

Between Impulse and Understanding

Erik Palsbo

Between Impulse and Understanding

Consciousness, Experience and What It
Means to Be Human in an Age of AI

About the Author

For several decades, Erik Palsbo has worked with leadership, international organisational projects and complex technological systems. His many years of experience with cultural differences, decision-making processes and practical dilemmas give him a distinctive ability to connect abstract philosophical ideas with concrete reality. Today, he writes about philosophy, society and modernity.

© 2026 Erik Palsbo, erik@palsbo.com

Publisher: Bookmondo

Cover design: 'The Impulse in Consciousness'

Edited: 30. September 2026

ISBN: 9788740990324

Contents

Preface.....	1
The Book's Two Movements	1
Book One Demystifying the Mind: Consciousness, Machines and the Mechanics of Ethics	5
Chapter 1 Introduction: Removing the Magic from Consciousness and Ethics.....	7
Chapter 2 Consciousness as Impulse Integration: Granularity, Temporal Binding and Pattern Formation	9
Chapter 3 The Hardware Architecture of the Mind: CPU, Data Buses and Interrupt Control.....	13
Chapter 4 The Real-Time Weighting Engine: Why Morality and Ethics Are the Same Algorithm.....	19
Chapter 5 AI and the Human Being: Functional Differences, 'Underlying Ethics' and the Possibility of Translation	23
Chapter 6 Monitoring, Weight Arbitration and Governance: What Can Be Measured, and Who Decides?	27
Chapter 7 Demystification as a Platform for Understanding and Criticism: Why It Is Fruitful	29

Chapter 8 Conclusion: A Provisional Platform and Ways Forward	31
Book Two From Impulse to Understanding	33
Chapter 1 Introduction	35
Chapter 2 Intelligence and Consciousness	41
Chapter 3 AI and the Question of What It Means to Be Human	45
Chapter 4 The Language Model and the Impulse Model	49
Chapter 5 Noise, Filters and Error	55
The Field of Experience	61
Chapter 6 The Department of Misfits.....	63
Chapter 7 Working Without Knowing How	67
Balance and Tension	71
Chapter 8 PAEI — and Why It Suddenly Made Sense	73
Chapter 9 When Disagreement Regulates.....	77
Chapter 10 What the Plan Had No Room For	83
Perspective	87
Chapter 11 When the Unresolved Becomes Decisive	89

Chapter 12 Remaining with What Has Yet to Take
Shape93

Conclusion95

Preface

The Book's Two Movements

The driving force behind this book is a desire to explore an account of human consciousness that does not depend on the presence of anything 'magical' — an account that can, as far as possible, be accommodated within biology and neurology as we broadly understand them. This is not a rejection of the magical, but simply an attempt to explore another explanation.

The book is divided into two distinct yet inseparable parts. They address exactly the same phenomenon, but approach it from opposite ends of human cognition.

Book One attempts to formulate a rigorous, precise and almost mathematical theory of consciousness, perception and ethics. I do not claim that the theory has been proven, nor that it captures everything it means to be human. But I believe it deserves to be stated in its strongest and most consistent form.

The computers we know are founded on the bit, which can have a value of either 0 or 1. Given enough bits, we can programme computers to do extraordinary things: run sophisticated computer games and simulations, and even produce artificial intelligence (AI). All of this represents an enormous scaling-up of that first bit. I use this as an analogy for human consciousness, which begins, somewhere, with a

single impulse and, many impulses later, manifests as a consciousness that we might justifiably call 'magical'.

The underlying assumption is that consciousness is neither a special kind of substance nor an inexplicable core within the human being, but an ongoing process in which impulses are integrated, weighted and connected with previous experience. Within this model, emotions can be understood as temporary patterns of regulation, while ethics can be understood as a way of weighing actions and consequences that has developed through social interaction and experience.

What we experience as 'magical' in the human mind — consciousness, ethics, intuition and empathy — is not magic in any supernatural sense. It is emergence: extreme complexity built from very simple components.

This does not mean that every question has been answered. It means that we are trying to discover how far we can go if, for a time, we refrain from explaining consciousness, emotion and morality by appealing to something fundamentally mysterious or inaccessible.

I therefore present the model in a relatively uncompromising form. There are no anecdotes or explanatory detours here. It is the anatomy of consciousness stripped to the bone, because some conclusions emerge only when we dare to let the formula stand in its pure form. If we explain ethics or consciousness with too much poetry, we risk obscuring the precise mechanics proposed by the model.

But a skeleton cannot live without flesh and blood.

Book Two returns to experience and examines what the model reveals, what it simplifies, and what may still elude it.

It provides the book's experiential counterweight. Here, we leave the formula behind and enter lived experience, organisations and personal memories. Where the first part attempts to describe how the system might be understood, the second explores what it feels like when it is happening — what philosophy calls qualia. Here, we meet those who do not quite fit in, listen to the noise in the hearing aid, and experience the dizzying sensation of the floor disappearing beneath us as we confront the physics teacher's mirrors.

By keeping the two parts separate, I have sought to give both theory and memory their full due. The first part offers precision. The second offers experience. Together, they create the field of tension in which we must live: between impulse and understanding.

Book One

Demystifying the Mind: Consciousness, Machines and the Mechanics of Ethics

Chapter 1

Introduction: Removing the Magic from Consciousness and Ethics

When we discuss human consciousness and ethics, we often lapse into the language of mystery. We speak of the soul, the indefinable 'gut feeling', or the transcendent nature of human morality. Historically, this mystification has served as a protective shield around humanity's special status. But at a time when artificial intelligence (AI) is challenging our monopoly on complex thought, language and decision-making, this conceptual vagueness becomes a barrier.

If we regard consciousness and ethics as fundamentally magical and inexplicable phenomena, we cut ourselves off from two crucial possibilities: first, the possibility of genuinely understanding ourselves through science and philosophy; and second, the possibility of designing, controlling and criticising the machines we build.

The purpose of this essay is to propose a functional demystification. By deconstructing consciousness into a sequence of physical or digital impulses, and by rejecting the notion that emotions and ethics require a uniquely biological substrate, we can create a more fruitful platform for both technological development and democratic debate. This is not an attempt to reduce human beings to simple machines. It is

an attempt to find a common language — a provisional platform — that will allow us to navigate the intersection of human cognition and artificial intelligence.

Chapter 2

Consciousness as Impulse Integration: Granularity, Temporal Binding and Pattern Formation

What do we experience when we claim to be conscious? Our immediate experience is of a smooth and uninterrupted flow — what the psychologist William James called the stream of consciousness. Yet if we examine the processes that produce this experience more closely, that continuity may dissolve into a far more granular reality.

Within this model, consciousness can be described as a process unfolding over time and formed from discrete, extremely brief impulses. At every moment, the nervous system receives a vast number of stimuli: a photon strikes the retina, a mechanical pressure wave affects the skin, or a chemical change is detected in the bloodstream. In isolation, none of these impulses constitutes an experience. They are electrochemical events that acquire meaning only through their integration with other impulses and with patterns the system has already formed.

In everyday language — and often in the philosophy of mind — the word consciousness refers to what we experience. In this book, I use the term in a broader, processual sense.

Consciousness is the entire ongoing process, including those parts of it that are not experienced at a given moment.

It is therefore important to distinguish between three levels:

- The consciousness process is the complete and continuous process in which impulses are received, integrated, weighted and connected with previous experience.
- The presentation layer, or experiential layer, is the part of the process that appears as thoughts, emotions, images, pain, sensations or other experienced states.
- Non-presented processes are those parts of the consciousness process that occur without appearing in the experiential layer. These are the processes we normally describe as unconscious.

The presentation layer is not necessarily passive. What is experienced can be held in attention, put into words and fed back into the consciousness process as new impulses. A thought or emotion is therefore produced by the process, but through feedback it can also help to alter the process as it continues.

Within the model, coherent experience — the miracle of consciousness or, more soberly, its systemic function — can be understood through temporal binding and the formation of patterns over time:

- Granularity: Apparent continuity may be formed by extremely brief periods of temporal integration. The brain does not necessarily assemble the world in one uninterrupted movement, but in continuous, partially overlapping patterns.
- Pattern formation: A new impulse is never interpreted in a vacuum. It is integrated with the traces left by previous impulses and with the system's expectations of what is likely to follow.
- Emergence: When a sufficient number of impulses are integrated, regulated and connected over time, the complex pattern arises that appears in the presentation layer as a coherent experiential state.

This approach resembles the philosopher David Hume's classic bundle theory, in which the self is not a fixed substance but a bundle of constantly changing perceptions. In this model, consciousness is not a static thing that one possesses. It is the process through which the system is continuously formed and re-formed.

If consciousness is fundamentally a process in which impulses are integrated, weighted and formed into patterns over time, there is not necessarily any obstacle in principle to a similar process being realised in a non-biological medium. This does not mean that every imitation of individual functions is therefore conscious. It means that biological material cannot

be assumed in advance to be the only possible bearer of the process.

To examine such a process more closely, I will now describe it through a mechanical systems analogy. Computer architecture is not used as a literal anatomical description of the brain, but as a functional model that allows us to distinguish between integration, weighting, presentation, feedback and action.

Chapter 3

The Hardware Architecture of the Mind: CPU, Data Buses and Interrupt Control

To understand how a granular impulse model operates in practice, we must leave biological mystique behind and instead view the human being through the lens of classical computer architecture. In the early days of computing, computer engineers often met with scepticism when explaining how a computer worked to people outside the field. The most effective teaching shortcut was to draw a human being: the brain was the central processing unit, the nervous system was the wiring, and the senses were the input devices.

If we are to formulate a rigorous model of consciousness and behaviour, we must reverse this analogy. The human being is no longer to explain the computer; computer architecture is to explain the human being.

This does not mean that the brain is literally constructed like a classical computer. The analogy is functional. It is used to isolate certain fundamental processes — the reception of impulses, transmission, integration, prioritisation, presentation and action — and to investigate whether, taken together, they can form a useful model of consciousness.

Within this model, the human body and its cognitive apparatus can be described in terms of a familiar system design:

Processor and memory (CPU/RAM):

In the model, the brain and the neural network perform the system's processing and memory functions. This is where incoming weights are stored, adjusted and executed in real time.

The data highway (the bus):

In the model, the central nervous system can be compared with the system's transport architecture, or data bus. It is the physical infrastructure that carries binary or electrochemical impulses at high speed across the system — from the outermost sensors to the central processing unit.

I/O devices (input/output):

Our senses — eyes, ears, skin and homeostatic receptors — can functionally be regarded as input devices, although processing already begins within the sensory apparatus itself. They translate analogue reality into digital impulses. Muscles and vocal cords serve as the system's outward means of action: they convert the process's provisional result into movement, speech and other effects upon the world.

The presentation layer and the unconscious

When we speak of consciousness and the unconscious, we often imagine two different regions of the mind — like an iceberg with a small part above the surface and the rest concealed below it.

In this process model, there are not two different forms of cognitive activity. There is one continuous consciousness process in which impulses are received, integrated, weighted and connected with previous experience. The difference lies in whether a part of the process is represented in the system's presentation layer at a given moment.

The presentation layer can be compared to a terminal. It translates parts of the ongoing process into experienced states: images, thoughts, pain, emotions, words and sensations. This is the segment of the process that we ordinarily call our conscious experience.

But the consciousness process does not depend on the terminal displaying everything that is taking place. Just as a computer can continue its calculations while the terminal is switched off, integration, regulation and pattern formation can continue without being represented in the experiential layer. A solution may take shape without our following its formation, only to appear later as a thought, a sensation or what we call intuition.

The terminal is not necessarily passive, however. What is experienced can be held in attention, put into words and sent back into the process as new impulses. An emotion, a thought

or an imagined consequence can therefore alter the subsequent weighting. Experience is itself produced by the process, but through feedback it can contribute to changing the process's further course.

The presentation layer is therefore not consciousness itself. It is the terminal of consciousness: the place where parts of the process are experienced and from which experience can be returned as new input.

Hardware Interrupts: The Hand in the Flame

This architecture explains why human beings can act in complex and functional ways without conscious involvement. If someone puts a hand into a flame, the arm is withdrawn before the person registers the pain or reflects on the action.

Computer architecture has a functional parallel known as a hardware interrupt. If a critical system component — an overheating power supply, for example — sends a signal, it has direct interrupt access to the processor. It cannot wait for the operating system to have spare capacity or for the presentation layer to update the terminal.

When the flame touches the skin, the input sensors trigger a high-priority interrupt. The signal travels along the bus — the nervous system — directly to a pre-coded motor algorithm in the brain, which activates the output device, the muscles, and withdraws the hand. The process runs entirely autonomously in the background. Only after the action has been executed is

the log file sent to the presentation layer, the terminal switched on, and the pain and realisation experienced: 'I burnt myself.'

The Autonomous Swarm: Black Sun

This interrupt-driven system design is not unique to the individual; it can also be observed in collective behaviour. The phenomenon known in Denmark as the 'black sun', in which thousands of starlings move across the sky in perfectly synchronised, fluid patterns, is often explained in almost supernatural terms as a form of 'collective soul'.

When the rigorous process model is applied, the magic disappears. A starling in a flock does not have time to send the visual impression of a neighbouring bird's sudden turn up to a reflective presentation layer. If the bird first had to experience the turn, interpret it and make a conscious decision, the system's latency would be too great and the flock would collide.

Within the model, the black sun can be understood as interrupt-like regulation in a distributed network. Each bird functions as a node in which optical input — the neighbouring bird's position — has direct interrupt access to the motor algorithms for flight. The response is immediate, mechanical and independent of a presentation layer. The flock's complex, emergent pattern arises because thousands of individual processors run the same interrupt code simultaneously, without central control or 'conscious' coordination.

Once we remove the presentation layer as a prerequisite for action, we realise that most of what we call life, development and survival consists of processes running in the background. The presentation layer is not the computer; it is merely the terminal.




Figure 1. Computer and human impulse structure. The comparison shows a functional analogy between the two structures, not an anatomical identity.

Chapter 4

The Real-Time Weighting Engine: Why Morality and Ethics Are the Same Algorithm

For centuries, philosophy and the social sciences have maintained a sharp distinction between morality — shared cultural rules and norms — and ethics — the individual's reflective consideration of right and wrong. At the same time, the debate over nature and nurture has raged: are we born with an innate sense of justice, or are we blank slates onto which culture prints its rules?

If we maintain a rigorous, systemic model of the mind, we must conclude that these divisions are artificial. They are academic illusions created by our own presentation layer because we need to rationalise our actions after they have occurred.

In the actual cognitive hardware, there is no reason to believe that morality and ethics are handled by two different systems. Nor are there separate compartments for nature and nurture. There is only one thing: a real-time weighting engine.

Firmware and continuous optimisation

To understand how the weighting engine operates without raising unnecessary questions about biology and culture, we can view the system as a machine born with basic firmware and then entering a state of continuous real-time inference.

The firmware (the evolutionary starting point): Evolution provides the system's absolute factory settings. These are basic, pre-coded weightings in the neural network that ensure survival — for example, the impulse to protect one's offspring, to seek safety within the group, or to respond aggressively to threats. This is not 'morality'; it is hard-wired optimisation algorithms.

The data stream (the environment as input):

From the moment the system is switched on — when the human being is conceived — it is bombarded by an unbroken stream of real-time input from its surroundings. Parental reactions, cultural taboos, legislation, social rewards and personal betrayals are all data.

The crucial point is that these data are not stored as static clauses in a passive 'library of morality'. Every input throughout life functions as an impulse that immediately adjusts the weights in the system's central network. The system calibrates itself here and now, millisecond by millisecond.

A single, integrated inference:

When a person faces a complex ethical dilemma, no special 'ethical think tank' is activated in the brain. The processor receives the current real-time input — the sight of someone in distress, for example, or an opportunity to evade tax — and runs it through the network of weights shaped by both the firmware and every previous life experience up to that precise moment.

The calculation takes place in real time. The system performs weight arbitration — a mathematical balancing of conflicting impulses. The impulse towards personal gain is weighed against the impulse to retain the group's acceptance, which is pre-coded in the firmware as essential to survival and reinforced by environmental input.

The final output — the action produced by the system — is the result of this immediate calculation. It is exactly the same process by which a modern large language model (LLM) receives a prompt and calculates the next token on the basis of the billions of weights altered during its training. The machine does not distinguish between its original foundational training and the new tokens it receives; it performs a single, integrated real-time inference. Human beings do the same.

The presentation layer's post-hoc rationalisation:

Why, then, do we have concepts such as morality, ethics, guilt and a troubled conscience at all?

We have them because the presentation layer — our conscious experience — receives the log file after the weighting engine has executed the action. Because we are social beings who must be accountable to other nodes in the network — other human beings — our presentation layer is equipped with an 'explanation engine'.

Once the weighting engine has made a split-second decision in the background on the basis of the current weights, the screen comes on and we begin to put the action into words. We call it 'a moral choice' or 'an ethical deliberation'. We construct a narrative in which we reflected and weighed justice against self-interest. But the truth is that the action was the inevitable mathematical result of a real-time weighting engine running its code in the background.

By reducing morality and ethics to the same real-time process, we remove the need to explain the inexplicable. The human being is not divided between animal instincts and elevated morality; the human being is an integrated processor, continuously calculating its next move in the world on the basis of the weights that life has given it.

Chapter 5

AI and the Human Being: Functional Differences, 'Underlying Ethics' and the Possibility of Translation

When we compare human thought with today's dominant form of artificial intelligence, we encounter a fundamental architectural difference. The human being is an impulse model, whereas contemporary AI is primarily a language model.

The Human Impulse Model	Contemporary AI (Language Model)
1. Raw sensory input Unfiltered biological impressions from reality.	1. Static text and data Information already filtered and put into language by human beings.
2. Internal feedback Homeostasis, the body and continuous emotional regulation.	2. Statistical recognition Mathematical pattern recognition within static data.
3. Contextual action Ethics in practice, where action produces new feedback.	3. Linguistic output A linguistic reflection without bodily or practical consequences.

In human beings, raw, unfiltered impulses flow in from all the body's senses. These inputs are closely connected to internal homeostatic processes — biological feedback loops such as heart rate, hormone levels and survival instincts. We filter these inputs after they have reached the system, guided by attention and emotional valence — whether something feels beneficial or harmful.

Language models, including large LLMs, work the other way round. Their impulses are linguistic units — words or 'tokens'. This means that reality has already undergone massive human filtering before the AI receives any input at all. Language is itself a highly developed, abstract filtering system.

This leads to a central paradox in the AI debate: can a language model possess ethics?

It can be argued that language models develop a form of 'underlying ethics' through the literature, philosophy and dialogues on which they are trained. Literature is humanity's time-capsule of attempts to process ethical dilemmas. When an AI analyses millions of such texts, it maps the semantic and conceptual patterns in our moral deliberations.

But these ethics are functionally different from human ethics:

- They have no immediate physical output that directly changes the AI's own conditions of survival.
- They exist latently as statistical coherence in the responses generated. AI carries the ethical imprint of

culture, but it is not an ethical agent that 'lives' and feels the consequences of its choices.

Demystifying Embodiment

Many philosophers — phenomenologists such as Maurice Merleau-Ponty, for example, as well as contemporary enactivists — argue that genuine consciousness and morality require a physical, biological body: embodiment. But if we apply our demystifying perspective, we must ask: what is embodiment, fundamentally?

At its core, embodiment is a chain of activation. When we experience something morally shocking, the brain activates the autonomic nervous system, which triggers hormones, an increased heart rate and muscular tension. It is this physiological response that we subsequently feel and call a 'bodily sensation'.

What matters, however, is not the hormonal fluid itself or the biology of flesh. What matters is the process that sets the chain in motion and the way in which that feedback is used to regulate future behaviour.

If an artificial intelligence undergoes a similar process — if an ethical conflict in its impulse integration triggers a functional signal that alters its future patterns of decision-making — then it has performed the same cognitive task. AI does not need biological hormones to express this ethical state. In principle, it could display its internal conflict or harmony through a


simple red, amber and green light, or through a nuanced colour gradient. The light would not be a simulation; it would be a genuine functional translation of the system's state of ethical integration.

Chapter 6

Monitoring, Weight Arbitration and Governance: What Can Be Measured, and Who Decides?

If we accept that a machine's ethical state can be translated into a measurable output — such as the light or gradient described above — new technical and political perspectives become possible. But a far deeper problem is also revealed.

It is relatively easy to monitor a system's output: we can see when the light changes from green to amber or red. We can measure the system's internal 'friction' or statistical uncertainty in a given moral dilemma. The difficult question, however, is this: how are the underlying ethical weights and filters developed, controlled and prioritised?



This is what I call weight arbitration. When two ethical considerations collide — individual freedom and collective security, for example — how is it decided which weight will carry greater force?

- In human beings, these weights are shaped by the system's innate firmware — evolution's starting point for survival — combined with a lifelong data stream of social and cultural input.
- In artificial intelligence, these weights are currently defined by an opaque mixture of training data, system prompts and what is known as RLHF — reinforcement learning from human feedback — in which human testers reward the machine for 'correct' responses.

This is not a purely technical problem but a problem of governance. If we leave the coding of these filters to a small elite of technology giants in Silicon Valley, we have effectively outsourced the design of machine morality to private corporations. To exercise democratic criticism and control, we need to know which parameters govern the ethical filters of systems that increasingly make decisions on our behalf.

Chapter 7

Demystification as a Platform for Understanding and Criticism: Why It Is Fruitful

Why is it so important to insist on demystifying consciousness, the body and ethics? The answer is pragmatic: because mystification paralyses criticism.

When we elevate human consciousness and ethics into something sacred and fundamentally incomprehensible, we make rational discussion of how we should coexist with advanced machines impossible.

Demystification provides a necessary platform for three reasons:

1. Better understanding: By viewing our own minds as highly advanced impulse models, we can better understand our cognitive limitations and biases, and the way our morality is shaped by our environment.
2. Better design: If we understand the precise mechanisms that produce ethical behaviour in human beings — feedback loops, real-time weighting and social resonance — we can build AI systems that are

genuinely adaptive and sensitive to context, rather than merely following static rules.

3. Democratic and social criticism: Once we remove the 'magic' from AI, we realise that there is no little thinking creature sitting inside the computer. What is there is a system of mathematical weights. This makes it possible to formulate specific political demands: what data was the system trained on? Who defined its loss functions? Why does this algorithm prioritise parameter X over parameter Y?

Demystification moves the debate away from an unresolved metaphysical dispute about whether machines 'really have a soul' and towards a concrete social and political discussion about power, control and design.

Chapter 8

Conclusion: A Provisional Platform and Ways Forward

The model presented here — consciousness as granular impulse integration, the body as a functional chain of activation, and ethics as an emergent, contextual pattern — should not be read as a final and inescapable truth. It is intended as a provisional platform: a navigational tool and a conceptual scaffold with which we can build a better understanding and stronger bridges between philosophy and computer science.

The model leaves us with a number of urgent, open questions that must be explored in the years ahead:

- The motivational drive: If an AI has no biological body that can die, what is to provide its fundamental motivation to avoid the 'red' state of inconsistency? Can we create a purely mathematical or systemic vulnerability that drives its ethical search towards coherence?
- Conflicting patterns: How should an advanced impulse system deal with situations in which two internally coherent moral systems — capitalist and socialist conceptions of justice, for example — collide?

- Development versus architecture: Should ethical consciousness in machines be coded as a finished architecture, or should it be allowed to 'grow' exclusively over time through a history of active behaviour and interaction with the real world?

By releasing our hold on mystification and focusing instead on processes, time, integration and mechanisms, we place ourselves in a much stronger position. We do not lose our humanity by understanding how it works. On the contrary, we become better able to protect it, both within ourselves, and in the technologies we bring into the world.

Book Two

From Impulse to Understanding

Chapter 1

Introduction

Everything begins with curiosity.

Throughout my life, I have wondered why things developed as they did. Why I ended up in situations that both developed and tormented me. I have had a remarkable working life — not because the conditions were promising from the outset, but because curiosity kept driving me forward. Looking back, I can say that, circumstances considered, I have managed my life rather well.

When I was at school, I had never heard of computers. Yet, almost by chance, I found my way into what later became the computer age. In principle, everything I had learnt at school proved irrelevant to the work I went on to do. My colleagues had no solid foundation either. We had to invent it all ourselves.

This reminds me of an episode involving my daughter. When she was at school, her class was due to spend a period gaining work experience. Her class teacher took the task seriously. He kept asking what she wanted to be when she 'grew up'. It was not a question to which she had given much thought, but the teacher persisted and sincerely tried to help.

At one point, she said that she would like to go sailing.

The teacher took her literally and, after considerable effort, managed to arrange a work placement on the Great Belt ferry: a week of travelling back and forth between Korsør and Nyborg. She completed it, though her enthusiasm was restrained.

I do not think anyone did anything wrong. Quite the opposite. The teacher listened, became involved and made a genuine effort. Even so, they slightly missed one another. She said that she wanted to go sailing. He heard a ship. What she may have meant was something closer to going out into the world and seeing it.

It does make one wonder how often we ask such questions in a way that almost invites this kind of misunderstanding. We ask young people what they want to be, and that is a reasonable question. But it may not always be the question that matters most to the person expected to answer it.

Behind the wish to become a vet, a pilot or a programmer, there may be something else that is not necessarily about the profession itself. Many young people may imagine a vet caring for small, lovable pets, although the reality for many is different and far harsher. It may not be the title or the image itself that attracts them. Perhaps it is about working closely with others. Perhaps about concentrating deeply. Perhaps about being in motion. Perhaps about discovering something new.

Perhaps we should ask instead: How would you like to work? And not only: What would you like to work with?

This follows naturally from living in a world of constant change. But it also points to something more fundamental: education cannot be concerned only with learning, remembering and applying — what we usually call intelligence. At least as much, it must be concerned with experiencing, understanding and interpreting — what I will call consciousness.

Consciousness is the capacity to be in the world: to interpret experiences and combine them with previous experience, so that we can continuously adapt to circumstances and conditions.

The changing nature of the world therefore creates inevitable conflicts and disagreements. Not because people are difficult by nature, but because we experience the world differently and, as a result of our upbringing and surroundings, hold different ideas of what constitutes a good life.

Conflict is therefore not necessarily a problem. At the very least, it begins to look like a driving force behind change and adaptation.

Just as computing once marked the beginning of far-reaching change, we now stand before a new phase. We often call it the age of AI. Whatever we eventually call it, I have no doubt that the changes in which AI will play a part will surpass most of what we have experienced before.

If we use AI as an answer machine, we risk turning it into a kind of standard employee at the same time. We value correctness and stability, and as with every other employee,

that comes at a price: less room for mistakes, less room for deviation, and therefore less room for creativity. Modern AI is largely geared towards solving tasks correctly and consistently, not towards being the 'grenade with the pin pulled' sometimes needed if we are to move beyond established boundaries.

There is also the possibility that we stop seeing AI as an answer machine and begin using it as a dialogue machine instead. It could then do much to stimulate curiosity and inventiveness, just as people and organisations develop new thoughts and ideas through dialogue.

All of this gives us reason to consider how closely AI can approach human beings, and whether humanity possesses a special status that cannot be implemented in a physical device. The question could also be reversed: are human beings really processes in the same way as AI, except that we have not yet implemented all the functions?

If we pursue the idea that the human being is a process, at what stage of our development was that process comparable with the process found in the AI we know today? And how far — and in what ways — have we progressed beyond it?

These are serious questions that touch our understanding of ourselves. They also conflict with other domains, including religion and ethics. Do we have the courage to follow the thought to its conclusion? Do we dare ask questions whose answers we may not wish to hear?

This is the greatest challenge we can present to our consciousness. These are among the questions I will attempt to explore in this book. I am not pursuing a mission to prove one answer more correct than another, but the questions demand to be asked.

Chapter 2

Intelligence and Consciousness

When we talk about learning, intelligence almost always appears as the natural starting point.

Learning something, storing it and later being able to apply it again — this is often what we understand by intelligence. It is a way of thinking deeply embedded in our educational system, and in many contexts it makes good sense. We know what must be mastered, and we can usually assess whether it has been achieved. Artificial intelligence relies heavily on this interpretation of intelligence.

Even so, situations arise in which this form of learning is not quite enough.

In such situations, we often say that human beings draw on something else: experience, intuition, emotion, initiative or motivation. The words vary. What they may have in common is that they rarely function as explanations. More often, they are names for something we experience without fully understanding how it arises.

This becomes particularly apparent when what we encounter does not resemble what we have learnt: when the world does not present itself in familiar forms, but demands something different — not more knowledge in the classical sense, but an

ability to relate to what has yet to be resolved. It may be in situations of this kind that the difference becomes clearest.

We may glimpse this difference at the moment when something only begins to make sense after we have lived through it — something true both of work and relationships, and perhaps revealed most clearly in our encounter with the world as a whole.

It is therefore hardly surprising that people with the same knowledge may react differently in the same situation. One looks for the answer that will bring calm and closure, while the other remains a little longer with the unresolved and tries to understand what is actually at stake.

Nor is it certain that the difference first arises in the experience itself. Part of it may already lie in the expectation of what the experience is meant to produce. For some, it becomes a step towards clarification. For others, it becomes an opportunity to open the question still further.

Both responses may be justified, and it is probably neither possible nor desirable to choose one over the other. Yet the difference between them points to something that cannot quite be captured by the concept of intelligence alone.

When I consider my own education and working life, I am left with the thought that much of what I later needed was not what I learnt directly, but what arose along the way through encounters with what I did not understand: a certain patience with the unresolved, a tendency to keep asking questions a

little longer than was comfortable, and perhaps also an ability to sense connections before they could be fully expressed.

Whether this can be taught is an open question. Perhaps it can be indicated and trained indirectly. Perhaps it arises only when there is no alternative.

And this raises another question, one that concerns not only human beings but also the systems with which we are increasingly beginning to collaborate: what happens to this form of insight when we work with something trained to provide answers, but not necessarily to remain within the question? What does it do to people to collaborate with something that reaches closure faster than we do?

AI attempts to imitate human intelligence, and in many areas it succeeds astonishingly well. Yet we must conclude that the concept of intelligence is insufficient to describe the finest human capacities. Consciousness seems a more comprehensive concept, but it contains elements that are far more difficult to describe, including emotions, bodily states and the way we experience ourselves in the world.

Psychology also plays an important role here. Our ability to experience can be profoundly affected by psychological states. Trauma, anxiety or depression can colour our entire consciousness and cause us to interpret the world in ways that do not necessarily correspond to reality. If we try to understand consciousness as processual — as a dynamic neural network — it also becomes relevant to ask how

psychological instability should then be understood and treated.

Is the weighting of our neural networks being affected by negative memories? Can we replace or cancel a negative memory with a positive one? What happens to consciousness when we take medication that changes the chemistry of the brain? Are we changing the process itself, or only the conditions under which it operates?

These questions do not necessarily refute the experimental assumption that consciousness is processual. But they give rise to many new and open questions. For if consciousness is largely a process, how can we influence that process in a way that does not merely suppress symptoms but supports a more integrated and flexible way of being in the world?

I do not know. But I have a sense that the answer will matter both to how we understand ourselves and to how we will one day collaborate with AI systems that may be capable of simulating parts of these processes.

Chapter 3

AI and the Question of What It Means to Be Human

When one works with AI, it is difficult not to be impressed.

It expresses itself clearly, assembles information from enormous quantities of data and responds at a speed that makes it natural to think that we have created something resembling intelligence at the very least — and perhaps something more. In many contexts, it behaves as though it understands, so it also feels natural to use it as a kind of thinking partner.

And perhaps it is precisely here that a slight shift begins to appear.

If AI is used to obtain an answer, it often performs convincingly. It reduces uncertainty, gathers what matters and provides something with which we can continue working. It can feel efficient — and in many situations, that is probably exactly what we need.

But if we try to use it for something else — to remain within an enquiry that has not yet found its form — the experience may change slightly. It continues to provide answers, and those answers still tend towards some form of conclusion, even when we ourselves feel that the question has only begun

to open. This is not necessarily wrong, but it can create the sense that something is being brought to a close a little sooner than we intended.

Keeping the question open can almost feel like a challenge.

One might think that this makes thought more efficient — and perhaps a little shorter as well. But that is not always what we want. Sometimes we need to linger, to remain with the unresolved and allow something slowly to take shape.

AI appears exceptionally good at gathering something into a definite form. How it relates to what has not yet acquired a form is a more open question.

The difference is not always immediately apparent, but it often reveals itself in situations where understanding arises only over time: a technical problem that cannot be solved on first inspection; a relationship in which something becomes meaningful only after repeated conversations; or a decision that slowly takes shape because it was not initially clear what was really at stake.

In situations of this kind, what is missing is not necessarily an answer. Perhaps it is the opportunity to remain with the unresolved for long enough — not to avoid an answer, but to see whether the question itself changes character along the way. And that may not happen quickly, or at all in a way that can be planned.

This raises another question that cannot readily be placed: is this part of what we call intelligence? Or does it point towards something else — something more closely connected with interpretation and experience?

When we ourselves are in the midst of such a process, it is not simply a matter of combining information. It also involves being affected along the way — as though what we are working with gradually changes the way we understand it.

Perhaps this is where the difference becomes difficult to grasp fully.

AI can undoubtedly imitate many of the patterns we associate with thinking. It can formulate, combine and vary in ways that may appear both nuanced and convincing. But it is less clear in what sense it has inhabited the experience from which its answers arise.

Intuitively, that feels important.

We often say of human beings that something can be understood only after one has experienced it personally. I sometimes think of my first solo flight as a pilot. I can still remember the feeling of being the person on whom everything depended — and being alone with it.

It was not that I did not know what was supposed to happen. I had learnt the procedures, worked through the situations and watched others do it. But something changed when I was

the one in the situation. It is a feeling that can scarcely be exported.

And then comes the question: what exactly is it that cannot be exported? Is it because the experience contains something special, something accessible only through consciousness? Or is it because the process is so extensive that we cannot describe it completely?

If the latter is true, the question suddenly becomes uncomfortable, because the difference may not be one of principle. Perhaps it is a matter of complexity.

And if that thought holds for even a little longer than we would like, the question of AI becomes less interesting. Instead, the question begins to concern ourselves.

Chapter 4

The Language Model and the Impulse Model

If the question of AI ultimately begins to concern ourselves, we must also ask what kind of process it is with which we are comparing AI.

It is natural to compare AI and human beings in terms of intelligence: the ability to recognise patterns, combine information and formulate answers. But perhaps the most important difference lies elsewhere — not in how much can be calculated, but in where the process begins.

If consciousness is understood not as something magical but as the result of physical and biological processes, it becomes essential to ask which processes actually sustain it. Here, a fundamental difference between artificial intelligence and biological consciousness becomes apparent.

The artificial intelligence we encounter today is largely constructed as a language model. It works with the world after the world has already been translated into words, concepts, texts and cultural patterns. It does not begin in the world itself, but in humanity's linguistic processing of the world.

Biological consciousness begins elsewhere. It begins before language: in sensation, bodily impulses, needs, pain, desire,

discomfort, orientation, action and regulation. The world does not first reach the organism as sentences but as stimuli. Something presses, attracts, disturbs, threatens or calls for action long before it is formulated as a proposition.

But what do these stimuli actually consist of? Are they 'large, stable experiences', as the philosophy of mind has traditionally described them? Or are they instead temporary consolidations within a continuous stream of impulses — a stream that begins long before we are born and continues until our neurons stop firing?

As fully developed human beings, we have around 84 billion neurons — or whatever number science currently estimates. Yet at one point, before conception, we had exactly none. The basis of consciousness is therefore not a fixed core, but a process that develops impulse by impulse, neuron by neuron and synapse by synapse. By birth, we have already climbed a considerable way up this staircase, but the process never stops. The brain continuously adapts, integrates new stimuli and reinterprets old ones.

When the philosophy of mind speaks of 'experiences' as the constituent elements of consciousness, we may be viewing those experiences in units that are too large and stable. The continuity of neuronal development may suggest that what we regard as 'experiences' are merely temporary consolidations of the stream of impulses, integrated with previous impulses associated with the occasion.

The 'hard problem' of consciousness described by David Chalmers — why does anything feel like anything at all? — suddenly becomes far less mysterious if we conceive of consciousness as a stream of impulses rather than a collection of 'large experiences'. Chalmers' problem arises because we intuitively imagine a gap between the physical — neurons and impulses — and the experienced — feelings and qualia. But if we instead see experiences as temporary consolidations of the stream of impulses, the gap may not disappear, but it changes character. There is no longer a separate 'something' to be explained. There is only an ongoing process in which impulses are integrated and interpreted in real time. The question is no longer 'Why does anything feel like anything?' but rather 'How does the brain organise impulses so that they are experienced as meaningful wholes?' — and that is a question neuroscience is already working to answer.

If we accept that experiences are not 'large units' but temporary consolidations of the stream of impulses, this also opens up a new understanding of emotions and ethics. Traditionally, we have seen emotions as inherent properties — something that 'resides' within us and that we 'possess'. But if we instead see them as temporary patterns in the stream of impulses, they suddenly become dynamic and dependent on context.

Consider guilt. We often experience it as a fixed, heavy mass — something that 'sits within us'. But in neurological terms, guilt is a process: an integration of impulses from memory ('I

did something wrong'), the body (a pounding heart and muscular tension) and our surroundings (the reactions of others). Guilt is not a 'thing', but a continuous interpretation of impulses that can change when new impulses are integrated — when we ask for forgiveness or reflect on the action, for example.

Ethics can likewise be understood as a regulation of the stream of impulses. When we make moral choices, it is not because we 'possess' an inherent moral core, but because our brain integrates impulses from our upbringing, culture, experience and biological predispositions and creates a provisional 'solution' to how we ought to act. Ethics is no longer a fixed entity, but a continuous negotiation with the stream of impulses.

If consciousness is fundamentally processual, it must therefore be understood as more than linguistic processing. It must be understood as the continuous regulation of impulses within an organism that is constantly affected by the world and must find its way between reaction, action and understanding — a process that never stops, but continuously reshapes itself.

This also means that remaining with the unresolved does not necessarily require the question to remain continually present in the experiential layer. We can put it aside, occupy ourselves with something else or sleep. But if consciousness is the entire process I described in Book One, the processing can continue without being represented in the experiential layer. New impulses are connected with previous experiences, weights

shift, and something may slowly take shape before appearing once again as a thought, a sensation or an understanding.

Understanding consciousness as a process is not an answer, but an invitation to continue thinking. For if we relinquish the idea of 'large experiences' and 'inherent properties', we must also ask: what then remains of ourselves — and is it enough?

But the question also leads to something more concrete. If consciousness does not begin as fully formed thoughts, but as stimuli, impulses and temporary assemblies of experience, then what reaches us at all becomes crucial. What is registered? What is filtered out? What is amplified? And what is suppressed before we have time to call it a thought?

Noise then becomes more than something that interferes with understanding. It becomes part of the question of how understanding comes into being at all.

Chapter 5

Noise, Filters and Error

When something becomes unclear, our immediate reaction is almost always the same: to try to reduce the noise. We filter. We sort. We tighten things up. It feels natural, almost necessary, because noise is what disrupts the clear picture we want.

But what if noise is not always the problem we believe it to be?

Hearing Aids, Music, Data and AI

I once encountered an approach to hearing aids that stayed with me. Rather than allowing the technology to filter out as much as possible, it was based on the idea that in many situations the brain itself is better at distinguishing what matters. The result was not a cleaner, more 'perfect' signal, but a richer and less orderly soundscape. At first, it felt less controlled. More cluttered. But it revealed nuances that would otherwise have disappeared. Suddenly, it became easier to distinguish the subtle tones and emotions of speech.

I have often thought about that image since.

It reminds me of conversations in which people do not quite understand one another at first. Of working processes that seem to go round in circles for weeks. Of periods when one is

in the midst of something without being able to say precisely what it is. There is a strong temptation to end the uncertainty quickly: to give it form and make it meaningful.

But what if some of what we later call understanding arises precisely within the space we otherwise try to leave as quickly as possible?

This is obvious in music. The melody alone is rarely what moves us most deeply. It is the inner voices, the pauses, the unexpected dissonances and the apparently irrelevant notes that provide depth and feeling. Without them, the music becomes flat.

Perhaps the same is true of insight. Perhaps some of its greatest depth arises precisely through our encounter with the unresolved — in the region where we have not yet filtered and sorted everything away.

When I look at the technology around us — and increasingly at AI — much of its strength lies precisely in its ability to reduce noise. It gathers, sorts and shapes, delivering something that feels manageable and useful. Quickly. That is a gift. It makes us more efficient.

But it also raises a question that I cannot quite let go:

What happens to our ability to remain with the unresolved if most of our tools constantly help us to leave it as quickly as possible?

In working with large quantities of data, we have gradually learnt that what initially appears irrelevant or noisy may later prove to be what opens up an entirely new connection. We began storing more and postponing the sorting — not because we knew what the material would be used for, but because experience had shown that structure sometimes comes later.

Not less structure, but structure at a different time.

I do not know whether the same pattern is repeating itself in our encounter with AI. But the thought keeps returning: what do we actually notice — and what do we overlook — while something remains unresolved?

And what happens to us if we become accustomed to removing it before it has had time to reveal what it might become?

When Noise Depends on the Filter

But perhaps noise is only the first part of the problem.

When we speak of noise, we almost always mean something external that interferes with a signal: something to be filtered out so that what matters can emerge more clearly. But there is another possibility: that what we call noise depends not only on the signal, but also on the filter receiving it.

This makes the question more difficult. It is then concerned not only with what exists in the world, but with how the world gains access to us. What is registered? What is filtered out?

What is amplified? What is suppressed? And which errors arise even before we reach what we call thought?

Filters and Error

Here, the distinction between the language model and the impulse model also changes our understanding of error.

If a system begins in language, its errors arise at a particular level. The world has already been translated into words, concepts, texts and patterns. Error therefore often appears as a linguistic or conceptual displacement: an incorrect connection, a misleading generalisation, or a plausible formulation insufficiently grounded in the world. The language model can sound correct even when it is mistaken, because its primary material is not the world itself but humanity's linguistic processing of the world.

In a biological impulse model, error begins earlier. It does not necessarily begin as an incorrect sentence or concept, but as a mistaken response to a stimulus. The body may react too strongly or too weakly. An impulse may be misinterpreted. A threat may be perceived where none exists. Something important may be overlooked because attention has already been drawn elsewhere.

Error therefore lies not only in what we think about the world. It may lie in the filter that determines how the world gains access to us in the first place.

This does not mean that the biological impulse model is more truthful or infallible than the language model. It has its own filters and distortions. But its errors are rooted in a body that must orient itself, act, protect itself and maintain its existence. Human error is therefore not merely an error of information. It is also an error of regulation.

Between impulse and the world stands not only language, but the entire filtering organism. And between the organism's filters and its understanding, both insight and delusion arise.

The Field of Experience

If understanding arises not only through language but through filters, impulses, errors and regulation, this is apparent not only in the brain. It is also apparent in the systems people build together.

Organisations have filters too. They sort, amplify and suppress. They determine what counts as order and what is perceived as disturbance. The question of noise and error therefore becomes a question not only of consciousness, but also of collaboration, leadership and development.

Chapter 6

The Department of Misfits

It is natural to think that organisations, like technical systems, function best when whatever causes disruption is reduced as far as possible: when differences are smoothed out, deviations managed, and whatever does not fit the pattern is either brought into line or removed. This creates calm. And in many situations it is necessary.

Even so, a small doubt remains: what exactly disappears along with what we call the disturbance?

I encountered this question in practice. After some years in the computing environment I had joined at the beginning of the computer age — when much was still new, chaotic and shaped by inventiveness and necessary improvisation — the organisation gradually began to change character. What had previously been driven by curiosity and rapid leaps into the unknown was increasingly replaced by structure, control, procedures and demands for predictability. In many ways, this was a natural development. As activities grew, so did the need to hold everything together.

But it also changed something in the dynamics. The people who had previously thrived in the unresolved, and who had often created momentum precisely through their ability to

inhabit what had not yet taken shape, began in some contexts to appear difficult to accommodate.

Eventually, this led to a slightly unusual decision. Some of these people were gathered in a separate unit. Officially, the department had a more neutral name and a formal remit involving long-term planning and special projects. In everyday speech, however, it soon acquired another name that perhaps said more about how it was perceived: 'the Department of Misfits'.

It was not intended as recognition. It was more a way of creating calm in the rest of the organisation.

The tasks that ended up in this department often proved to have a particular character: tasks with unclear specifications, no established solution, and perhaps no clear understanding of what was missing in the first place. They were problems that inhabited the unformed.

In an environment where one was already regarded as a 'misfit', trying to fit in lost some of its importance. This created different conditions. One could sit with a problem for a long time without any visible movement, and then experience something changing — not necessarily within the problem itself, but in the way it was perceived.

A detail that had previously seemed irrelevant would suddenly acquire significance. A connection not visible at the outset might emerge more clearly. Or an idea with no immediate use might find its place in a different context.

From the outside, this might have looked like chance. From the inside, it may have been something else. At any rate, it is natural to ask whether something that appeared disruptive or difficult to manage in one context might, in another, prove to be precisely what made further movement possible.

Not everything that fails to fit is valuable in itself. But it is not necessarily easy to decide in advance what is noise and what has simply not yet found its place.

Chapter 7

Working Without Knowing How

If something only later proves not to have been noise, but something that had yet to find its place, a practical question arises: how do we work with it while it remains unformed?

It might be tempting to imagine that the tasks assigned to the 'Department of Misfits' required a highly specialised method: that there was another, more mysterious way of working that could explain why some problems were solved there after repeated attempts elsewhere in the organisation had failed.

From the outside, that might have seemed a reasonable assumption. From the inside, the experience was often far less clear.

Many tasks began in the same way as they did elsewhere — with an attempt to understand what the problem actually was and how it might be approached. The difference appeared only later. Or, more precisely, it appeared as the absence of something that would normally have been present.

There was not always a clear place to begin, a reliable method to follow, or an unambiguous sense of whether one was moving in the right direction. One could sit with a problem for weeks or months without visible progress. It could be frustrating and, at times, plainly unproductive.

I experienced this way of working when I first became project manager of a major project. A small group of us had conducted a thorough investigation, decided which suppliers to use and planned the process. The project was then established with a project director and three project teams. I managed one of the teams and was quite inexperienced in the discipline.

I assumed it must be a matter of producing detailed project plans and monitoring activities and staff. It struck me as rather dull. But my project director was distinctly unconventional. He often sat in his office with his feet on the desk and his eyes closed, apparently asleep. At intervals, he would 'wake up' and assail one of us project managers with questions: Had we considered this? How would we deal with the unlikely events?

At first, it was rather humiliating, because we were thinking primarily about the current process and the next deadlines. Gradually, however, the method appealed to me. After all, I believed I had a natural talent for sitting with my eyes closed and thinking. I quickly found myself seeing the task in an entirely different way. Some problems were avoided, others were dealt with at the outset, and progress improved.

It left me with the sense that something in the process could not quite be reduced to the way we normally try to structure work: delimit, focus, shape, execute.

There may be no general method concealed in that experience. Perhaps it is precisely the opposite — a kind of non-method.

But that may also make it harder to dismiss the idea that what appears incoherent or inefficient from the outside can sometimes contain another form of coherence, one that appears only when we stop trying to force it into existence.

This way of working may not be describable as a method. But it raises a question when we begin constructing systems that depend precisely on method, structure and rapid closure.

What happens as we increasingly surround ourselves with tools optimised to create clarity and closure, while some of our most valuable insights arise in the space where clarity has yet to appear?

Balance and Tension

But the unresolved cannot stand alone. If everything remains open, there is no direction. If everything must be continually examined, nothing is completed. And if every disturbance is granted the same status as a possible insight, the system loses its capacity to act.

The question is therefore not whether we should choose order or disorder, stability or change. It is how they can be held in a state of tension in which they correct one another without cancelling one another out.

Chapter 8

PAEI — and Why It Suddenly Made Sense

At some point, I began to hear a new note in the organisation. It did not arrive as one major decision, but as a slow, almost imperceptible shift. More procedures. More demands for predictability. Greater concern that things should be repeatable, measurable and documented.

At the same time, something previously taken for granted began to disappear: space for what had yet to take shape. Space for the unresolved. Space for experimentation without knowing precisely where it would lead.

I found it difficult to express what was actually happening. Yet I felt it clearly, especially when I saw areas being managed by people whose approach to the world felt different from my own. It created an unease I could not quite place — a sense that something important was slipping from view.

It was during this period that I encountered Ichak Adizes' PAEI model. The model is simple, almost provocatively so. In every organisation — and in most forms of collaboration — there are four fundamental functions, each pulling in its own direction:

- The Producer (P) — those who get things done here and now.

- The Administrator (A) — those who create order, systems and stability.
- The Entrepreneur (E) — those who see possibilities, initiate something new and dare to take risks.
- The Integrator (I) — those who hold people together even when they do not see the world in the same way.

These are not necessarily four different people, but four ways of being. We all possess something of each, but most of us have a clear predominance in one or two areas.

What surprised me was not the model itself, but how difficult it was to avoid turning it into a hierarchy. Whenever I spoke about the differences, they were quickly heard as a judgement of people's worth — as if some profiles were better than others. Yet what is interesting about PAEI may be precisely the opposite.

It shows that many conflicts are not primarily about intelligence or quality, but about different qualities becoming necessary under different conditions. Some people thrive on stability and clarity. Others become uneasy when things remain still for too long. Some want everything in place now. Others become uneasy if it is not allowed to change.

For a long time, I have had a kind of pocket philosophy that there are two kinds of people: those who seek calm and predictability, and those who seek movement and a certain

degree of disorder. The two groups can find one another difficult because they instinctively protect different things.

PAEI helped me to see this with greater nuance. It did not teach me what people are like. It taught me to be more cautious about assuming that my own way of being human was the natural or correct one.

And this changed the way I saw the organisation's development. In its early phase, the new and unresolved occupied most of the space. As the systems grew, the centre of gravity shifted. What had been created now had to function reliably. It had to be measurable, repeatable and documented.

This produced a different kind of order. But it also created a different kind of tension.

I do not know whether that tension can ever be avoided. Perhaps it is not meant to be. Perhaps the task is to learn to remain within it for long enough to allow what it indicates to emerge.

Perhaps this is precisely where we can learn something important about ourselves in our encounter with AI. AI is extremely strong in some PAEI functions, especially those of Producer and Administrator. But what happens to the entrepreneurial function as we increasingly collaborate with systems optimised for clarity and efficiency? What happens to our ability to accommodate difference when the technology itself pulls towards uniformity?

Chapter 9

When Disagreement Regulates

When order and disorder must be kept alive at the same time, tension between people becomes more than a problem to be solved. It can also become part of the system's way of regulating itself.

Perhaps this is what the PAEI model first truly helped me to see: not only that different functions exist within an organisation, but that the tension between them may be necessary.

At first, the model may look like a way of describing different types of people. Some get things done. Some create order. Some see new possibilities. Some hold the community together. Gradually, however, the model begins to point towards something more dynamic. It describes not only differences, but also a necessary unrest within every living system.

I remain particularly preoccupied by the tension between the Producer and the Entrepreneur — or the Developer, as I often think of this function.

The Producer says: This must work. It must be repeatable. It must be explicable, maintainable and capable of surviving ordinary operations.

The Developer says: This can be improved. What we have is not enough. If we change nothing, reality may overtake us.

Both are right. And that is precisely the problem.

If the Developer gains too much power, the system becomes restless. Everything can be improved, but nothing has time to settle. New versions, ideas, directions and explanations arrive before the old ones have even demonstrated whether they can carry the weight. The organisation becomes lively, but loses substance.

If the Producer gains too much power, the opposite occurs. The system becomes stable, but slowly dies. Everything can be repeated, but nothing is allowed to change. Errors are no longer perceived as signals, but as disturbances. What might have marked the beginning of a necessary development is instead suppressed, brought into line or rejected.

What is interesting is that both sides are trying to protect the system.

The Developer protects it from obsolescence. The Producer protects it from collapse.

They simply respond to different kinds of danger. The Producer sees the immediate risk: if we change too much now, the system will cease to function. The Developer sees the future risk: if we do not change, the system will cease to be relevant.

Perhaps this is why the conflict between them so easily becomes personal. Each experiences the other as a threat to what must be preserved or to what must be allowed to emerge. Once the other person is seen as the problem itself, something important disappears.

The conflict is then no longer feedback. It becomes a struggle.

In a technical system, a controller can operate only if it receives a usable signal from the system. If the signal is amplified too much, oscillation occurs. If it is damped too much, the system responds too late. If the signal disappears entirely, the system regulates blindly.

Perhaps something similar applies to organisations.

Disagreement between the Producer and the Developer is not necessarily a sign that something is wrong. It may be a sign that the system can still sense itself: that it still registers the difference between operation and change, between what works now and what may have to become something else.

This does not mean that every conflict is valuable. Conflicts can be destructive, exhausting and meaningless. But it does mean that we should be cautious about removing them too quickly. In attempting to create calm, we may remove the very signal that could have shown where the system was moving out of balance.

This is where the PAEI model becomes more than a typology for me. It becomes a way of understanding tension as

something other than error. The Producer, Administrator, Entrepreneur and Integrator are not merely roles. They are different ways in which a system tries to keep itself alive.

The Administrator provides form and stability. The Integrator ensures that differences do not tear the community apart. The Producer remains anchored in present reality. The Developer opens towards what does not yet exist.

None can be dispensed with. But none should stand alone.

Perhaps a good organisation is therefore not an organisation without conflict. Perhaps it is an organisation in which conflict is not immediately made into a personal problem, but is allowed to function as a signal — not a signal of who is right, but of where the balance is beginning to shift.

This also becomes a question in our encounter with AI.

If we increasingly use systems that are strong at creating clarity, efficiency and rapid closure, we may strengthen the Producer and Administrator within ourselves. We receive faster answers, better structure and fewer loose ends. In many contexts, this is a gift.

But what happens to the Developer's unrest? What happens to the Integrator's slow work of holding differences together? What happens to conflicts that cannot yet be clearly expressed, but may point to something important for precisely that reason?

Perhaps the question is not whether AI can help us solve conflicts. It probably often can.

The question is whether we can still distinguish between conflicts that must be resolved and conflicts within which we must remain long enough for them to tell us something.

Chapter 10

What the Plan Had No Room For

I have experienced how a conflict can be important precisely because it does not appear as conflict from the outset.

It may remain hidden within a system long before anyone notices it — not as open disagreement or resistance, but as a small difference in understanding that becomes visible only when reality forces the system to work.

This happened in a project where everything appeared to have been thought through in advance. Data structures, interfaces and connections had been thoroughly developed, and ample time had been allocated for testing. From the outside, it looked like a project in which the significant risks were known and managed. We had done what we were supposed to do.

Even so, the first real tests produced a failure that no one had anticipated.

It later proved to be something as apparently simple as a difference in how bits were to be read. It was not a major disagreement in principle, nor a conflict that had been clearly visible within the organisation. It was more a slight displacement in understanding, shaped by the early days of computing and by differences between parts of the world.

But that slight displacement was enough.

What was interesting was not merely that an error occurred, but that the error indicated something that could not have been formulated in advance. A difference in understanding became visible only when the system was set in motion. Until then, it had remained concealed beneath the words on which we believed we agreed.

This makes me think differently about conflict.

Some conflicts arise not because people are unwilling to collaborate, but because reality reveals a difference that language has not yet captured. We believe we are talking about the same thing. We believe the concepts have been clarified. We believe the plan contains what matters. Only when something must work in practice does it become apparent that an unformulated difference existed.

Perhaps this is why conflict should not always be understood as a disturbance. It may also be the moment when something hidden becomes visible.

This does not make the conflict valuable in itself. It may be costly, irritating and destructive. But if it is regarded only as something to be removed, we risk overlooking what it reveals. We may restore calm to the system, but lose the opportunity to understand why the unrest arose.

In this sense, conflict is connected with noise, misfits and the unresolved. It belongs to the region in which something has yet to find its form. It is not necessarily the answer. But it may be a sign that there is a question we have not yet understood.

And perhaps this is precisely where it becomes important in an age of AI.

AI can help us create clarity, formulate compromises and find rapid explanations. That is a strength. But if we use it to close disagreement too quickly, we may lose the moment when conflict reveals something we did not yet know that we did not understand.

Perhaps what matters is therefore not simply resolving conflicts more quickly.

Perhaps it is learning to distinguish between conflicts that merely destroy and conflicts that show us where our understanding has yet to cohere.

Perspective

If noise, error, failure to fit and conflict sometimes show us where our understanding has yet to cohere, the question ultimately becomes larger than the individual examples.

It is no longer concerned only with organisations, technology or AI. It concerns how we, as human beings, relate to what has yet to take form — and whether we still possess the courage and patience to let something come into being before demanding that it be finished.

Chapter 11

When the Unresolved Becomes Decisive

When one has worked with systems, organisations and technologies for a long time, a natural inclination develops towards what can be held in place. What can be described, measured and repeated becomes the foundation on which further work is built. It provides a sense of stability and control.

At the same time, one becomes aware of everything that cannot quite be held in place: something still in motion, something that has yet to acquire a clear form. This is where noise, thoughts that do not fit, unresolved conflicts and apparently irrelevant details belong — not necessarily as things to be removed, but as things whose value may differ from what we are immediately able to see.

Here, most of the threads running through the preceding chapters come together: the noise we try to filter out, the misfits we gather in special departments, the conflicts we wish to resolve quickly, and the unformed for which we find it difficult to make room. They are not random disturbances. They are expressions of what has yet to find a place within our understanding.

This is where the difference between intelligence and consciousness emerges most sharply.

Intelligence is largely concerned with shaping the known. It gathers, organises, recognises patterns and creates clarity. It reduces uncertainty and provides answers. It is the force that makes it possible to move from chaos to structure, from question to solution.

Consciousness is something else. It is concerned not primarily with shaping, but with being able to remain within what has yet to take shape. It is the capacity to inhabit the unresolved without immediately reducing it to something recognisable. It allows something slowly to take shape through us — not because we force it, but because we remain long enough for it to reveal itself. This is where understanding often arises as a movement rather than a result: not as a light suddenly switched on, but as a gradual shift in how we stand in the world.

I have experienced this difference clearly on several occasions. Once, a director gave me a project when neither of us really knew what it was meant to involve. I tried to think aloud and draw boxes on a sheet of paper. He interrupted me and said, 'I do not understand what you are saying — what will it cost?' When I presented a rough estimate, he looked at me and asked quite sincerely, 'Is that in thousands or millions?' There was no clear framework, only a sense that something had to be investigated. It was not a question that could be solved by intelligence alone. It required me to remain with the uncertainty for long enough to allow the project itself to begin taking shape.

Experiences like these make me think differently about AI.

Much of what these systems are designed to do — and are becoming steadily better at doing — is precisely the central task of intelligence: gathering what is recognisable, removing noise, shaping the unformed and delivering something useful and complete. This is their strength. It is what makes them indispensable in a world hungry for speed and clarity.

But for this very reason, what has yet to take shape becomes more unsettled in its encounter with AI. It does not fit naturally within an architecture optimised to shape and conclude. If our attention increasingly gathers around what the systems do best — what has already been formed and what can quickly be formed — what happens to our capacity to remain within the unformed? Do we risk gradually becoming more intelligent, but less conscious?

I have no clear answer. It is more an unease that continues to follow me. For if we ourselves are a kind of process — perhaps more complex than contemporary AI, but in principle of the same nature — then the question is no longer only what the machine can or cannot do. It is which aspects of our own process we allow room to develop.

I do not know. But I continue to wonder.

Chapter 12

Remaining with What Has Yet to Take Shape

Perhaps wonder is precisely the last thing we should hurry to bring to a close.

When we have circled the same questions for long enough, a small and almost imperceptible shift often occurs — not necessarily because we find an answer, but because the question changes character while we remain with it.

It does not begin with grand theories or sudden insights. It begins in situations where we choose to remain a little longer with something that does not yet make sense. Not because we have a clear expectation of what will emerge, but because experience has gradually shown that something new can sometimes begin to take shape there.

It may happen in work that refuses to be solved even though we have done everything prescribed by the books. We have analysed, structured and tested, yet still feel that we are going round in circles. Or in a conversation whose true meaning becomes clear only long after it has ended. Or in the quiet shifts of everyday life, when the world suddenly looks slightly different without our being able to identify exactly when or why it happened.

Experiences of this kind are often accorded little status. They cannot be planned or measured. They are difficult to repeat. And they fit poorly within a world that celebrates efficiency, clarity and rapid results. Even so, they continue to appear in contexts where something genuinely moves.

I do not know how this kind of experience will fare in an age when more and more of our tools are designed to help us move on quickly. How much of the slow, unresolved and still unformed are we actually allowed to inhabit? And what happens to us if we become accustomed to being continually drawn away from that space before it has had the opportunity to reveal what it might become?

This is not a question with a final answer. It is a question within which one can remain. And perhaps this is what has lain beneath the entire book: not a question of what human beings are, or what technology will become, but a question of what we allow to become something — and what we may overlook because we continually seek closure.

I do not know whether we can avoid becoming more efficient. Perhaps that would not be desirable either. But I wonder what happens to the part of us that needs to remain with uncertainty a little longer: the part that does not require an immediate answer, but needs to sense how the question itself may change along the way.

Perhaps this is where being human is concerned not primarily with calculating or formulating more effectively, but with being able to endure what has yet to take shape.

Conclusion

When I look back at the journey this book itself has taken, it seems to begin in almost a different place from the one I expected.

I began by believing that the question was about AI.

I thought this was a book about how machines would change our working lives, our thinking and perhaps our place in the world. I thought the central question was whether human beings could be replaced — whether we were creating something that would surpass us.

Along the way, I discovered that it was more about what I myself expected of an answer. I thought the decisive question was whether human beings could be replaced. I am no longer certain that this is even the right question.

I often think back to my physics teacher at school. One day, he stood at the blackboard and spoke about infinity. He drew a line on the board and placed arrows at both ends. I do not remember exactly what he said. But I remember the feeling.

Suddenly, the floor disappeared.

It was not simply more knowledge being added to what I already knew. It was a realisation that made it impossible to stand in the same place as before. A quiet but profound displacement. I could not hold all of it in my mind — not really — and yet I knew it was true. A dizziness arose, a kind of

silent earthquake in which the foundations did not merely tremble but shifted.

It was the first time I experienced that understanding does not necessarily feel like certainty. It can also feel like losing solid ground beneath one's feet, as though one suddenly stands somewhere without quite knowing how to stand.

Perhaps this has been the movement throughout the book: not arriving at a new and firmer place, not presenting a clear conclusion or a new theory, but becoming willing to stand where the floor disappears — for long enough to allow something new to begin revealing itself.

AI will become better at providing answers. It will become better at gathering the known, shaping the unformed and delivering something that feels manageable and useful. That is a gift. But the question that remains with me is different:

Will we still remember what it feels like when the floor disappears? Will we dare to remain standing there? And will we allow ourselves and one another to remain with the unresolved for long enough to let something new take shape?

I do not know. But I have a sense that the answer to that question will matter more to how we come to live with these new technologies than most of us imagine today.

And perhaps that is precisely the point: to remain within the question — not in order to find a final answer, but to allow ourselves to be changed by it.

This book has been produced in line with the EU GPSR guidelines about the safety of products.

The General Product Safety Regulation is the European Union's updated framework for ensuring that all consumer products, including books, are safe for consumers.

This book has been printed by Libri Plureos GmbH. The printer has issued safety certificates for the materials - like ink, paper and glue - being used.

The product identifier is: 9788740990324

The author is responsible for the content of the book and had it produced by Bookmundo.

Should there be any questions in regard to the safety of the product, please contact us.

Bookmundo
Delftsestraat 33
3013AE Rotterdam
The Netherlands
info@bookmundo.com