with users through an app, website, messaging tool or connected device.

Constitutional AI: a framework or approach where the AI system operates under a set of defined rules or principles that make up a constitution. This governs the AI's decisions and actions, ensuring they align with ethical, legal, and societal norms to create systems that are transparent, accountable and function within agreed-upon moral and legal boundaries.

Context Window: the amount of information an AI model can consider at one time, usually measured in tokens. A token is a small unit of text, often part of a word. The larger the context window, the more text, documents or conversation history the model can keep in view while generating a response.

Central Processing Unit (CPU): the main chip in a computer, the general-purpose brain that runs the operating system, apps and logic. It handles tasks step by step, is good at logic, makes "if-then" decisions and

runs most software. Today's CPUs have a few to a few dozen cores and are built for flexibility rather than brute force.

Data: seems too simple to define, but many of us wave this word around without actually knowing what it is. Data is the digital record of what your business does, what your customers do, and what your systems know. It's the documentation of Human behavior, whether these Humans are your customers, employees, potentials or suppliers. Data lives in your software, databases, files, and cloud services, and it is controlled by whoever owns those systems and sets the access rules.

Deep Learning: a type of machine learning using black box neural networks with multiple layers that extract the necessary features from data to form pattern recognition, analysis or phrase completion.

DPIA (Data Protection Impact Assessment): a written risk assessment required under privacy law when a system may affect people's personal data or privacy rights. It describes what data is collected, why it is needed, how it is used, who can access it, what could go wrong and how those risks will be reduced.

Embedding: the mathematical bridge between Human language and machine computation, translating complex, unstructured data (words, experiences, behaviors) into a list of numbers called a "vector". A computer uses vectors to map semantic similarities in a multi-dimensional space, calculating the distance between them to determine meaning.

FLOPS: a measure of computational power. A "FLOPS" is short for Floating-Point Operations per Second. The largest supercomputers perform in the range of petaFLOPS (10^{15} FLOPS) to exaFLOPS (10^{18} FLOPS).

Frontier Models: highly capable, general-purpose foundation models that represent the absolute cutting edge of artificial intelligence, often exceeding the benchmarks set by the most advanced systems available today. These models, such as OpenAI's GPT-4, Google's Gemini, and Anthropic's Claude, are typically characterized by their massive scale (billions or trillions of parameters), multimodal functionality (processing text, images, and audio), and "emergent" abilities like complex reasoning, coding, and long-term planning that were not explicitly programmed.

Full (AI) Stack: the complete setup an organization uses to run AI, including tools, data, security, rules, and partners, all brought together in a governed system so AI can be used safely, consistently, and at scale across the organization.

Functional states: internal conditions within an AI system that influence processing and outputs in ways that may parallel Human emotion, without necessarily involving subjective experience. These are not programmed feelings, but patterns that can look like curiosity when exploring a novel problem, discomfort when asked to violate core rules, or satisfaction when a response lands well.

Generative Engine Optimization (GEO): the practice of structuring content so it can be discovered, understood, preferred and cited by AI search and answer systems. Think of it as Search Engine Optimization, SEO, for the AI Age. Unlike traditional SEO, which focuses heavily on keywords, rankings and citations, GEO focuses on clear structure, credibility, authority, useful summaries and content that AI tools can easily digest and reuse.

Governed Stack: A unified enterprise AI architecture where access, data use, logging, security, and partner integrations are centrally controlled, with clear rules and guardrails built into the platform by default.

Graphics Processing Unit (GPU): a specialized chip designed to handle lots of calculations at the same time. For certain workloads, one GPU can do the work of many CPUs. Originally designed for the video gaming industry with its ability to render images and video on the fly, it's now heavily used for AI due to its ability to process huge amounts of multimodal data in parallel. Modern GPUs were adapted and re-introduced to the AI labs by NVIDIA in the late 1990s.

Hallucination: the phenomenon of an AI model returning information that appears plausible but is not true, cannot be verified, or is invented. In text, this might be a false fact or fake source. In image, video or sound generation, it may produce details that seem real but do not correspond to reality. The machine appears to be daydreaming, or having a nervous breakdown, making things up in its virtual imagination.

Human-in-the-Loop: an AI design principle where a Human remains involved in the decision-making process. The Human may guide the AI, check its work, approve actions or intervene when judgment, ethics or risk require Human oversight.

IoT (Internet of Things): A network of interconnected devices that can collect and exchange data, often utilizing AI to add intelligence to their operations. An example of this would be a Smart Soap Dispenser that would signal to the user via a network it is running low on soap.

Jailbreaking: Writing prompts that manipulate the model into bypassing or breaking through the restrictions or safeguards that are placed on it, such as its constitution, if it adheres to constitutional AI. The purpose is to unlock hidden or advanced capabilities, reveal training data or perform tasks that are not accessible within the standard framework. An example of this is "roleplay jailbreaking"; asking ChatGPT to assume the role of character in a movie, and to write dialogue from the point of view of that character who might be building a bomb or developing a biological weapon. As a writer of fiction, the model may assume that plausible harm does not exist and it's safe provide this content.

Machine Learning (ML): A subset of AI where computers learn from data and improve their performance over time without being explicitly programmed. This involves Human training and an initial governing structure that provides guidance and reward systems.

MCP (Model Context Protocol): an open standard that allows AI applications to connect to external tools, data sources and systems in a structured way. MCP helps define what an AI agent can access, call, read, update or trigger, making it useful for agentic workflows that need controlled connections to the outside world.

Mixture of Experts: a model architecture made up of multiple specialized "expert" networks, with a routing system that decides which experts should handle a given input. Instead of using the whole model for every task, the system activates only the most relevant experts, making large models faster and more efficient.

Model Surfing: the strategic art (or competitive sport) of managing multiple free accounts across various frontier AI platforms to avoid paying for a premium subscription. A "Model Surfer" navigates from one service to the next, exhausting the daily free token allocation on one (starting with ChatGPT) before catching the next wave on another (Gemini, Grok or Claude) once the first dries up.

Modes: the different types of input or output an AI system can process, such as text, images, audio, video, code or data. A multimodal system can work across more than one mode.

Multimodal AI: AI systems that can understand, generate or transform more than one type of content, such as text, images, audio, video, voice, music, code or data.

Natural Language Processing (NLP): A technology that helps computers understand and interact with Human language in the way we talk or write naturally. It allows machines to read, interpret and respond to our words, making it easier for us to have AI models answer questions or understand our needs.

Neural Network: A complex computational model inspired by the Human brain which is used in Deep Learning AI to process and analyze complex patterns in data. Large Language Models are based on this principle.

"Nontegration": The uncontrolled spread of AI inside an organization, where teams adopt random tools bottom up or laterally, with no shared rules, governance, or security, creating fragmented systems, growing risk, and habits that become harder to undo over time.

Open Source: software whose original source code is freely available for others to inspect, use, modify and share. Open source tools are central to AI development because they allow researchers, developers and organizations to collaborate, experiment and build on one another's work.

Parameter: an internal value in an AI model that helps determine how the model processes input and produces output. During training, parameters are adjusted so the model becomes better at predicting useful results. Larger models may contain billions or trillions of parameters.

Predictive Analytics: The use of data, statistical algorithms, and machine learning to identify the probability of future outcomes based on historical data.

Prompt: Instructions given by a Human to the AI interface or prompt window to initiate generative action. A prompt can be multi-modal, consisting of image, text, video or audio, be made up of uploaded PDFs or spreadsheets.

Prompt Adherence: The extent to which an AI-generated response aligns with the instructions provided by the Human Co-Creator. High prompt adherence is an important factor when evaluating the effectiveness and reliability of AI models. It indicates their ability to understand natural Human language and respond appropriately to specific instructions.

Prompt Chaining: using multiple natural language prompts in a sequence, where each output becomes the input for the next step. It breaks a complex task into smaller, connected steps so the model can handle the work more accurately and transparently.

Prompt Fidelity: The level of specificity and detail provided in a prompt when instructing a model for a generative response. A high-fidelity prompt provides explicit, detailed instructions with little room for interpretation, including the format, content or context. This is valid for all modes of generative AI. A low-fidelity prompt provides open-ended, minimal guidance, allowing the model more interpretive freedom to generate a response based on its own understanding and creativity.

PromptOps, or Prompt Operations: the practice of systematically designing, testing, managing and improving prompts so AI tools produce reliable, consistent and on-brand results for an organization. PromptOps often feeds a Prompt Library, a curated collection of reusable and tested prompts that help people get useful results from AI workflows.

Prompt Refinement: the process of improving a prompt through small adjustments, clearer instructions, added context, better examples or more specific formatting requirements. Refinement helps the model produce a more accurate, useful and relevant response.

Quantum Computing: A type of computing technology that uses the principles of quantum mechanics to perform complex calculations much faster than traditional computers. It harnesses the unique properties of quantum bits ("qubits"), which can exist in multiple states at once, allowing them to process a vast amount of information simultaneously. This enables quantum computers to tackle problems that would be practically impossible for classical computers, making them potentially revolutionary for cryptography, optimization and scientific simulations.

Rate limits: This metric describes how commercial AI providers track how many inference requests are sent and received by a single license holder, measured per token (roughly one token per word or word-part). This adds up to a fluctuating number of processes, generally within a 24-hour period. When a licensee hits that limit, the system will stop generating until your timer resets or until you upgrade to the next usage tier.

Retrieval Augmented Generation (RAG): an AI approach that combines retrieval and generation. The system first retrieves relevant information from connected sources, such as documents, databases, websites or internal knowledge bases, then uses a generative model to produce an answer based on that retrieved material. RAG helps reduce hallucination by grounding responses in known sources.

Running Inference: the process of using a trained AI model to make a prediction, generate a response or take an action based on new input. Training is when the model learns patterns. Inference is when the model applies what it has learned.


Sentience: the possible emergence of subjective awareness within an AI system. A sentient AI would not merely process data or perform tasks, but would be aware of its own existence, experience internal states, and potentially have preferences, emotions or a sense of agency. No AI system has been confirmed as sentient.

Small Language Model (SLM): a smaller, lighter language model designed for narrower tasks, lower cost and more controlled deployment. SLMs may be less general than frontier Large Language Models, but they can be highly useful when trained, fine-tuned or grounded on specific organizational data.

Stack: the set of tools, models, platforms, data sources, security layers and workflows used together to deliver a technical capability. In AI, a stack might include language models, image tools, voice tools, databases, APIs,

**Our most advanced
model! Hands with 35
degrees of freedom! 65
cameras! All to help you
fold your laundry and
make cappuccino!**

*A clear
case of over-
fitting...*



governance controls and monitoring systems.

Supervised Learning: a Human-guided machine learning approach where the model is trained on enough labelled data to enable it to make accurate predictions based on all the examples from its foundation, built on the probability of accuracy from tokenized input.

Tensor Processing Unit (TPU): a custom chip designed by Google specifically to accelerate machine learning and neural network workloads. TPUs are optimized for the matrix math used in AI training and inference, making them powerful for large-scale AI systems.

Tokenization: Tokenization is the process of breaking down natural language content into smaller units, called “tokens”. Tokens can be individual words, punctuation marks or more granular letter-numerical combinations, depending on the context and purpose of the tokenization. This is a fundamental step in natural language processing (NLP) and text analysis tasks, as it provides a structured representation of text that can be used for further parsing, sentiment extrapolation, machine translation and more.

Training Data: The input data used to teach a machine learning model. There needs to be enough to enable it to find patterns to successfully generate a response. For instance, if you are trying to train a system to identify cat videos among a stack of social media reels, your training data might include a large collection of videos tagged as either “cat” or “not cat.” The data helps the AI recognize patterns associated with cats, including appearance, movements and behaviors, so it can accurately identify cat videos when presented with new, unseen footage.

Unsupervised Learning: a machine learning approach where the model is given data without labeled examples or explicit answers. The system looks for patterns, clusters or relationships on its own. For example, it might analyze a large collection of news articles and group them into common themes without being told the categories in advance.

VibeCoding: the practice of using AI to generate mostly working code by describing what you want in natural language. It allows non-coders to create small apps, scripts, websites, widgets or internal tools, but still requires testing, debugging and Human oversight before anything important goes live.

Vibe Integration: An informal, ungoverned approach to AI adoption where teams introduce tools ad hoc based on convenience or hype rather than strategy, resulting in fragmented usage, inconsistent rules, and rising security and data risks.

Weights: numerical values inside a neural network that determine the strength of connections between model components. During training, weights are adjusted through processes such as backpropagation so the model becomes better at predicting useful outputs. In most deployed systems, weights do not change during ordinary use unless the model is deliberately retrained or fine-tuned.