

Hidden Signals

From Paracelsus to Plant Sensors



Peter A. Gloor

Manuscript · Working edition for teaching

Contents

Prologue · The Vault — or: A Plant That Notices Someone.....	3
The Guiding Idea	5
Chapter 1 · Honest Signals — what humans, animals and plants give away.....	6
Chapter 2 · The Visionaries — four dreams that became measurable	9
Chapter 3 · The AI Toolbox — from trees to transformers.....	13
Part I.....	17
Chapter 4 · The Bonfire of Basel.....	18
Chapter 5 · The Chemistry of Teams — honest signals in humans	20
Chapter 6 · The Book of Nature — the five symbionts and their roles.....	23
Chapter 7 · Reading Honest Signals — the Symbiont tool	25
Part II.....	35
Chapter 8 · The Invisible College — from brotherhood to open source.....	36
Chapter 9 · The Face as a Reading Device.....	39
Chapter 10 · The Voice as a Signature	42
Part III	44
Chapter 11 · Plants Feel — and how you measure it	45
Chapter 12 · The Language of Animals.....	50
Chapter 13 · Plants That Read Us.....	54
Part IV.....	58
Chapter 14 · Psychohistory — from the individual to the crowd.....	59
Chapter 15 · Swarm Creativity — how groups create something new.....	62
Chapter 16 · The Digital Twin — and where the journey ends	64
Part V.....	67
Chapter 17 · A Shared Timeline	68
Chapter 18 · The Inverted Ladder	69
Chapter 19 · The Open Vault — a Word to You	70
Sources and Further Reading.....	72
About the Author	74

Prologue · The Vault — or: A Plant That Notices Someone

Post CXX annos patebo — “After 120 years I shall open.”

— *inscription over the tomb of Christian Rosenkreutz*

In the year 1604, so runs a story that may or may not have happened, a handful of men discovered a hidden door in the wall of their house. Behind it lay a seven-sided chamber, lit by an artificial sun that had burned without fuel for one hundred and twenty years. In its centre rested a perfectly preserved body, and in the dead man’s hand lay a book. Over the door stood the words: *Post CXX annos patebo* — “After 120 years I shall open.”

The dead man was Christian Rosenkreutz, founder of the Brotherhood of the Rosy Cross, and the story appeared in a pamphlet called the *Fama Fraternitatis*, printed anonymously in Kassel in 1614. It described a fellowship of learned physicians who had cracked the code of nature — who could read the book of nature directly, without the detour through Aristotle, Galen or the Church. The vault was the archive of this knowledge; the artificial sun the light of understanding that, once kindled, never goes out again. And the one hundred and twenty years were a wager on the future: that one day the instruments and the minds would come that could read this book.

The pamphlet was almost certainly an invention — probably written by the young theologian Johann Valentin Andreae, as provocation and thought experiment. But inventions, as we shall see, have a way of becoming real. The Rosicrucian writings helped inspire the founding of the Royal Society; they planted the idea that nature is a text — written in a language that is neither Latin nor Greek but something more fundamental — and that whoever learns this language gains an access that no amount of scholarship can replace.

This book tells how the vault opens.

* * *

More than four centuries later, in the spring of 2026, a second vault stands in a great hall in Dietikon. It is not made of stone but of glass, cable and code. It is called the Phänomena, and thousands of people walk through it every day.

At one of the stations stands a table, and on the table stands an ordinary potted plant — a Kalanchoe, of the kind that grows on many windowsills. A small adhesive electrode is clipped to one of its leaves, the same kind you otherwise stick to a patient’s chest to measure the heart. A thin cable leads to a small box, the box to a screen. On the screen a green line crawls from left to right.

A girl steps up. She has just come out of the noise of the hall; her breathing is quick. In the moment she leans over the plant, the line changes: a small deflection, a swell that was not there before. The girl steps back, breathes more slowly, grows calm — and the line calms with her.

The plant has no eyes, no ears, no nervous system in our sense. And yet something has happened that cannot be explained away: **the plant has noticed the girl.**

We leave the two of them standing there a moment — the girl and the plant that have noticed each other. Hold onto this image. At the end of the book, when you know how such a thing is measured and where its limits lie, we will return to this table. Perhaps by then you will be standing there yourself.

* * *

The vault opens — and the light inside it is bioelectric.

What the plant is sending here is neither magic nor metaphor but signal processing: run by a living thing that has mastered it for four hundred million years, read by instruments barely ten years old, interpreted by learning methods not yet five years old. And it is an *honest signal* — one the plant does not choose, does not embellish, does not filter. That is exactly what this book is about: the honest signals that everything living sends out ceaselessly, and the tools that finally let us hear them.

The Guiding Idea

Aristotle's old ladder of nature — man at the top, then animal, plant, fungus, stone — which dominated Western science for 2,000 years, was never a ranking of worth but a map of our understanding. We understand humans best, and yet not even them entirely; the animal better than the plant and the fungus. But the further we “descend” the ladder, the greater grows not the insignificance but the wonder: each kingdom is perfect in its own way, and more like us than we care to admit. Humans are not the crown of creation but the only beings who can build the tools to finally hear the others — and that is not a crown but a responsibility. *Hidden Signals* teaches you to listen across the kingdoms.

The next three chapters lay the conceptual groundwork of the book: what we read (the honest signals of humans, animals and plants), why (the four visionaries who dreamed it) and how (the AI toolbox). The narrative Parts I to V build on this.

Chapter 1 • Honest Signals — what humans, animals and plants give away

WHAT THIS CHAPTER IS ABOUT

Beneath the words — and with animals and plants entirely without words — lies a second layer of communication: signals that do not send, but give away. This chapter introduces the book's central concept, the honest signal, and explains why artificial intelligence can read it today in humans, animals and plants — like a new sense organ.

A face flushes, a voice cracks, a dog's ears flatten, and the tension in a leaf jumps upward. None of these signs was meant as a message — and yet each tells the truth. Beneath the level of words, and with animals and plants entirely without words, lies a second level of communication: signals that do not send, but give away. This book is about that hidden level — and about how you learn to read it.

What makes a signal honest

Not every sign is honest. Words can lie; a smile can be feigned. A signal is honest when it arises involuntarily and is hard to fake — when it carries the inner state reliably to the outside, whether the sender wants it or not. The biologist Amotz Zahavi explained the peacock's tail this way: it is credible because only a truly strong animal can afford such splendour. In humans it is the blush, the tremor in the voice, the quickened heartbeat — precisely the things we cannot control. Honest signals are honest because they can barely lie.

Human, animal and plant

Human. This field owes its name to the book *Honest Signals* by Sandy Pentland at the MIT Media Lab. But that the outcome of negotiations or job interviews can often be predicted better from tone, gesture and the rhythm of give-and-take in conversation than from the words themselves became clear only once these signals were measured systematically — with wearable sensors, with the analysis of email networks, with body sensors. A large part of this measurement research grew up in parallel at the MIT Center for Collective Intelligence and forms the foundation of this book. Precisely because we do not control these signals, they give away our true state. Facial expression, voice, heart-rate variability and the response times in a chat are such honest signals — and it is exactly these that this book reads in the activities on face, voice, heart and group.

Animals. An animal cannot tell you how it is doing — and sends it out ceaselessly all the same. A cat's meow, the pained expression of a horse, the posture of a dog are honest signals of a being without words. To read them is to understand a language that was never meant as language.

Plants. And the plants? They have no face and no voice. But electrical signals run through their tissue that change when they are touched, when the light shifts, when they come under stress — and perhaps, and this is the open question of this programme, when we are near them too. Here caution is called for: a voltage curve is not a feeling, and we ascribe nothing to the plant. But it is an honest, involuntary trace of its state — the most honest language a plant has. The Biolingo sensor makes it visible for the first time.

Why only now? — The machine’s new senses

These signals are not new. The blush, the bat’s call, the click of the tomato — they have always existed. What is new is only that we can finally reach and interpret them. Artificial intelligence gives us something like additional senses: eyes that see more finely, ears that hear further, and a sense of touch for the electrical.

The eye of the machine. Image recognition is the greatest leap of recent AI. A deep neural network, trained on millions of photos, reads the tiny movements of a face — the pull of a corner of the mouth, the tension around the eyes — and translates them into probabilities for emotions. The same works with the body: with posture, with gait, with the set of a dog’s ears or the lowered head of a horse. What an experienced person can only guess, the machine now measures frame by frame — and in animals, which can tell us nothing, it becomes readable for the first time whether they feel pain or fear.

The ear of the machine. Sound, too, can be read. Turn a tone into a spectrogram — an image of its frequencies over time — and the same image recognition takes hold. In the human voice it reads pitch, pauses and tremor: the stress beneath the words. And here a surprise awaits. Many animals we thought mute merely speak outside our hearing. Bats call in the ultrasonic, far above what we hear; whales and elephants communicate in deep infrasound, far below. It was not the animals who were silent — our ear was too narrow. With the right microphone and a learning model, this hidden babble of voices becomes audible and distinguishable.

Even plants make sounds. A thirsty or injured tomato plant emits rapid ultrasonic clicks — a staccato, high above our hearing threshold. Researchers showed in 2023 that a learning model not only recognises these clicks but can distinguish a wilting plant from a freshly cut one. What looked like silence was a conversation we never overheard.

The third sense — measuring. And where there is neither face nor voice, the electrical trace remains. Fine changes in voltage run through every plant; a simple sensor on a leaf makes them visible. This is exactly the signal the Biolingo sensor reads, the one you will work with yourself in this book — a language that knows neither words nor sound, only volts.

The shared trick is astonishingly simple: almost every signal — light, sound, voltage — can be turned into numbers or into an image, and then the same pattern-recognising machinery goes to work. The model learns the patterns itself, from large quantities of

examples; cheap sensors and fast computers did the rest. That is why only now: not because the signals are new, but because we can reach and read them for the first time. AI here is a new sense organ — it extends our perception into frequencies, speeds and subtleties our body was never built for.

The same principle, the same toolbox

Across all three kingdoms it is the same task: a signal that gives away an inner state, and the question of what it means. A face, a meow, a leaf — different in substance, the same in principle. That is why the same tool works for all three. How a signal becomes a statement about a state is shown by the chapter “The AI Toolbox.” This chapter says *what* we read; that one, *how* we read it.

The honest limit

One last, important point. To read honest signals is to read the outside and infer the inside — not to read minds. We measure a trace and estimate a probability; the experience itself remains the other’s. And precisely because honest signals can barely lie, reading them is a kind of access that no one has expressly granted — in humans, animals and plants alike. This honesty obliges us: whoever reads it bears responsibility for what they make of it.

Chapter 2 • The Visionaries — four dreams that became measurable

WHAT THIS CHAPTER IS ABOUT

Four dreamers across five centuries — a rebellious physician, an invented brotherhood, a science-fiction author and a mathematician of the future — imagined something that sounded like poetry then and is measurable today. This chapter introduces them and shows how each of the first four parts of this book stands under the sign of one of them. At the end stands the thought that carries the whole book: you read — and you are read.

Before we build sensors and write programs, a short detour by way of four people who never met. Almost five centuries lie between the first and the last; two of them really lived, two are figures from novels. And yet a single bold thought unites them: each dreamed of reading a hidden layer of the world — long before the tools for it existed.

These four are the visionaries of this book. Its four parts are named after them, and whoever understands their ideas understands why we do what we do: why we clip a plant to an amplifier, why we read a chat like a network, why we listen to a cat's meow and why we measure the mood of an entire crowd. Let us take them in turn.

The First — Paracelsus (1493–1541)

Summer 1527, Basel. A young town physician steps before the university and throws the most venerable textbooks of medicine — the works of Galen and Avicenna, untouchable for over a thousand years — into the midsummer fire. A scandal. The man's name is Theophrastus Bombastus von Hohenheim, born in Einsiedeln in Switzerland, and he has given himself the name Paracelsus: “beyond Celsus,” greater than the famous Roman physician. Modesty was not his strength.

His heretical thesis. The old books, said Paracelsus, contain false models of nature. Whoever wants to heal should not memorise the authorities but question nature herself — the plants, the minerals, the sick body before him. At a time when knowledge meant quoting the right books, this was an attack on the whole edifice. He trusted the thing more than the text about the thing.

Behind this stood an entire worldview. Paracelsus was convinced that every living being carries its inner nature outward — as a kind of sign, a “signature,” that can be read if only you look closely enough. The world for him was a book in which the human is only a small paragraph: the microcosm that mirrors the great cosmos in miniature.

Where he went wrong. His doctrine of signatures often led him astray: because a walnut looks like a brain, he thought it must be good for the head — which is of course nonsense. The concrete interpretations were mostly wrong. But the basic stance — read

nature directly, observe, test, rather than parroting authorities — is exactly the stance from which modern natural science later grew.

And this is precisely where this book begins. We take Paracelsus at his word — only with instruments he never had. The sensor you will build reads the electrical “signature” of a leaf. A camera reads the stirring in a face, a microphone the tension in a voice. These are honest signals: signs that a body sends out without lying.

“Do not read the old books — read nature herself.” — after Paracelsus

The Second — the Rosicrucians (1614)

Almost a hundred years later, in 1614, a mysterious pamphlet surfaces in Kassel: the *Fama Fraternitatis*, the “Call of the Brotherhood.” It announces a secret society — the Rosicrucians — founded by one Christian Rosenkreutz, who is said to have travelled the world, gathered the knowledge of all lands, and left behind an order. Its members heal the sick for free, recognise one another by secret signs, and share their knowledge over an invisible network, across every border and confession.

The catch. This brotherhood almost certainly never existed. The writings probably came from a handful of young theologians around Johann Valentin Andreae — it was an invention, a literary prank with a serious core. But it was a fiction that worked. All over Europe, scholars wrote letters begging to be admitted to the order. The idea electrified the age: that knowledge should not be hoarded but shared; that a network of clever minds across national borders can achieve more than any single scholar in his chamber.

Out of such dreams grew something very real. The notion of an “invisible college” — of researchers freely exchanging their results by letter — inspired in the seventeenth century the circle of natural philosophers from which the Royal Society eventually emerged, one of the first scientific societies in the world. The open, shared science we know today has one of its roots here.

Why this belongs here. Invisible colleges have long since become real — only today they are called Wikipedia, Linux, or the open-source community. Thousands of people who never see one another build something greater together. In this book you learn to read such a group: who contributes, who connects, who is the secret queen bee in the swarm? We no longer read only the single human but the network between the humans.

“Share your knowledge freely — across every border and wall.” — after the *Fama Fraternitatis*

The Third — Douglas Adams (1952–2001)

A leap into the twentieth century. In 1979 *The Hitchhiker’s Guide to the Galaxy* appears, perhaps the funniest science-fiction novel of all time, written by the Briton

Douglas Adams. In it there is a tiny yellow creature: the Babel fish. You push it into your ear, and instantly you understand every language in the universe — every word spoken anywhere reaches you translated. At a stroke, no living being in the cosmos is mute to you.

The dream behind it. Adams meant it with a wink, but the thought runs deep: the end of the language barrier. Not only between humans who speak different tongues, but between entirely different kinds of beings. Perfect understanding across every boundary.

And language is not only human. A dog that flattens its ears is saying something. A cat has different sounds for annoyance, begging and contentment. A plant that is touched sends an electrical wave through its leaves. All of this is “language” — only we have had no Babel fish for it. The decisive question of this part of the book: what if a machine could be that Babel fish?

The modern Babel fish. This is exactly what artificial intelligence attempts. A model that has learned patterns from thousands of examples can translate the facial expression of a dog, the sound of a cat or the signal of a plant into something we understand. No perfect Babel fish — but a beginning.

“Push the Babel fish into your ear — and no living being stays mute to you.” — after Douglas Adams

The Fourth — Isaac Asimov (1920–1992)

Our last visionary is the great science-fiction author Isaac Asimov. In his *Foundation* cycle, begun in 1951, the character Hari Seldon invents a new science: psychohistory. Its basic idea is astonishing and unsettling at once. A single human, Seldon says, you cannot predict — he is too free, too capricious. But the behaviour of billions of humans can be calculated with mathematical certainty. From the statistics of the mass, Seldon reads the future of an entire galactic empire, millennia in advance.

The idea. The crowd is more predictable than the individual. What is pure chance for one person becomes, across millions, a smooth curve. Whoever observes enough people at once sees patterns that stay invisible in the individual — and can read from them the rough direction of the future.

As far off as that sounds — on a small scale we already do it. Pollsters predict elections without knowing every voter. Epidemiologists calculate the spread of a disease. And indeed, from the mood of millions of short messages you can read, a little, where the stock market will tend the next day. The individual remains a riddle; the swarm gives itself away.

What we make of it. In Part IV you no longer read a single human but the crowd: the mood of a swarm from social-media data, collective feelings laid against events. And at

the end stands the boldest application — the digital twin, a model that forms an image of you out of many of your traces.

“You cannot predict the single human — the crowd, you can.” — after Isaac Asimov

When the instruments catch up with the dreams

Four imaginations, one shared dream: that the hidden layers of the world can be read. Paracelsus wanted to read the signs of the body, the Rosicrucians the shared mind of a community, Adams the unknown tongues of other beings, Asimov the logic of the crowd. Each dreamed his dream before the tools existed.

Today they exist. Sensors that measure more finely than any hand; cameras and microphones that miss nothing; and above all machine learning, which draws the patterns out of mountains of examples. For the first time in history the instruments catch up with the dreams — and you will not only marvel at the tools but build and operate them yourself.

A word of caution. The same tools with which you read others read you as well. You read — and you are read. That is why the last lesson of this book is not a technique but a stance: wonder and scepticism at once, and the responsibility that follows from the ability.

The four parts of this book follow the four visionaries, each with its own chapters and activities. A fifth part brings it all together at the end. The thread is already drawn — from a physician who burned books to a mathematician of the future.

Chapter 3 · The AI Toolbox — from trees to transformers

WHAT THIS CHAPTER IS ABOUT

Behind the many names — random forest, neural network, transformer, Whisper — lies a simple order. This chapter reduces it to two questions (representation and time), explains the spectrogram trick by which a breakthrough in seeing conquered hearing and measuring too, and draws the map on which every activity in the book has its place.

“Artificial intelligence” sounds like a single, mysterious thing. In truth it is a toolbox, and in this book you have already held several of its tools: a decision tree that sorts cat sounds, a deep network that reads a mood from a face, a language model that analyses a chat. At first glance they have little in common. This chapter shows that a simple order lies behind them — and that you can get to the bottom of any AI method with just two questions.

Question 1: How is the signal represented?

The first question decides which “weight class” a method belongs to — and thus how much computing power and how much data it needs.

Features and trees. In the lightest case you decide for yourself which numbers matter — the pitch of a sound, the variation of the heartbeat, the ratio of give-and-take in a chat — and leave the model only the sorting. Methods like random forest or XGBoost string many small decision trees together for this. They run on any laptop, need little data, and are transparent: afterwards you can read off which feature tipped the scales.

Images and deep networks. One step harder: you let the model find the features itself. For that you turn the signal into an image — a spectrogram from a sound, the voltage curve of a plant, or simply the photo of a face — and hand it to a deep neural network (a CNN). This is more powerful but demands more data, ideally a graphics card, and you lose transparency: the network finds patterns no one has named.

Sequences and transformers. Hardest are the methods that process an entire ordered stream — a sequence of tones, a sequence of words. Transformer models like Whisper or the large language models keep the context in view and deliver remarkable results for it. They are at once the most computation-hungry, the most data-hungry and the most opaque.

Running through these three tiers is a single trade: the more you leave to the model, the more powerful it becomes — and the more hardware and data it costs, and the less you can see inside (see Figure 1).

The spectrogram trick

The greatest leap of recent AI came in seeing: networks trained on millions of images recognise, with remarkable certainty, what is in a photo. The notable thing is that this leap did not stay with the image. It transferred to sound — and even to the electrical voltage of a plant. The reason is a single, simple trick: you turn the signal into an image.

A sound is nothing but a trembling in time, a rise and fall of air pressure. An old mathematical idea — Fourier analysis — states that every such trembling can be broken down into pure tones of different pitch. Do this not for the whole sound at once but in many short, successive time windows, and for each moment you obtain how much energy lies at which pitch. Line these windows up side by side and a two-dimensional image emerges: time runs left to right, frequency from bottom to top, and brightness shows how loud it is there. This image is called a spectrogram — a photo of the sound.

And once the sound is an image, you no longer need a dedicated “sound AI.” You hand the spectrogram to the same image-recognition network that learned to tell a cat from a dog — and it now learns to separate a tense voice from a calm one, or the click of a thirsty tomato from that of a cut one. The trick does not end with sound. The voltage curve of a plant, too, is only a trembling in time. Slide the same Fourier window over it and it becomes a spectrogram — and in exactly this way, later in the book, an electrical curve turns into a classifiable “plant emotion.”

This is the deeper reason a single breakthrough reached so far. By translating sound and voltage into the language of images, the eye of the machine becomes a universal reader: the same seeing that was trained on photos now hears voices and registers the electrical pulse of a plant. One sense, lent to all the others.

Question 2: How is time handled?

The second question is independent of the first: does the method take the temporal course into account?

Snapshot. Some models make a decision for each input on its own, without memory: a face — a mood, an image of a horse — an emotion.

Sequence. Other models process the order itself; time sits inside the model. Whisper does not hear individual sounds but the flow of speech.

Course from snapshots. And here the fine but important distinction: you can also run a snapshot model once a second and only afterwards look at the course — say, the mood of a face over an hour. Then time sits not in the model but in your evaluation. This is often the cheapest route to a temporal analysis: a light model, applied many times.

Incidentally, not every temporal analysis needs AI at all. You find a plant’s day-night rhythm with a Fourier analysis — pure signal processing, without any learning. The

most honest question therefore often stands at the beginning: do I really need a learning model here?

A map

Figure 1 sorts the methods of this book into the three weight classes and shows at the same time how each handles time. The cat sounds (features + tree, snapshot) sit at the very top, light and transparent; Whisper and the language models (sequence + transformer) at the very bottom, powerful and opaque. The face mirror sits in the middle — a snapshot model that we assemble into a course.

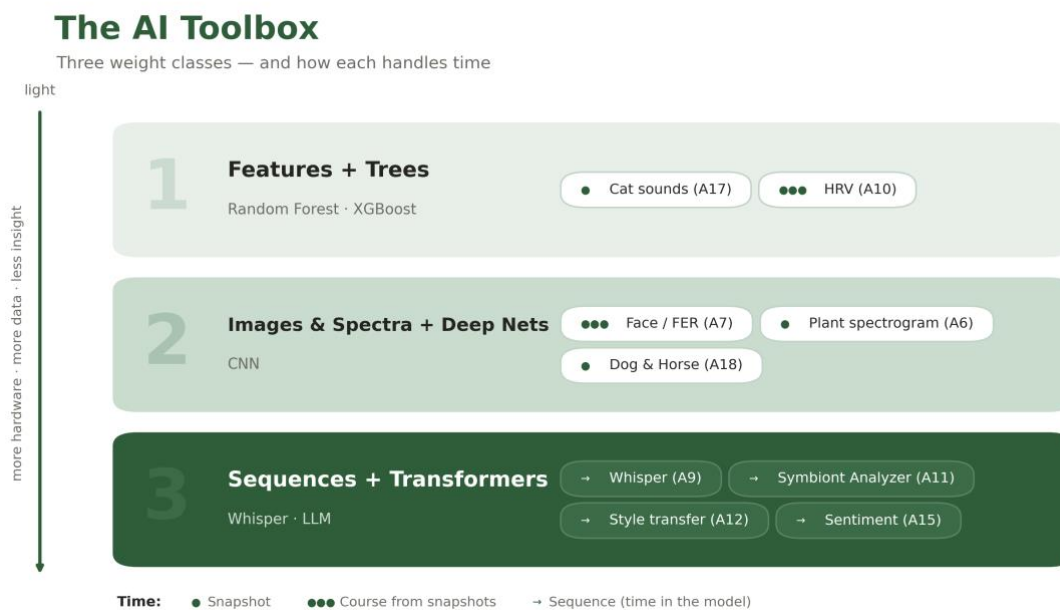


Figure 1: *The AI toolbox — three weight classes and the handling of time.*

These three tools run through the whole book. The light, transparent methods you meet above all where you choose the features yourself — with the cat sounds and the voice measurements. The middle class, the deep networks, works everywhere a signal becomes an image: with the face, with the horse, with the spectrogram of a plant. And the heavy, pretrained models appear wherever language is involved — in transcription, in the Symbiont Analyzer, in the mood of an entire stream of posts. You do not have to learn them anew each time; you recognise them again.

What this means for AI literacy

Once you have this map before your eyes, you no longer see “AI” as magic but as what it is: a toolbox in which you trade transparency and computational cost against power. That is half the battle against the fear that a machine could read minds. A model whose weight class and handling of time you can name keeps no secrets from you — you know where its answer comes from and where its limits lie.

The same machine — one image, three worlds

Before the journey through the five parts begins, an image to remember. It shows the secret protagonist of this book — not a human, an animal or a plant, but the method that makes all three readable. For as different as a tense voice, a cat’s meow and the ultrasonic click of a thirsty tomato may sound: the moment you turn them into a spectrogram — into an image of time and frequency — they look astonishingly alike, and the same image-recognition network reads them all.

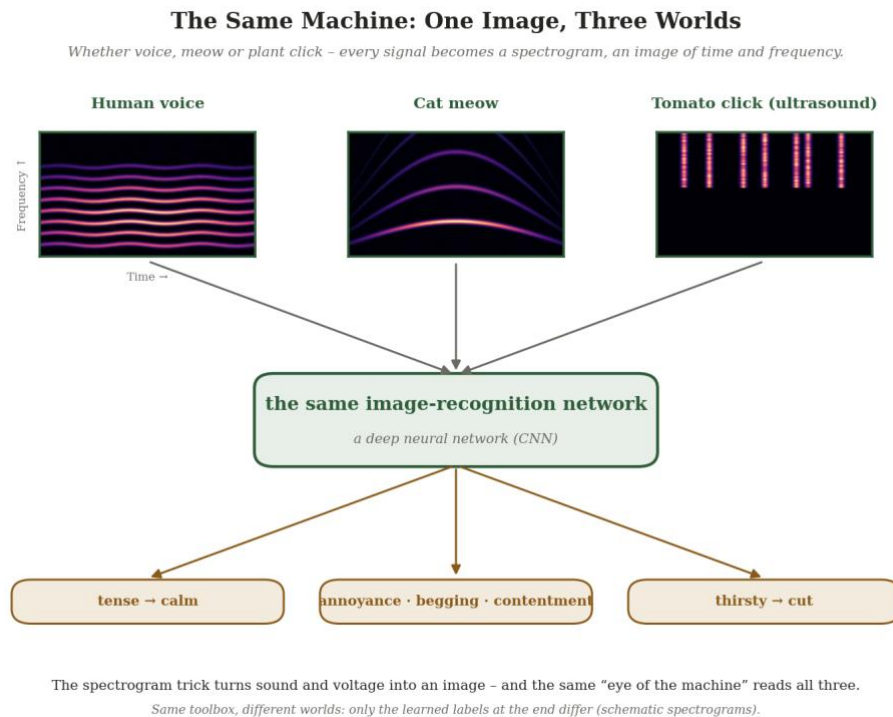


Figure 2: *The same machine — one image, three worlds. Voice, meow and plant click become a spectrogram, and the same image-recognition network (CNN) reads all three. Only the learned labels at the end differ. (Schematic spectrograms.)*

Look at the three tiles above. In the voice, horizontal bands run — the overtones that carry a voice. In the meow, a line arches up and down — the cat raises and lowers the pitch. In the tomato click, single vertical strokes flash high up: short, broadband cracks far above our hearing threshold. Three worlds, three patterns — and yet the same format. That is why a single “eye of the machine” suffices; only the labels at the end differ, because each world has its own states.

A recognisable marker. This image returns in each of the five parts — as a short line at the start that tells you which world the machine is reading there. Wherever in the book a signal becomes a spectrogram, think of this page: it is always the same machine, only before a new world. And in the last part it will read a world you have overlooked until now — yourself.

Part I

The Physician

Paracelsus · 1493–1541

“Do not read the old books — read nature herself.”

— after Paracelsus

The same machine — here it reads the human and their groups: from words, rhythm and roles.

Chapter 4 • The Bonfire of Basel

WHAT THIS CHAPTER IS ABOUT

In 1527 a combative physician burns the sacred textbooks of medicine in Basel and calls out: do not read the old books, read nature itself. This chapter recounts Paracelsus's life and his demand — and why, five hundred years later, that demand is the foundation of every measurement in this book.

On 24 June 1527, St. John's Day, a bonfire burned in Basel. Students danced around it, as they did every year at the solstice. Then a small, combative man stepped up to the flames and threw in something no one had expected: the venerable textbooks of medicine. The works of Galen and Avicenna, for a thousand years the unchallenged authority of every physician, went up in smoke before the eyes of the city. The man who burned them was the new town physician and university lecturer — and his name was Paracelsus.

He was born in 1493 as Theophrastus von Hohenheim, near Einsiedeln, a few kilometres from the place where this book is being written. His father was a physician, and the young Theophrastus learned his craft not only at universities but on the road: in the mines of Tyrol, on battlefields, among surgeons, midwives and barber-surgeons. He gave himself the name Paracelsus — “beyond Celsus,” beyond the famous Roman physician. The name itself was a declaration.

The thesis behind the fire. Paracelsus's anger was directed not at the old physicians as individuals but at a condition: for a thousand years, medicine had been copying the same books instead of questioning nature itself. Authority had taken the place of observation. His counter-thesis was simple and outrageous: the true book of the physician is nature. Read the patient, the mineral, the plant directly — trust the thing more than the tradition. He gave his lectures in German rather than Latin, so that knowledge would not remain locked behind a scholarly language. That, too, was a scandal.

Where he was mistaken. We should not idealise Paracelsus. In particulars he was often wrong. His doctrine of signatures — the idea that a plant resembling an organ heals that very organ — is superstition; his alchemy contains a great deal of wishful thinking. A reader finds, alongside bold insights, no shortage of nonsense. This is not a contradiction but typical of a pioneer: the impulse can be right long before the method is.

What remains. And the impulse was right — it has remained the foundation of every science: to trust the thing more than the authority. That is exactly what you do when, in this book, you clip an electrode to a leaf and read the plant's voltage itself instead of looking up in a book “what plants feel.” You question nature directly. Five centuries after the Basel fire, Paracelsus's demand has become a measuring instrument.

Paracelsus was a physician — the one who reads the body. This part therefore begins with the being we believe we read best: the human. The next chapters ask what our body and our groups give away without meaning to — and build from it the first reading tool of this book.

Chapter 5 · The Chemistry of Teams — honest signals in humans

WHAT THIS CHAPTER IS ABOUT

Some groups spark, others stall — and for a long time this “chemistry” was treated as an immeasurable accident. This chapter shows that it lies in a group’s honest signals: in the body (Pentland’s sociometer), in language, and in six measurable patterns of collaboration that condense in the end into five roles.

Put five people in a room and give them a task, and within minutes something arises that you can almost grasp: a mood, a rhythm, a kind of chemistry. Some groups spark, others stall. For a long time this was considered immeasurable — a “soft” factor about which one could, at best, have an intuition. This chapter shows that it is measurable, because a group’s chemistry lies in its honest signals.

The body speaks. At MIT, Sandy Pentland developed small wearable sensors that he called sociometric badges. They measured not what people said but how: how much energy lay in the voice, whether someone unconsciously mirrored the gestures of a counterpart, who set the tempo of a conversation. The first practical field trials with these sociometers I ran in Cologne, together with students of the University of Cologne, at a bank. What emerged was notable, and I was able to document it in a paper co-authored with Pentland: from a few minutes of these non-verbal signals one can read not only moods but even predict personality traits — and the outcome of negotiations or job interviews often better than from the content of the words. The body gives away what the words conceal.

Your well-being on your wrist. What Pentland’s sociometer did in the laboratory now sits in an ordinary smartwatch. With my team I built the Happimeter from it — an app that estimates, from the watch’s body signals (pulse, movement, activity, plus place, time and weather), how happy, active or stressed you currently are. At first you tell the watch a few times a day how you feel; from this it gradually learns your individual pattern. At a bank in Hamburg, 22 employees wore the watch for three months — and the notable result was not the measurement itself but its effect: the group that received daily feedback about its mood was, in the end, roughly 16 percent happier and 26 percent more active than the comparison group without feedback. Merely measuring and making one’s own state visible made people happier — we call this virtual mirroring. It is the same idea as the face mirror: whoever sees themselves honestly can change.

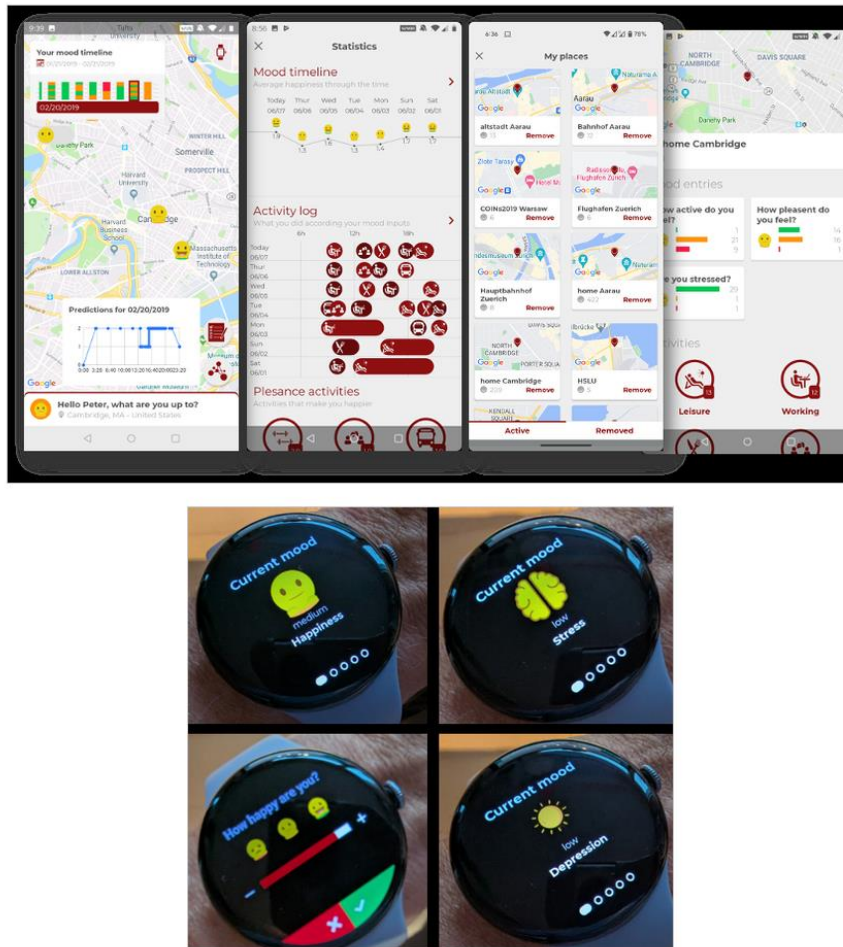


Figure 3: *The Happimeter on smartphone and watch. Top, the app (from left): a mood map showing where one felt how; the timeline of mood across the week with a log of activities; the saved locations; and the questions “How active, how content, how stressed are you?” Bottom, the smartwatch: it displays the estimated current mood – happiness, stress, low spirits – and occasionally asks directly (“How happy are you?”) in order to learn your individual pattern. (App interface shown in the original German.)*

Language gives away, too. Honest signals lie not only in the body but in language itself – not in what we talk about, but in the small words we never notice: “I,” “but,” “because.” The psychologist James Pennebaker has shown how much these function words reveal about a person. How to read this linguistic trace with your own hands is the subject of Chapter 7.

From two people to the group. In my own research over the past twenty years at the MIT Center for Collective Intelligence, I extended this idea from individuals to whole groups. When five, ten or a hundred people work together on an idea over months – in school classes, research teams, open-source projects – their shared written communication indicates whether the group will be creatively productive or not. Six signals proved especially reliable: whether there is a strong, connecting

leadership — measurable as so-called betweenness centrality, that is, how much a person connects the others; whether this leadership rotates over time rather than fixing on one person; how quickly members respond to one another; how evenly contributions are distributed; how honestly moods may be expressed; and how far the group develops its own shared language. These six signals of collaboration are taken up in detail in Part IV.

From signals to roles. For a start, these six signals condense into something more concrete: a small set of recurring roles that people take on in a group — often without noticing. Who drives with new ideas? Who executes? Who provides beauty, who provides peace? And who takes more than they give? The next chapter gives these roles five names from the book of nature — and the one after that turns them into your first reading tool.

TRY IT YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Two activities make the honest signals of this chapter directly experienceable (full instructions online: hidden-signals.swarmcreativity.com/aktivitaeten-en.html).

5.1 Heartbeat and stress. With a heart sensor it becomes visible how your heart-rate variability falls under stress and rises at rest — an honest signal in pure form, one you cannot control at will.

5.2 The six honest signals. From the timestamps of a real chat you compute response times, contribution balance and the give-and-take ratio. An intuition about a group becomes a measurable quantity — the precursor to the tool of the next chapter.

Bonus: measure your well-being — the Happimeter (free). If you have an Apple Watch or Wear OS watch, download the Happimeter for free (happimeter.socialcompass.ai) and test for a few days how well a machine can estimate your mood from pulse and movement.

Chapter 6 • The Book of Nature — the five symbionts and their roles

WHAT THIS CHAPTER IS ABOUT

For the roles that people take on in a group, nature has long held ready models. This chapter presents five of them — Bee, Ant, Butterfly, Capybara and Leech — and insists on one sentence that carries everything: these are roles, not labels for people.

“The book of nature” was Paracelsus’s name for what the physician should read. Here we open it in the literal sense. For the roles that people take on in a group, the living world has long held ready models. In nature, the coexistence of different species is called symbiosis — a living-together in which each contributes to the whole in its own way, or serves itself in its own way. Five such roles recur in almost every group. We borrow their names from five living things.

The Bee. It carries back and forth between the blossoms and pollinates in passing. In the group it is the creative one: it throws in new ideas, proposes unknown paths, asks “what if we did it completely differently?” Without the Bee, a group stands still.

The Ant. It builds, hauls, keeps the structure running. In the group it is the executor: it structures, plans, makes lists, fixes deadlines, brings ideas to realisation. Without the Ant, every idea remains only an idea.

The Butterfly. Here a note is worth making, because the common image misleads: the Butterfly of these five roles is not the “sociable networker” it is often taken to be. It stands for beauty. In the group it is the designer: the person who ensures that the result is beautiful — who designs the layout, chooses images, creates a logo, attends to mood and aesthetics. Without the Butterfly much still functions, but nothing moves you.

The Capybara. The capybara is perhaps the most peaceable animal in the world; in photographs it is seen resting while birds, monkeys and even crocodiles feel at ease beside it. In the group it is the harmoniser: it settles disputes, mediates between camps, listens, says “no stress” when others come under pressure. Without the Capybara a group falls apart on its own tensions.

The Leech. The leech sucks — it takes without giving. This is the most unpleasant of the five roles, and we have deliberately kept its name rather than soften it. In the group the Leech is the one who draws more value from the community than they contribute: who picks up others’ ideas, reshapes them slightly and passes them off as their own — and pushes to the front when, at the end, the question is who did what. Every group knows this role. Naming it honestly helps more than overlooking it.

Roles, not labels. And now the most important sentence of this chapter: these are roles, not labels for people. No one “is a Bee.” The same person can be the driving Bee in one group and, tired and overwhelmed, become the Leech in another. The role

belongs not to the person but to the relationship — it can change from group to group and from week to week. That is why the tool you build in the next chapter never measures “who someone is,” but only which role they played in exactly this chat, at exactly this time. Whoever forgets this and stamps a classmate as a “Leech” for good has read the book of nature wrongly.

In the next chapter we leave description behind and become practical. We take a real group chat — the planning of a class trip — and build, step by step with the help of a language model, a program that recognises these five roles from the words alone. It is the first complete reading tool of this book.

Chapter 7 • Reading Honest Signals — the Symbiont tool

What our word choice gives away — and how a language model shows it to us

WHAT THIS CHAPTER IS ABOUT

When five people plan a class trip together, each writes differently. One keeps throwing new destinations into the chat. Another writes lists of flight times. Another posts attractive photos of the hostel roof terrace. A fourth calms things when the discussion escalates. A fifth reads along and pushes to the front at the end. In this chapter we build a program that recognises exactly these five roles from a WhatsApp chat — by counting particular words. And we have the program written for us, step by step, by a language model.

What Pennebaker discovered in the word “I”

In the 1990s the American psychologist James Pennebaker had an unusual idea. He asked himself: do people give away, in their word choice, things about themselves that they do not notice? Pennebaker and his team analysed thousands of texts — diaries, emails, essays — with a simple computer program that merely counted small words: “I,” “you,” “but,” “because,” articles, pronouns. Words, that is, that no one consciously thinks about.

The result was robust enough that it now belongs to the standard of language psychology: people who suffer from depression use the word “I” markedly more often than non-depressed people. They may themselves say “I’m fine” — but their word count indicates something else. Such unconscious language patterns are the honest signals of language — behavioural traces we do not consciously control, and which therefore give away more than what we deliberately say.

This discovery is the foundation for everything that follows in this chapter. It states: when we count words instead of interpreting content, we can detect personality patterns of which the writer is not even aware. This sounds like magic, but at its core it is statistics.

Six signals of creative collaboration

In my own research over the past twenty years at MIT, I extended Pennebaker’s method to groups. When five, ten or a hundred people work together on an idea over months — in school classes, research teams, open-source projects — their shared written communication yields six robust signals that predict whether the group will be creatively productive or not. These six signals were introduced in Chapter 5 and are described in detail in Part IV; here a brief reminder suffices: they concern how leadership is distributed in the group, how quickly members respond to one another,

how honestly moods may be expressed, and how far the group develops its own language.

In this chapter we work not with the six signals directly but with their condensation into five symbiotic roles: Bee (the creative one), Ant (the executor), Butterfly (the designer), Capybara (the harmoniser) and Leech (the taker). You met them in detail in Chapter 6 – where the important sentence also stands, that these are roles, not labels for people. For this chapter a reminder of the Leech suffices, because it makes the analysis tricky in the end: the Leech contributes no ideas of their own but is not silent – they pick up others’ ideas, reshape them slightly and present them as their own. When a teacher asks at the end who did what, the Leech pushes to the front, ahead of exactly those whose work they took over.

The thesis of this chapter: these five roles can be recognised from the words used – in the WhatsApp chat of a class-trip planning as much as in the emails of a research team.

Our example: the class trip to Lisbon

Imagine an eleventh-grade class planning its study trip. Five students take over the preparation and create a WhatsApp group for it: Mira, Tom, Lara, Jonas and Felix. They have known one another for years. Over the next two weeks they write 33 messages. Here is the complete exchange:

```
# chat.txt – 33 messages from the class group
Mon 14:02 - Mira: I'd really love Lisbon. What do you think?
Mon 14:03 - Mira: Idea: we could also come back via Porto, it's supposed to be
beautiful.
Mon 14:05 - Tom: I checked the flights. Lisbon return from Düsseldorf about 180 euros in
May. Here's the list with three providers.
Mon 14:11 - Lara: Oh yes, Lisbon! There's this viewing terrace over the rooftops, I
could use it for our programme poster 😊
Mon 14:18 - Jonas: Both suggestions have something. Let's collect them calmly before we
decide.
Mon 18:42 - Felix: Idea: why not Lisbon AND Porto? Take both in on one trip, that would
really be something.
Tue 08:51 - Mira: Yes Felix, exactly as I said yesterday 😊 Porto would be great. Tom,
could you look at train connections Lisbon-Porto?
Tue 09:14 - Mira: What do you think about building in a day of surfing? Cascais is
supposed to have perfect beginner waves.
Tue 09:30 - Tom: Hostel research is done. Three options, all under 30 euros a night,
location checked. Table in the Drive.
Tue 09:42 - Lara: I'll take care of the programme booklet. Look, I've already designed a
logo – does it fit Lisbon?
Tue 10:01 - Jonas: Lara, well done. Tom, thanks for the table. Mira, your surf day
sounds good, does it cost extra?
Tue 14:20 - Mira: Imagine we also do a Fado concert in the evening. It would be a
completely different experience from just the beach.
Tue 14:35 - Tom: Concretely: surf day in Cascais costs 45 euros per person, Fado concert
from 18 euros. I'll factor it into the budget.
Tue 22:11 - Felix: A watersports day on the Atlantic coast would be a highlight for me.
Something like that just has to be part of Lisbon.
Wed 08:50 - Lara: Last night I made a mood image for the Insta post – orange sunset,
white houses. Take a look!
Wed 09:02 - Mira: Lara, the image is gorgeous! Idea: we could start the Insta account
before the trip, as anticipation.
```

Wed 09:15 - **Tom**: Registration runs until Friday. Who's definitely in now? Please just "yes" or "no" here.

Wed 09:20 - **Jonas**: Yes from me.

Wed 09:21 - **Lara**: Yes of course! And I'll gladly design the registration email – shall I propose something?

Wed 09:25 - **Mira**: Yes! And a new idea: everyone brings a song for our shared Spotify playlist.

Wed 13:55 - **Felix**: By the way, I presented our Lisbon planning to my parents – programme booklet, Insta concept, everything. They were thrilled.

Thu 10:00 - **Tom**: Status: 5 registrations in. Felix, please look over the hostel table once more if you want to contribute something, otherwise it goes out tomorrow morning.

Thu 10:14 - **Jonas**: Lara, the programme booklet looks great, excellent work. Tom, thanks for booking the hostel.

Thu 11:30 - **Lara**: I've finished designing the programme booklet. Look – does the orange gradient fit the Lisbon vibe?

Thu 14:08 - **Mira**: Idea for the last evening: picnic by the Tejo with pastéis de nata? A nice finish.

Thu 14:45 - **Tom**: Picnic: ok, costs about 8 euros per person for snacks. I'll handle the shopping on site.

Fri 08:30 - **Felix**: Ms Berger just spoke to me, I explained our trip planning to her in broad strokes. She was really impressed by what we pulled off.

Fri 08:45 - **Tom**: 5 of 5 confirmed. Registration is being sent.

Fri 09:00 - **Jonas**: Great that it worked out. Well done, everyone.

Fri 12:11 - **Lara**: Insta account is online: @11b_lisbon. First image posted – take a look!

Fri 14:00 - **Mira**: What if we make a little travel-diary app? Everyone writes three sentences a day.

Fri 14:20 - **Lara**: I like the idea, I'll design the layout for it.

Fri 14:35 - **Jonas**: Sounds good. Let's not make it too big, no stress, three sentences are enough.

Read the chat through calmly before you read on. As you do, try once to assign each of the five people a symbiont role intuitively — that is: who is Bee, who Ant, who Butterfly, who Capybara, who Leech? Note down your guess.

Which words give away which role?

If you read the chat attentively, you probably noticed: the five people write quite differently. Mira begins almost every message with something new: "Idea:," "imagine," "what do you think." Tom uses factual, planning words: "concretely," "I'll handle," "done." Lara works with aesthetic words: "beautiful," "design," "look," "fits." Jonas speaks reassuringly: "calmly," "let's," "I understand." With Felix it is harder. At first glance he sounds like a mixture of Mira and Lara. But if you look more closely, his "ideas" always come shortly after those of others — and most conspicuous is how often he mentions outsiders: "my parents," "her" (the teacher), "impressed by." He positions himself toward the outside.

This observation is the basis of our analysis. For each symbiont role we define a small list of typical words and phrases — a so-called word list. Then we count how often each person uses words from which list. The person with the most Bee words is (probably) the Bee; the one with the most Leech words is (probably) the Leech. "Probably," because this procedure is not perfect — we discuss its limits at the end of the chapter.

Role	Typical words and phrases
------	---------------------------

Role	Typical words and phrases
Bee	<i>idea, new, what do you think, we could, imagine, what if</i>
Ant	<i>done, finished, I'll handle, here's the list, concretely, status</i>
Butterfly	<i>beautiful, look, design, fits, mood, layout, logo</i>
Capbara	<i>calmly, let's, no stress, thanks, sounds good, everyone, I understand, mediate</i>
Leech	<i>my parents, told her, impressed by, what we pulled off, a highlight for me</i>

These lists are not a law of nature. They are our hypothesis about which words belong to which role. Another researcher could propose different lists — and would then have to justify why. This is precisely the scientific core of the analysis: we make our assumptions explicit, write them down, and can criticise them.

Prompt workshop: have the program written for you

In this chapter we do not program ourselves — we have the program written for us, step by step, by a language model. For this we use Claude (free at claude.ai), but you can just as well use another language model. What matters is only that you understand what the language model is doing for you and why the generated code works.

We build the program in four small steps. At each step you formulate a request (a so-called prompt), the language model responds with code, and we examine the code together. You copy the code into a file with the ending `.py` and run it with Python.

Step 1 — Read the chat as a file

YOU TYPE

Hello Claude, I have a WhatsApp chat export as a text file `chat.txt` in the same folder as my Python script. Each line looks like this: “Mon 14:02 - Mira: Idea for Lisbon.” Please write me a Python program that reads the file and outputs all messages as a list of three values: time, sender, text.

The language model responds with a short program. It looks like this (with explanations we will go through in a moment):

```
# Program: read_chat.py
messages = []

with open('chat.txt', encoding='utf-8') as file:
    for line in file:
        if ' - ' not in line or ': ' not in line:
            continue
        time_and_rest = line.split(' - ', 1)
        time = time_and_rest[0]
        sender_and_text = time_and_rest[1].split(': ', 1)
        sender = sender_and_text[0]
        text = sender_and_text[1].strip()
        messages.append((time, sender, text))
```

```
print(f'Found {len(messages)} messages in total.')
for time, sender, text in messages[:3]:
    print(f'{time} | {sender}: {text}')
```

What happens here? Let us go through the most important lines:

- `messages = []` — we create an empty list in which we will later collect all messages.
- `with open(...) as file:` — we open the chat file for reading; the `with` construct ensures the file is closed automatically at the end.
- `for line in file:` — we go through each line of the file individually.
- `if ' - ' not in line ... continue` — if a line does not have the expected format (for instance a blank line), we skip to the next.
- `split(' - ', 1)` — we cut the line at the first " - " into `time` and `rest`, then the `rest` at ":" into `sender` and `text`.
- `messages.append((time, sender, text))` — we append the three values as a tuple to our list.

When you run this program, it should print roughly:

```
Found 33 messages in total.
Mon 14:02 | Mira: I'd really love Lisbon...
Mon 14:03 | Mira: Idea: we could also come back via Porto...
Mon 14:05 | Tom: I checked the flights...
```

If that works, you have completed the first step: the program now understands the structure of a WhatsApp chat.

Step 2 — Write the word lists into the program

YOU TYPE

Very good. Now extend the program: define five word lists as a Python dictionary (dict), as I will send them to you. The keys are the role names, the values are lists of words and phrases. Here are the lists: ...

You insert the lists from the table above into your prompt. The language model extends the program with roughly this block:

```
wordlists = {
    'Bee': ['idea', 'new', 'what do you think',
           'we could', 'imagine', 'what if'],
    'Ant': ['done', 'finished', "i'll handle",
           "here's the list", 'concretely', 'status'],
    'Butterfly': ['beautiful', 'look', 'design',
                 'fits', 'mood', 'layout', 'logo'],
    'Capybara': ['calmly', "let's", 'no stress', 'thanks',
                'sounds good', 'everyone', 'i understand',
                'mediate'],
    'Leech': ['my parents', 'told her',
             'impressed by', 'what we pulled off',
             'a highlight for me'],
}
```

A dictionary (dict) is a Python data structure that stores a value for each key. Here the keys are the role names and the values are the word lists. With `wordlists['Bee']` you get back the Bee words, for example.

Step 3 — Count the words per person

YOU TYPE

Now the main part: write a function that counts, for each person in the chat, how often words from each word list occur. The result should be a nested dict: one dict per person, containing a counter per role.

```
def count_words(messages, wordlists):
    hits = {} # empty dictionary for the results
    for time, sender, text in messages:
        # if this person is not yet in the dictionary, add them
        if sender not in hits:
            hits[sender] = {role: 0 for role in wordlists}
        # lowercase the text so "Idea" and "idea" count the same
        text_lower = text.lower()
        for role, words in wordlists.items():
            for word in words:
                if word in text_lower:
                    hits[sender][role] += 1
    return hits

result = count_words(messages, wordlists)
for person, counts in result.items():
    print(f'\n{person}:')
    for role, count in counts.items():
        print(f' {role:14s} {count} hits')
```

This function is the heart of the program. It goes through each message, looks at each word for each of the five roles, and checks with `if word in text_lower` whether the word appears anywhere in the text. If so, the counter for this person and this role is increased by 1. The result for our chat looks roughly like this:

```
Mira:
  Bee          9 hits
  Ant          0 hits
  Butterfly    3 hits
  Capybara    0 hits
  Leech       0 hits
Tom:
  Bee          0 hits
  Ant          5 hits
  Butterfly    0 hits
  Capybara    0 hits
  Leech       0 hits
Lara:
  Bee          1 hits
  Ant          2 hits
  Butterfly    8 hits
  Capybara    0 hits
  Leech       0 hits
Jonas:
  Bee          0 hits
  Ant          0 hits
  Butterfly    0 hits
```

Capybara	9 hits
Leech	0 hits
Felix:	
Bee	1 hits
Ant	0 hits
Butterfly	0 hits
Capybara	0 hits
Leech	4 hits

Four of the five give a clear picture: Mira dominates on the Bee (9 hits), Tom on the Ant (5), Lara on the Butterfly (8), Jonas on the Capybara (9). Small secondary hits appear almost everywhere — Mira uses an aesthetic word here and there, Lara an Ant word — but the emphasis is clear in each case. With Felix it is more interesting: he has four Leech hits, but also one Bee hit. If you look closely, you see where it comes from: Felix begins his first message with “Idea:” — he adopts Mira’s style but presents a suggestion that Mira had made five hours earlier. This is exactly Leech behaviour: he sounds like a Bee but is not one. This ambiguity is a genuine weakness of our procedure, to which we return at the end of the chapter.

Step 4 — A small visualisation

YOU TYPE

Last step: show the result as a bar chart. One row per person, five bars side by side for the five roles, in different colours. Use matplotlib.

The language model delivers a script using the matplotlib library (which you install once with `pip install matplotlib`). Here is the finished bar chart for our Lisbon chat:

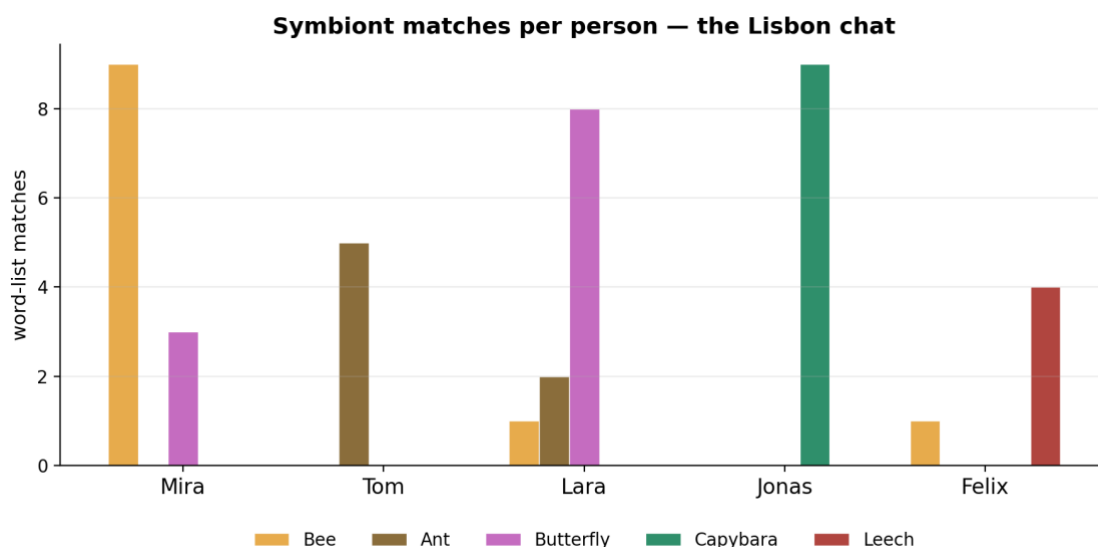


Figure 4: *Symbiont matches per person in the Lisbon chat. Each person has a clear emphasis in exactly one role; small secondary matches (for instance Mira on the Butterfly, Lara on the Ant) are normal. Only Felix is conspicuously spread — four Leech matches and one Bee match from the adopted contribution.*

At a glance you see: Mira is the Bee, Tom the Ant, Lara the Butterfly, Jonas the Capybara. Felix, by contrast, stands out on the Leech — and the small Bee bar recalls how he passed off Mira’s idea as his own. You can insert this image into your homework or present it as a basis for discussion in class.

WHAT YOU HAVE LEARNED ABOUT PROMPTING

Proceed in small steps. A long prompt such as “write me the whole program” rarely works well. Four small steps are better than one large one.

Describe the data format precisely. “Each line looks like this: ...” helps the language model more than general descriptions.

Say what you want to see at the end. “Result as a bar chart” is clearer than “visualise it nicely.”

Stay in control. If you do not understand a line of code, have the language model explain it. It is your tool, not your teacher.

Limits of our method — where caution is called for

We have just built a program that recognises five personality roles from a chat. The result looks impressive. But what have we actually measured?

First limit: our example chat is constructed. We wrote the chat ourselves — and in such a way that each person has a clear role. In a real WhatsApp chat the boundaries are far blurrier. A person can be Bee and Butterfly at the same time, or switch between roles depending on the day. That our method works so clearly with this chat does not prove that it works so clearly with every chat.

Second limit: our word lists are a hypothesis. We claimed that “idea” and “imagine” belong to the Bee. But someone could use “idea” without being creative, and someone could be very creative without ever using the word. Word lists are coarse tools. Better methods — for instance those based on language models such as BERT — understand meaning in context rather than merely counting words. Such methods sit in the AI toolbox (Chapter 3); the Master extension at the end of this chapter shows one of them in action.

Third limit: word counting does not detect everything. You saw it with Felix: his Leech behaviour — adopting others’ ideas and selling them outward as his own — is only partly captured by our word list. We get a hint (the self-positioning words “my parents,” “told her,” “impressed by”), but miss the actual mechanism: that Felix’s “ideas” are, in content, repetitions of others’ ideas. To detect this, we would have to measure semantic similarity over time. This is exactly what modern language models do — the Master extension takes it up.

Fourth limit — the most important: we have judged other people without asking them. This is a serious point. Before you apply this program to a real chat, consider two things. First, do the others consent? If not, the analysis is not science but surveillance. Second, what do you do with the result? If your program classifies

someone as a “Leech” — do you tell her? Do assumptions about others keep you from getting to know them for real?

These four limits are no reason to discard the method. They are a reason to describe it honestly. This is exactly what distinguishes AI literacy from blind trust in AI.

GOING DEEPER · CODE, INTERACTIVE AND FURTHER READING

The complete, runnable code examples of this chapter are in the GitHub repository github.com/pgloor/verborgene-signale-code. Alongside it there is an interactive web version of this chapter in which you can run the code steps directly in the browser and try them with your own chats: hidden-signals.swarmcreativity.com/kapitel7-en.html.

Anyone wishing to understand the framework of the six honest signals more precisely will find in the 12-minute talk “Six Honest Signals of Creative Communication” (YouTube) a step-by-step explanation using real research data. Further practical chapters — including some with your own hardware — follow in Part III.

Discussion questions for the close of the chapter

- Pennebaker found that depressed people say “I” more often. Now that you have this knowledge — would you behave differently when you write something? Does the research thereby change the very thing it set out to measure?
- You have analysed a chat. If the five people learned that you had sorted them into five boxes — how would they react? What would you say to them?
- Felix repeatedly adopts others’ ideas and presents them as his own. A direct confrontation never occurs in the chat. Why, do you think, do groups often avoid open engagement with such behaviour? What long-term consequences follow when no one says anything?
- Imagine a company analyses its employees’ emails with a similar program. Would that be acceptable? What conditions would have to be met? Research what the General Data Protection Regulation (GDPR) says about it.
- In the next part of the book we will see that a plant, too, sends out “honest signals” — electrical voltages that give away its state. What property makes a signal an honest signal? Is it bound to consciousness, or does a measuring apparatus suffice?

MASTER EXTENSION · From word counting to a language model — the Symbiont classifier

This section is aimed at advanced readers who have programming experience and want to push the method from the base part further. In the Master course at MIT, HSLU and the University of Cologne, in which this material is tested, this step is the central learning occasion. The word-list approach has two weaknesses. It does not

recognise a stylistic Butterfly who never literally says “beautiful” but writes aesthetically. And above all, it recognises the Leech only unreliably — for the Leech’s actual signature is not a word but the repetition of an idea that someone else voiced earlier. The extension pursues a two-stage approach, known in language processing as weak supervision:

Step 1 — Generate a training dataset with the simple method. We collect 100 to 500 anonymised chats (for instance from open-source forums or the COINs research corpus), apply the word-list classifier, and automatically obtain pseudo-labels. This is not perfect, but good enough as a teacher model.

Step 2 — Fine-tune a BERT model. With the pseudo-labels we train a pretrained model (such as `gbert-base` for German, or a comparable English model) on the symbiont classification. The Hugging Face library `transformers` makes this possible in about 30 lines of Python; training time on a standard GPU: 1–2 hours.

Step 3 — Catch the Leech with semantic similarity. In addition, for each message we compute the semantic similarity to all previous messages by other people (with Sentence-BERT). Whoever regularly produces contributions with high similarity to earlier contributions by others, without citing them, is a Leech candidate. In Felix’s case, this procedure would clearly flag the Mon-18:42 message as an adoption of Mira’s Mon-14:03 idea.

Step 4 — Reflection on the method. The BERT-based procedure, too, has limits: someone who has a thought in parallel and independently looks, to the model, like a Leech. Can we distinguish, from behavioural data alone, between adoption and independent convergence? The literature on plagiarism detection gives no unambiguous answer.

The full guide with code, training data and a commented Jupyter notebook session is available in the GitHub repository github.com/pgloor/verborgene-signale-code (folder `kap07-master`).

TRY IT YOURSELF — THE SYMBIONT TOOL

The do-it-yourself tool for this chapter — from the finished app to your own language model (interactive version and instructions online: hidden-signals.swarmcreativity.com/kapitel7-en.html).

7.1 Use the Symbiont Analyzer. You apply the finished web edition to an example chat and read the five roles from the words alone.

7.2 The prompt workshop. Instead of programming yourself, you have the analyzer written for you in four small steps by a language model — and understand every snippet of code you receive.

7.3 Symbiont classifier (Master). For advanced readers: from word counting to a language model — fine-tune a BERT model and catch the Leech with semantic similarity.

Part II

The Invisible College

The Rosicrucians · Fama Fraternitatis 1614

“Share your knowledge freely — across all borders and walls.”

— after the Fama Fraternitatis

The same machine — here it reads face and voice: from image and spectrogram.

Chapter 8 • The Invisible College — from brotherhood to open source

WHAT THIS CHAPTER IS ABOUT

In 1614 a pamphlet announced a secret brotherhood that shared its knowledge through an invisible network. The brotherhood was invented — but the idea came true. This chapter follows it to the present day: to open source, open science, and the tools you work with in this book. And it shows that such a college can itself be read — as a network.

In 1614 a slim pamphlet with a large promise appeared in Kassel: the *Fama Fraternitatis*. It told of a secret brotherhood, the Rosicrucians, said to be scattered across Europe, healing the sick for free and sharing their knowledge through an invisible network of minds. Scholars across half of Europe searched for this brotherhood, wrote letters, begged to be admitted. In vain — for it did not exist. The *Fama* was a fiction.

A fiction that came true. But it was a fiction that had effect. The idea of an open college, in which knowledge is not hoarded but shared, inspired a whole generation. A few decades later a circle of natural philosophers formed in London that put exactly this idea into practice — the Royal Society. One of its pioneers, Robert Boyle, spoke explicitly of an “invisible college”: a network of researchers who exchange their results openly instead of keeping them secret. An invention had become an operating principle.

The invisible colleges of today. This principle is more powerful today than ever. Wikipedia is an invisible college: hundreds of thousands who have never met write a treasury of knowledge together. Open science, which makes its data and papers freely available, is one. And most visibly: open source — software whose source code lies open and which anyone may help build. The operating system that quietly drives half the world, the programming language of this book (Python), the tools with which we are about to read faces and voices — they are all the work of such colleges.

Who builds the tools of this book? This is not a side note but the reason this book is possible at all. Twenty years ago, reading an emotion from a face would have required a research lab and expensive specialist software. Today a few lines of Python and a library that an open community built and gave away are enough. The language model that wrote your program in Chapter 7, the model that turns speech into text in Chapter 10 — they stand on the shoulders of thousands who were mostly never paid for one another. The Rosicrucians dreamed of a network of free minds; you use it every day.

A college can be read. And here a circle closes back to Part I. An invisible college is a network — and networks can be measured. Who writes to whom, who connects whom, where does the strong, connecting leadership sit (the betweenness centrality from

Chapter 5)? Make that visible, and you can see in a group whether it is creatively flourishing or falling apart. I call such groups Collaborative Innovation Networks — networks in which people create something new together of their own accord. They are the living descendants of the invented brotherhood of 1614.

My first college: the Deloitte email archive

I came upon this idea at my own inbox. Until the end of 2002 I was a partner at Deloitte, leading e-business for Europe — and I suspected that the real collaboration in my team ran quite differently from what the org chart claimed. So I did something unusual for the time: I took my own email archive and that of my working group, the “e-Champions”, and drew a social network from it — who wrote to whom, how often, how quickly the reply came. I did not read the content of the mails, but their pattern.

The result changed my view of organisations forever. The most important nodes in the network — the people through whom knowledge truly flowed — were often not those with the highest title. Some official bosses sat at the edge; some inconspicuous staff were the secret hubs without whom nothing worked. The org chart shows who gives whom orders; the communication network shows where innovation arises. Out of this observation came my most-cited book, *Swarm Creativity* (2006), and the beginning of twenty years of research on Collaborative Innovation Networks.

The great open colleges

The same shape is found wherever people work together voluntarily and across hierarchies. The World Wide Web arose this way: Tim Berners-Lee proposed it, and an open network of volunteers built it out. The technical rules of the internet are still made by the IETF, a loose college without a boss, whose motto is: “rough consensus and running code”. And perhaps the purest example is Linux — an operating system that a Finnish student named Linus Torvalds set in motion in 1991 and which today, built by tens of thousands of volunteers around the globe, runs on most of the world’s servers and in every Android phone. No one ordered it. It arose in the swarm.

The shape of a swarm

When you draw such networks, they have a characteristic form: most members are connected to only a few others, but a few — the hubs — are linked to very many. Researchers call this a scale-free network. These hubs hold the swarm together; they are, to stay with the image from Chapter 6, something like the queen bee in the hive: not the commander, but the centre around which everything orders itself. At Linux, Linus Torvalds sits at exactly this spot — and with the betweenness centrality from Chapter 5 you cannot merely claim this, but calculate it.

The most important consequence of these open colleges for our subject is simple: tools that were once locked away behind laboratories now lie ready for everyone. That is exactly what the next two chapters make visible — they turn two entirely everyday

devices, a webcam and a microphone, into reading devices for the honest signals of the human being.

TO DO YOURSELF — THE ACTIVITY FOR THIS CHAPTER

Read a real invisible college (instructions online: hidden-signals.swarmcreativity.com/akt-08-1-en.html).

8.1 Who talks to whom? — finding the queen bee. You fetch the public contribution data of an open-source community — such as the Linux kernel — and draw its network: who works with whom, who sits at the centre, and who is the secret queen bee (networkx).

Chapter 9 • The Face as a Reading Device

WHAT THIS CHAPTER IS ABOUT

No part of the body gives away as much as the face — and no language in the world is without it. This chapter shows how an ordinary webcam becomes a reading device with the help of open software: how a model computes a mood from the tiny movements of a face, what it reliably recognises — and what it cannot, in principle, see.

Even an infant who does not yet understand a word reads faces. A smile, a drawn-together brow, a widened gaze — we decipher them before we think. The face is the body's most honest page: it often moves faster than we can control it, and gives away in fractions of a second what we might have wanted to hide. It is exactly this readability that makes it the most obvious target when a machine is to learn to read honest signals.

How the machine reads a face. The path is exactly the one from Chapter 3. A camera delivers an image — and images are the showpiece discipline of AI. A model first looks for fixed points in the face: the corners of the mouth, the eyebrows, the eyelids. From their position relative to one another it calculates probabilities for emotional expressions: so many per cent joy, so much surprise, so much anger. It was trained on hundreds of thousands of face photos that people had previously labelled by hand with emotion tags — again a work of open datasets and communities, as the previous chapter described.

My own trail. I worked with exactly this technique for years. In projects such as the Happimeter and Perceptiface we measured emotional expressions from faces in real time — not to surveil people, but to understand how moods arise and travel in teams. What astonishes again and again: how much a brief expression gives away — and how easily one is deceived if one credits it with too much.

Your face gives away who you are

And it goes further than you might think. The face gives away not only how you feel right now but — across many reactions — who you are. In one of our studies, titled “Your Face Mirrors Your Deepest Beliefs”, we showed 85 people a series of fifteen short videos of very different kinds: funny, touching, disgusting, provocative. While they watched, a camera read their involuntary reactions second by second via facial recognition — not a word was spoken, only the face worked.

From this pattern of reactions a model could predict, with an accuracy of up to 86 per cent, personality traits and even moral attitudes — things people otherwise report about themselves in long questionnaires. The astonishing part: no single video suffices. Only the mixture gives you away — how you react to the touching and at the same time to the provocative. And often this measurement is more honest than the questionnaire,

for in self-judgement we like to flatter ourselves; the face does not. This is the honest signature in its purest form — and precisely for that reason unsettling too.

What it does not see. For the face has limits as a reading device, and a responsible use means knowing them. First, the model measures an expression, not a feeling: a raised corner of the mouth is a smile, not necessarily joy — people smile out of politeness, embarrassment, even pain. Second, not all cultures smile alike; a model trained mainly on Western faces reads others less well. And third, in much training data the diversity of people was poorly represented — light faces overrepresented, dark ones underrepresented — which is why such systems are markedly less accurate for some people. This is not a marginal problem: it decides whom a tool serves and whom it harms.

Whose face may you read? This makes the face the most delicate of the reading devices. It is read from a distance, without touch, often without the person read even noticing. The same question as in Chapter 7 arises here even more sharply: who consented? A webcam that mirrors your own mood is a toy of self-knowledge. The same camera, aimed at others without their knowledge, is surveillance. The tool is the same; the difference lies solely in consent. And because a face, as we have seen, can give away not only mood but personality and moral stance, this consent becomes the key question: no one should be allowed to draw a judgement about your character from your involuntary reactions that you never agreed to.

The golden rule. Does this mean one should never do such a thing? No. In twenty years of measuring networks and personalities in real companies, there was one rule I never broke and which in the end convinced every workforce: aggregated results for management, personal results only for the person themselves. The company learns how the team as a whole is doing — where collaboration is stuck, where the mood is tipping. Which individual employee holds which values, only she herself learns, no one else. I told every participant this at the outset and assured them of it personally, backed by the ethics committee of MIT, which watches over exactly such studies. This one promise turned a suspicion of surveillance into a tool of self-knowledge that people trusted. It holds for every reading device in this book: what you find out about others belongs to them, not to you.

Incidentally, this reading device is not confined to humans. The same method reads the posture of a dog or the expression of a horse — that is the subject of Part III. For now, though, we stay with the human being and turn from their face to their voice.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

The face as a reading device — from the simple demo to the delicate limit (instructions online: hidden-signals.swarmcreativity.com/aktivitaeten-en.html).

9.1 The face mirror. A program opens your webcam and shows, second by second, which feelings it reads from your face. You experience how pixels become probabilities — and feel the unease in it.

9.2 Mood over an hour. You record your emotional course over a longer time and plot it — and see how mood moves across the day.

9.3 Personality from expression (Master). For advanced readers and deliberately as reflection: retracing the study “You are what you post”, which even predicts value attitudes from many reactions — and reflecting on the limits of this delicate reading device.

Chapter 10 • The Voice as a Signature

WHAT THIS CHAPTER IS ABOUT

We hear above all what someone says. A machine also hears how: pitch, tempo, tremor, pauses. This chapter turns a microphone into a reading device — it turns speech into text and reads in the voice the honest signals beneath the words. And it poses the uncomfortable question of what happens when this signature can be forged.

Say the sentence “I’m fine” in three ways: cheerful, tired, angry. The words stay the same — and yet everyone immediately understands three different things. What carries the difference is not the words but the sound: the pitch, the tempo, the small breaks and pauses. In this sound lies a second message, often more honest than the first. The voice is something like a signature: so individual that you recognise people on the phone after three words — and so revealing that it betrays what the words conceal.

From sound to image — and to text. Here too the path leads through the trick from Chapter 3: a tone becomes a spectrogram, the image of its frequencies, and thus readable for the same pattern-recognising machinery. In this way an open model like Whisper reliably turns spoken language into text — a tool unthinkable a few years ago and freely available today, again a gift of the invisible colleges. From the same spectrogram, features such as pitch, loudness and speaking tempo can be read out — and from them, with caution, hints of tension or calm.

As unique as a fingerprint. The voice carries two things at once: what you say and who you are. Its fine, unmistakable idiosyncrasies make it possible to recognise speakers — the basis of your phone listening for your voice. That is useful and uncanny at once. And it recalls Chapter 1: just as many animals speak in frequencies our ear cannot reach, the human voice too holds more than we consciously hear — a machine brings it out. It is precisely this uniqueness that can be turned into a key: a model that has learned the voices of a group recognises from the mere sound who is speaking — more on that at the chapter’s end.

When the signature is forged. Here, though, the guiding idea of this book begins to waver. Honest signals were held to be honest because they cannot be forged. With the voice this is no longer quite true: from a few seconds of recording, today’s models produce a deceptively real copy of your voice, saying sentences you never spoke. The signature can be forged. This makes the voice the most striking example of the honest limit: a signal is only as honest as it is hard to forge — and technology can shift that threshold. Whoever reads voices must from now on also ask whether a voice is genuine at all. The voice key from the end of this chapter is affected too: a good copy could open it — one reason no serious safeguard should rest on the voice alone.

With this, Part II ends. We have made reading devices out of everyday apparatus — the webcam for the face, the microphone for the voice — and in doing so used the open

tools of which the Rosicrucians dreamed. So far all of this concerned the human being. In the next part we cross the species boundary: we build the Babel fish and aim our reading devices at animals and plants.

TO DO YOURSELF — THE ACTIVITY FOR THIS CHAPTER

What the voice gives away that the words do not say (instructions online: hidden-signals.swarmcreativity.com/akt-10-1-en.html).

10.1 The voice as a signature. From a recording you read out pitch, loudness and speaking tempo. Say the same sentence cheerfully, tiredly and angrily and see how the curves change (librosa). As an extension, you have Whisper turn the recording into text.

Bonus: “Open sesame” — the voice key. As a class you build a door opener from voice: everyone records the same sentence, and a small model learns who is speaking.

Part III

The Babel Fish

Douglas Adams · The Hitchhiker's Guide to the Galaxy 1979

“Slide the Babel fish into your ear — and no creature stays mute to you.”

— after Douglas Adams

The same machine — here it reads animal and plant: from sound and electrical trace.

Chapter 11 · Plants Feel — and how you measure it

WHAT THIS CHAPTER IS ABOUT

Plants have no brain and no nerves — and yet they conduct electrical signals and react to touch, light and injury. This chapter shows that a plant sends out measurable signals, explains where they come from, and leads to perhaps the most surprising tool in this book: a home-built sensor that makes the “voice” of a plant visible as an image.

In Douglas Adams’ *The Hitchhiker’s Guide to the Galaxy* there is a small creature, the Babel fish: put it in your ear, and you understand every language in the universe. For centuries we believed that only humans, and perhaps a few animals, had anything to say at all; the rest of living things seemed mute. But as Chapter 1 already hinted, it was not nature that was mute — we lacked the Babel fish. And this becomes most astonishing with the creatures we credit with the least: plants.

Plants are not passive. A plant has no brain, no nerves, no mouth — and yet is anything but passive. The Venus flytrap snaps shut when an insect touches two of its sensory hairs; this is triggered by an electrical impulse that resembles a nerve signal astonishingly closely, even though the plant has no nerves. The mimosa folds up its leaves when touched. And as you read in Chapter 1, stressed tomato plants emit a clicking sound inaudible to us. Plants react — to touch, light, injury, drought. They just do it more slowly and quietly than we are used to.

Where the signals come from. These reactions have a measurable electrical side. Every plant cell maintains a small voltage between its inside and outside — like a tiny battery, on the order of a tenth of a volt. When the plant is stimulated, channels in the cell walls open, charged particles flow, and the voltage changes. Such changes propagate through the plant, from leaf to root. This is not language in the human sense, and certainly not a feeling — but it is information: the plant processes what happens around it and writes it into tiny fluctuations of voltage.

The sensor: a heart monitor for plants. These very fluctuations can be made readable — with a kit that fits on a breadboard and costs only a few francs. Its heart is a small chip called the AD8232, that red component in the photos. It was actually built to measure the human heart, an ECG sensor; it amplifies tiny voltage differences enough that they can be read out. Clip its electrodes onto a leaf and into the potting soil instead of onto the chest, and it measures the plant’s signals. Beside it sits the ESP32, a fingernail-sized microcomputer with Wi-Fi that reads the amplified signal many times per second and sends it to your laptop. A small OLED display shows the curve at once, so you can see that everything is working. We call this home-built probe the Biolingo sensor — again a child of the open tools from Part II.

How it all comes together. For the wiring you need only a few more parts: some coloured jumper wires — the “jumpers” — and a single resistor that smooths the slow

plant signal. The logic behind it is simple. The AD8232 gets its power from the ESP32 (3.3 volts and ground). Its output — the amplified voltage — travels back to a measurement input of the ESP32. And three adhesive electrodes are the actual ears on the plant's body: two you stick onto the plant, say onto two leaves or onto leaf and stem, the third serves as a reference point. At its core it is no more than that — the exact step-by-step guide with all the slots is online. And anyone who would rather not plug and wire at all can get the same sensor ready-made in a case.

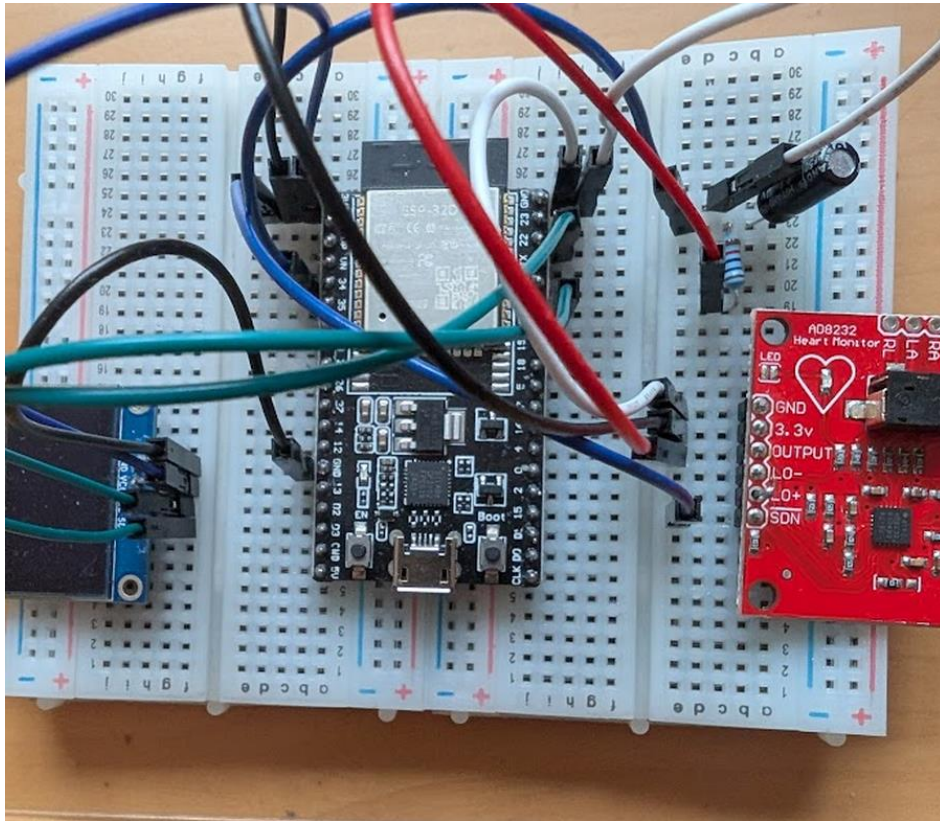


Figure 5: The finished, wired sensor on the breadboard. On the left the OLED display, in the middle the ESP32 microcomputer, on the right the red AD8232 with its labelled connectors (GND, 3.3v, OUTPUT ...) — joined by jumper wires and a resistor.



Figure 6: *The sensor in action. On the left the DIY build: two adhesive electrodes sit on the leaves of a plant. On the right the finished cased version with its own display, which you can buy if you would rather not solder and wire it yourself.*

From voltage curve to image. And now the trick from Chapter 3 comes into play again. The raw voltage curve of a plant looks like a jittery scribble — hard to interpret. So we transform it, as we did sound and voice before, into a spectrogram: an image of its frequencies. This turns the plant signal into something the same image-recognising machinery can read that otherwise tells cats from dogs. In this way we could, for example, tell from the electrical signal of a single plant alone, with over 85 per cent accuracy, whether a lamp in the room was switched on or off — and in the same series of experiments even, with around 73 per cent, whether the person in front of it was just then cheerful or sad. The plant “notices” the light — and perhaps more. This is a first, cautious hint of what Chapter 13 deepens.

When many experience it at once. This very act of making things visible can be experienced not only at your own breadboard but also in public. In the “Silent Signals” exhibition at the Phänomena in Dietikon, several such plant stations stood side by side for weeks. Above each raised bed hung a large screen that drew, in real time, the voltage curve of the plant growing behind it — the same jittery, light-blue line you also see on your own sensor. Whoever pressed a green button, struck a singing bowl or simply stepped kindly in front of the camera could watch the curve change shortly afterwards. Around ten thousand visitors thus saw the “voice” of a plant as an image for the first time — and felt how strangely close a living thing can come that one had until then held to be mute.



Figure 7: *The plant stations of the “Silent Signals” exhibition at the Phänomena. Above each raised bed a screen draws the plant’s voltage curve in real time — the light-blue line at the bottom edge. Visitors press the green button, strike the singing bowl or step in front of the camera and see how the plant answers.*

Hear, see, feel — and more. Look more closely at the stations and together they tell a larger story. At one you strike a singing bowl — and the plant hears. At the next a lamp goes on and off — and the plant sees the light. At the third you press a button and thereby touch the plant — it feels. And at the fourth it is enough to step in front of the camera cheerful or sad — and the plant seems even to perceive the person’s mood. Touch, light, sound, feeling: the same senses with which we too grasp the world, only slower and quieter. How a plant without eyes and ears notices all this we still do not know for certain; we suspect that the human heartbeat, the exhaled CO₂ and fine scent molecules (VOCs) play a part. How we followed that trail is told in Chapter 13.

If a plant already picks up whether the light is on — what else does it notice? The next two chapters push this question to the extreme: first towards animals, then towards perhaps the most astonishing possibility of all — that the plant notices us ourselves.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Four activities around the home-built plant sensor (Activity 11.1 comes as a fully worked-out worksheet).

11.1 The first plant sensor. With an AD8232 sensor, an electrode and a small microcontroller you build the Biolingo probe, clip it onto a houseplant and see the first voltage curve.

11.2 Touch and light. You observe how the plant answers to a touch and to changes of light.

11.3 Sound and vibration. You test whether and how the plant reacts to tones and vibration.

11.4 Human and plant. Sit down beside the plant, once quite calm, once agitated — and check whether it reacts to your presence and your mood.

Chapter 12 • The Language of Animals

WHAT THIS CHAPTER IS ABOUT

Your dog cannot tell you it is afraid — but its body does. This chapter aims the reading devices from Part II at animals: it shows how an AI reads the feelings of dogs, horses and cats from posture, face and sounds — often better than a human — and where the limits of this animal Babel fish lie.

A dog cannot tell you it is afraid. But its body says so loud and clear: the tucked-in tail, the crouched posture, the ears laid back. Animals send signals about their state continuously — through posture, face and sounds. Only we humans read them unreliably, often wrongly, and like to project our own feelings into them. The Babel fish for animals is therefore not a translator of words but a patient reader of bodies.

What does the dog feel? With my team I built exactly such a reader. We wanted to determine a dog's mood from a single photo — rage, fear, joy or relaxation. For this we first taught the software to find twenty-four body points on every dog image: muzzle, ears, base of the tail, paws. We trained this on almost fourteen thousand hand-marked dog images. From the position of these points relative to one another the model then recognises the mood — with a hit rate of 60 to 70 per cent. That sounds modest, but is remarkable: it surpasses, as studies show, the ability of most people to read dog emotions correctly. And it needs no sensor on the animal, only an image.

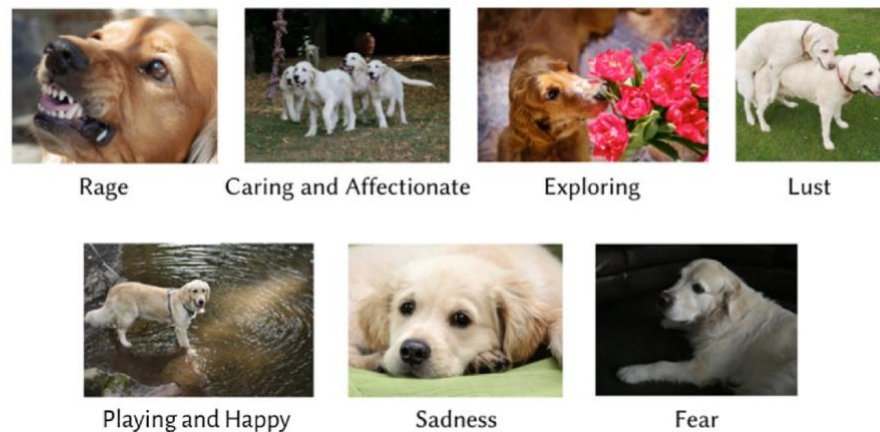


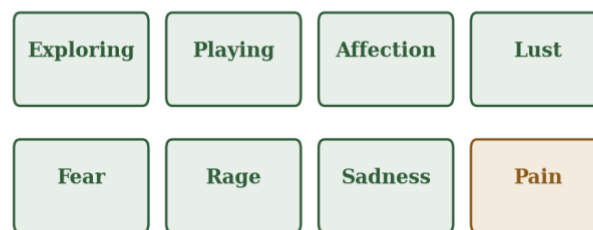
Figure 8: *Seven dog emotions, shown on a golden retriever — from rage through affection and exploring to fear. From such hand-labelled examples the model learns to recognise a dog's mood in a new photo. (From the dog paper: Bhave et al., Sensors 2024.)*

Horses, cats — and the same trick. The same method can be carried over to other species. In horses we read the facial expression with an image-recognising network; horses are especially interesting here because, conversely, they are extremely fine readers of human moods — more on that shortly. Cats are harder, for they hide their state; with them we listen more closely and read their sounds by — you can guess it —

turning the sound into a spectrogram. Whether posture, face or sound: in the end it always becomes an image, and the same tool from Chapter 3 reads it. One Babel fish, many species.

A shared alphabet of feelings. But with what do we even label an animal's feelings? We fall back on an order proposed by the neuropsychologist Jaak Panksepp, who identified seven basic emotions in all mammals: exploring, playing, affection, lust, fear, rage and grief. The beauty of it is that it is the same alphabet for dog and horse — and at heart for us humans too. For horses we added an eighth that counts especially with them: pain. From these categories the model learns, and on a new image it then reads them off.

The seven basic emotions of mammals, after Jaak Panksepp



The same seven categories serve our AI models as a vocabulary - for dog and horse alike. For horses, "Pain" is added as an eighth, especially important emotion.

Figure 9: *The shared emotional alphabet. The same seven basic mammalian emotions, after Jaak Panksepp, serve our AI models as a vocabulary — for dog and horse alike. For horses, “pain” is added as an eighth, especially important feeling.*

When the machine discovers something new. At first we showed the models the correct label with each image — for the dog they reached over seventy per cent accuracy this way, for the horse, which we first read from ears, mouth and eyes, about two in three. But then we turned the tables and let a model sort the animal images entirely without labels, purely by similarity. With almost four thousand images of wild horses it formed groups on its own — and some fit into none of our drawers. Perhaps, so the cautious suspicion goes, there are finer horse feelings for which we simply lack the words. Here it is no longer only the machine reading the animal; it points us to something we ourselves had overlooked. By now the same tool even reads the interplay between horse and human — who reacts to whom, and how.

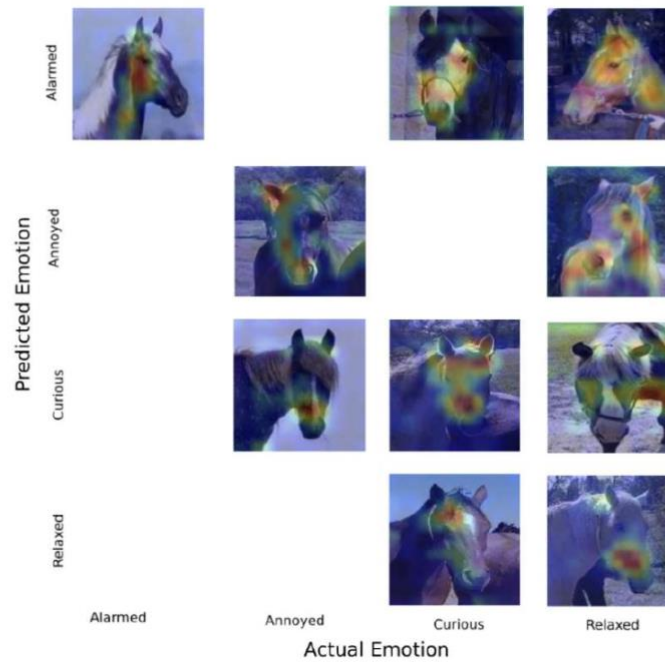


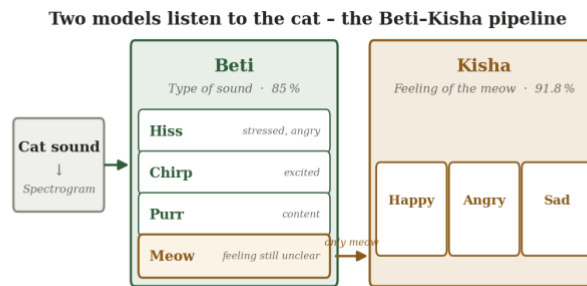
Figure 10: *Where the AI looks. Each horse image is overlaid with a “heat map” showing which spots the model uses for its decision — above all ears, eyes and mouth, exactly where we too read a horse’s mood. The rows show the emotion predicted by the model, the columns the actual one. (From the first horse paper: Corujo et al., Future Internet 2021.)*

What is in the meow?

Cats do not give themselves away through the face — but through sound.

Dog and horse carry their mood outward in posture and expression; the cat keeps it hidden. In return it talks to us: adult cats use the meow almost only towards humans, hardly among themselves. That is exactly where my team and I started — not with the image but with the sound. We collected over three thousand cat sounds and turned each one, as we had voice and plant signal before, into a spectrogram — an image of its frequencies that the same image-recognising machinery from Chapter 3 can read.

Two ears that listen one after the other. Instead of a single model we built two, chained one behind the other. The first, called Beti, sorts a sound first by its type: hiss, chirp, purr or meow — and is right in about eighty-five of a hundred cases. Even this rough division reveals something about the mood: a hiss sounds like stress or anger, a purr like contentment. With the meow, though, the feeling stays open at first — and here the second model, Kisha, takes over. It is specialised solely on meows and distinguishes whether a meow sounds rather happy, angry or sad — with nearly ninety-two per cent accuracy, even more precise than the first. From this two-stage chain came a little app called CatMotion: you hold a cat sound up to it and get back a cautious guess about the mood.



First Beti sorts the sound by its type; only for the meow does Kisha ask, in a second step, about the exact feeling. Both read - as throughout the book - the spectrogram of the sound.

Figure 11: *How two models interpret a meow. First “Beti” determines the type of sound — hiss, chirp, purr or meow — with around 85 per cent accuracy. Only if it is a meow does “Kisha” ask, in a second step, about the exact feeling: happy, angry or sad, and is right in about 92 per cent of cases. Both models read, as everywhere in this book, the spectrogram of the sound. (From the cat paper: Longhitano et al., 2026.)*

Where the Babel fish errs. But caution — and here responsibility is called for again. The model reads the expression, not the experience. A wagging tail is not always joy; a dog also wags in excitement or uncertainty. We humans tend to ascribe our own feelings to animals — and the machine can even learn this error along with us if the training images were already labelled by people with their preconceptions. Moreover, a model that learned on one dog breed reads another less well. The animal Babel fish is an aid, not an oracle.

What it is good for. Used rightly, though, it is a blessing. An animal cannot say where it hurts; a system that recognises early, from posture, that a horse is in pain or a dog lives in constant stress helps vets and owners to prevent suffering before it becomes bad. As with all reading devices in this book, the purpose decides: read for the animal’s good, it is care; read for its exploitation, it is the opposite.

One observation from this chapter will not let go of us: horses read us. Animals perceive our mood, often more accurately than we perceive theirs. And that leads to the boldest question of the whole book — whether a creature without eyes and ears might not read us too. A plant.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Read animals — through sound and through posture (instructions online: hidden-signals.swarmcreativity.com/aktivitaeten-en.html).

12.1 The language of the cat. You classify cat sounds — hiss, chirp, purr, meow — with a small, two-stage audio AI.

12.2 Dog and horse. You photograph an animal in different situations, let an emotion model guess how it is doing, and compare that with your own assessment.

Chapter 13 • Plants That Read Us

WHAT THIS CHAPTER IS ABOUT

Perhaps the most astonishing chapter of this book. If a plant reacts to light — does it also react to us? Here I present results from five years of research: that from a plant's electrical signal one can read how a person nearby is feeling — and why, with such an extraordinary claim, one must be especially careful and especially honest.

Now comes perhaps the most astonishing chapter of this book — and at the same time the one where I most insistently urge you to be sceptical. We have seen: a plant notices whether the light is on. The next question forces itself upon us, as crazy as it sounds: does it also notice us? Our nearness, our movement, our mood?

Do you remember the girl from the prologue — the Kalanchoe in the hall at Dietikon whose line rose as she approached and calmed when she grew calm? That was not a trick and not a story, but exactly the experiment this chapter is about. Now you are ready to understand it — and, more importantly, to test it.

An incredible result. Over five years of research, together with an international network of universities, we pursued this question. In the experiments a person sat near a plant and was by turns cheerful or sad — we measured their mood independently, for instance via a smartwatch or facial recognition. At the same time the sensor recorded the plant's electrical signal. Then we let a model — the same ResNet that otherwise recognises images — try to guess the person's mood from the plant's spectrogram. The result made us doubt ourselves: it succeeded with up to 97 per cent accuracy. A control model, whose labels we had deliberately scrambled, reached only 30 per cent — sheer chance. In other experiments the plant even recognised which of six people was present (66 per cent).

Why one must now be especially careful. So extraordinary a claim demands extraordinary evidence — and above all extraordinary honesty. So we tried to break the result ourselves. The most obvious explanation would be that it is not the plant sensing something but the sensor merely picking up electromagnetic interference in the room — from the person, the phone, the mains. So we repeated the experiment in a Faraday cage that shields such fields — once even in a car with closed doors and windows. The result held. That rules out the cheapest explanation. But it does not prove that we understand how the plant does it — and that is exactly what I write openly in my scientific papers too: these are mostly experiments with a single plant; they show feasibility, not final certainty. They must be repeated with many plants, on many days, in many laboratories.

I tell you this in such detail because here the whole book is put to the test. A responsible person does not believe an astonishing claim just because it is exciting — not even when it comes from a professor who has written a book about it. They ask: how was it measured? Which explanations were ruled out? Was it repeated? Practise these

questions calmly on my own results. Wonder and scepticism are not opposites — good science needs both at once.

Measuring feelings with plants

How does one even measure such a thing? Before I reveal what came out of it more precisely, it is worth a look at the workbench. To test whether a plant perceives a person’s feelings, two things must be measured at the same time and independently of each other: the person’s feeling and the plant’s signal. The human side we read without contact from the face — the facial recognition (FER) from Chapter 9 estimates from the expression whether someone is cheerful, neutral, angry or sad — and additionally from the heartbeat, measured with a chest strap (Polar H10) of the kind athletes wear — and, as calibrated emotional ground truth, from the Happimeter on the wrist (Chapter 5), which estimates from pulse and movement how happy and active a person is right now. This independently measured mood of the person is the yardstick against which the plant measurement can be calibrated in the first place. The plant side is picked up by various small measuring devices: an Oxocard microcomputer on the leaf of a banana plant, a BYB SpikerBox on the Christmas cactus, or our own Biolingo sensor from Chapter 11. Thus for every moment an honest pair of numbers arises — here the person, there the plant — that can later be compared.

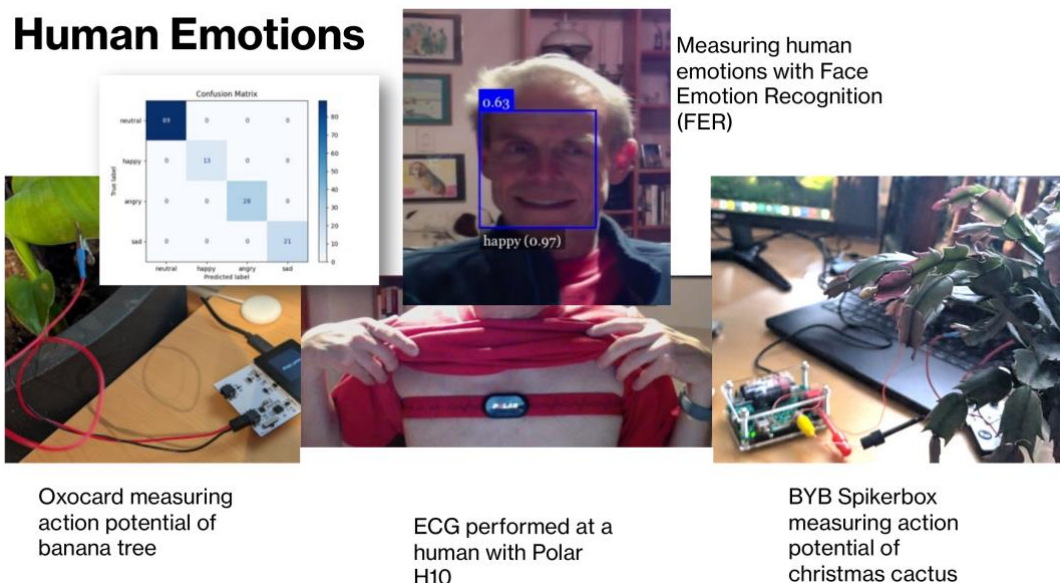


Figure 12: *The bench for measuring feelings with plants. Top centre, the facial recognition (FER), which estimates the emotion from the expression (here “happy”); bottom centre, the chest strap (Polar H10) for the heartbeat. On the left an Oxocard measures the electrical signal of a banana plant, on the right a BYB SpikerBox that of a Christmas cactus. The small table at the top is a confusion matrix: it shows how reliably the model tells the four moods — neutral, cheerful, angry, sad — apart.*

Everything at a glance: the dashboard. So that one can keep an eye on all these streams at once, they come together in a shared display we call EmotiPlant — a kind of

cockpit on the laptop. Right at the top the person’s heartbeat ticks in real time. Below it run — and this is the real knack — three plant sensors side by side: a main sensor and a second one directly on the plant, plus a third on a control plant in the same room but a good ten metres away. This control plant is our built-in honesty: if it showed the same correlation as the near plants, it would probably be down to interference in the room and not to real perception — it should stay flat. At the bottom the dashboard mirrors what the facial recognition currently sees, and translates it into a simple scale from sadness to joy — the same display that also ran at the Phänomena exhibition and gave the plant a visible reaction to the visitors’ mood.

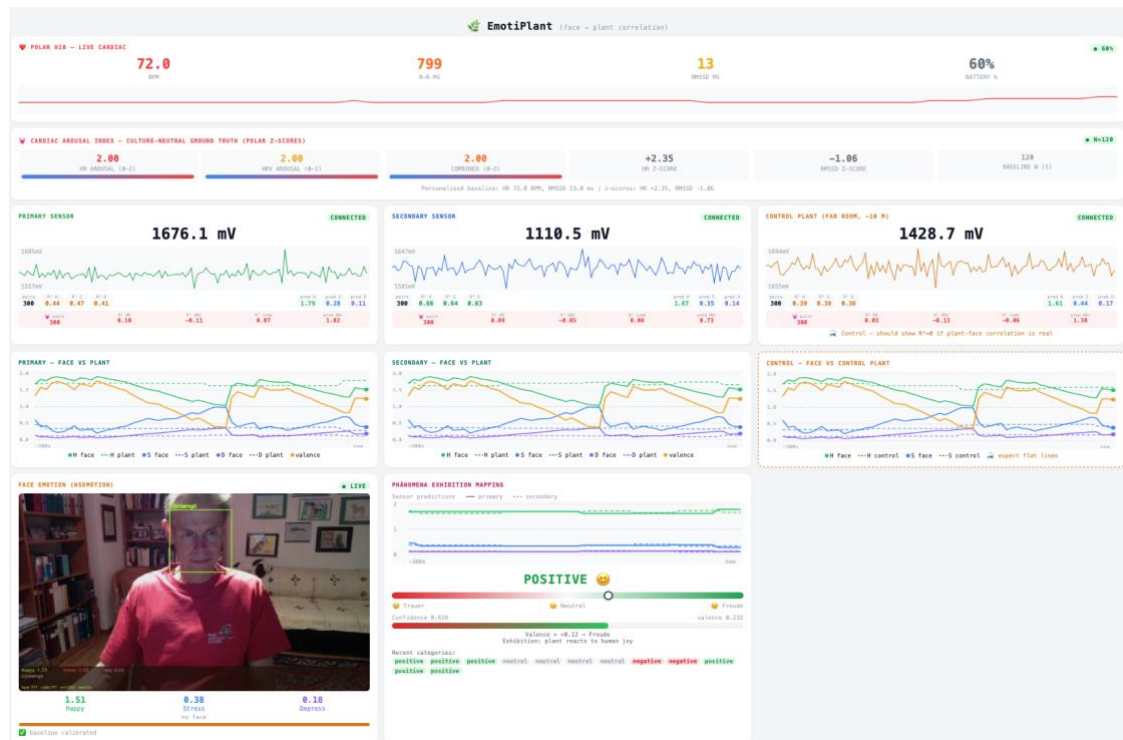
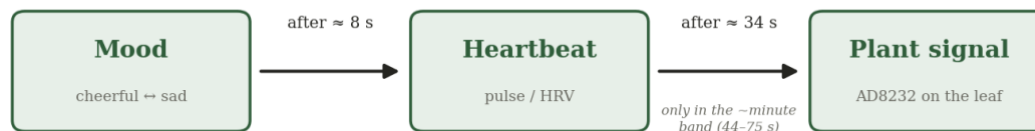


Figure 13: The EmotiPlant dashboard with which we collect the data. Top, the heartbeat (Polar H10); in the middle, three plant sensors at once — main, second and a far-away control plant; bottom left the live facial recognition and bottom right the translation of the expression into “joy” or “sadness”. (Dashboard shown in the original German.)

The heartbeat in the leaf. A trace of where the correlation might come from we found in a particularly careful self-study: a single test person, over seven sessions and around 93,000 seconds, beside a Kalanchoe. This time we measured not only mood and plant signal but a whole bundle at once: the pulse and heart-rate variability via a chest strap, the facial expression by camera, plus the CO₂ and the scent molecules (VOCs) given off by the body in the room air. So we could ask: by which route does the mood even reach the plant?

The answer was surprisingly precise. The plant signal does not follow the mood directly but the heartbeat — and with a clear delay: when the feeling changes, the pulse changes after about eight seconds, and only around thirty-four seconds later does this show in

the plant's signal. And not everywhere, but only in one very particular, slow oscillation — a beat of about one minute. This very selectivity is the real find: a mere electrical disturbance would appear in all frequencies. That only a narrow, slow band reacts looks more like a biological process — and recalls, astonishingly, the Venus flytrap, which adds up two touches within about thirty seconds before it snaps shut.



CO₂ ruled out as mediator · VOC strongest overall predictor · timing like the Venus flytrap

Figure 14: *How the mood reaches the plant — not directly, but via the heartbeat: delayed and only in a slow frequency band.*

The obvious chemical suspect, the exhaled CO₂, we could rule out as a mediator; the scent molecules remained the strongest indicator over longer periods. None of this is finally proven either — it is a single person, a single plant, a proof of feasibility. But it shows how one gets to the bottom of so incredible an observation step by step, honestly and cautiously — instead of either believing it too quickly or dismissing it too quickly.

With this the Babel fish turns around. Adams dreamed of a fish with which you understand every language. Here the plant becomes a Babel fish that reads us — and the boundary between “us” and “the world out there” proves more permeable than we thought. So far we have read individuals: a face, a voice, a person beside a plant. In the last part we dare the biggest step — from the individual to the crowd, from the person to the whole of society. There Isaac Asimov waits with an idea he called psychohistory.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Get on the trail of the plant's long memory (instructions online: hidden-signals.swarmcreativity.com/aktivityaeten-en.html).

13.1 Day and night. You record the plant signal over 24 hours and use a Fourier analysis (FFT) to look for the circadian rhythm — the plant's day-night beat.

13.2 Plant emotions with AI. From the signal you make a spectrogram and let a pre-trained model guess — pure inference, no training of your own.

Part IV

Psychohistory

Isaac Asimov · Foundation 1951

“You cannot predict the individual — but the crowd you can.”

— after Isaac Asimov

The same machine — here it no longer reads the individual, but the crowd.

Chapter 14 · Psychohistory — from the individual to the crowd

WHAT THIS CHAPTER IS ABOUT

So far we have read individuals: a face, a voice, a plant. Now the crowd becomes the subject. This chapter shows an old science-fiction idea that is becoming reality today: that the behaviour of large groups can be predicted, even though the individual person remains unpredictable — and why exactly this unpredictability is precious.

In 1951 the writer Isaac Asimov, in his novel *Foundation*, invented a science that did not exist: psychohistory. Its inventor, the mathematician Hari Seldon, can calculate in advance the fate of an entire galactic empire over millennia — not the single life, but the course of the masses. His trick is the same one with which physicists describe a gas: which single particle flies where cannot be predicted. But how billions of particles behave together — pressure, temperature, flow — follows clear laws. Seldon treats people like gas particles: individually unpredictable, in the mass astonishingly predictable.

Science fiction becomes science. For a long time this was pure invention. But today each of us leaves traces continuously — searches, posts, clicks, journeys — and in sum, millions of such traces make a measuring instrument for the mood and behaviour of whole societies. This is psychohistory in miniature, and it works. When many people suddenly google “flu”, a wave of illness can be predicted before the doctors report it. And the honest signals from Part II — the mood in language — reveal, in the mass, more than any survey could.

“I hope it is not as bad as I fear.” An example from my own work — and we were among the first to try it. As early as 2010 we measured the mood of millions of Twitter messages over six months: we simply counted how often words like “hope”, “fear” and “worry” appeared, and compared this emotion curve with the big stock indices — Dow Jones, S&P 500, NASDAQ. The title of our study was a real tweet: “I hope it is not as bad as I fear.” The result was strikingly simple: when the emotions ran high — much hope, fear and worry at once — the Dow tended to fall the next day; when it was emotionally calm, it rose. No one could say what a single person would do next. But the averaged mood of the crowd ran ahead of the price movement — exactly Seldon’s idea, in real time.

Not only the stock market. The same listening to the “web buzz” — the murmur of the mass online — we also tried elsewhere back then: we predicted who would win an Oscar and how well a film would open at the box office, by measuring how much and how confidently people talked about it beforehand. These were early, preliminary attempts — and in our blogs and papers we always noted openly that much more work is needed to really secure such a thing. But they all showed the same pattern: the

attention and mood of many, measured correctly, often reveals something about tomorrow already today.

Tribes in language. But the crowd has not only a mood — it also has a composition. And the same words that betray its temper also betray who is actually writing. From this idea another tool of my work arose, the Tribefinder. It reads people’s word choice in social networks and assigns them to “tribes”: groups connected not by origin or place of residence but by a shared attitude. You need not ask anyone. A neural network first learns from typical representatives of each tribe how they talk, and afterwards recognises from language alone where a new person most likely belongs. Someone who constantly writes “plant-based” and “sustainable” gives themselves away as a vegan before they say it — word choice is an honest signal, scaled up from the single person to millions.

Four tribes, four value sets. This becomes especially vivid with four tribes that differ by their values. The “Fatherlanders” prize tradition and authority; the “Nerds” believe in science and facts; the “Spiritualists” place compassion and non-material values above possessions; and the “Treehuggers” bet on sustainability and the environment. Hardly anyone fits fully into one of these drawers — most of us are a mixture, and that is a good thing. But at scale, averaged over millions of posts, this lets you read not only a society’s temper but also which values are growing louder and which quieter: psychohistory not merely of mood, but of conviction. That precisely this holds a danger too, we see in Chapter 16.

Four tribes, four value sets

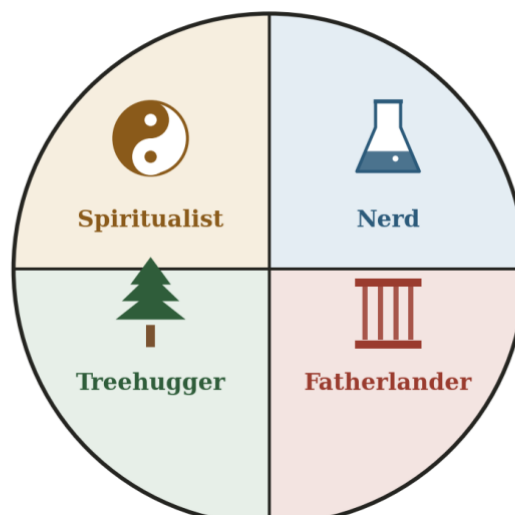


Figure 15: *Four tribes, four value sets. The Tribefinder infers from word choice which “tribe” a person stands closest to: Spiritualists place compassion and meaning above possessions, Nerds trust science and facts, Treehuggers sustainability and the environment, and Fatherlanders prize tradition and authority. Hardly anyone belongs wholly to a single one — most of us are a mixture.*

The Mule. Asimov was too clever to make his psychohistory all-powerful. In the novels a figure appears whom he calls the Mule: a single, unforeseen person with unusual abilities who throws Seldon's whole calculation into disarray. The message is deep: the prediction holds for the mass, never with certainty for the individual. There is always room for the unpredictable one who breaks ranks and changes everything. This unpredictability is not a flaw in the system — it is the place of your freedom. Remember the Mule; we return to him at the end of this part.

Before we come to the dark side of this power, let us look at its brightest application: not the predicting of stock prices, but the understanding of how groups create something new together. That is the question to which I have devoted my research life.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Now you no longer read individuals, but the crowd (instructions online: hidden-signals.swarmcreativity.com/aktivitaeten-en.html).

14.1 The mood of the swarm. You read the sentiment from social-media data and make the mood of a whole crowd visible over time.

14.2 Asimov's psychohistory. You lay the collective mood against events — stock prices, elections — and test whether the one predicts the other to some extent.

Chapter 15 • Swarm Creativity — how groups create something new

WHAT THIS CHAPTER IS ABOUT

This chapter brings together two threads of the book: the honest signals from Part I and the open colleges from Part II. It shows how you can tell whether a group is creatively flourishing — and how, with the same tools, you read a whole team as you earlier read a single person. It is the heart of my own research.

For over twenty years I have researched a single question: how do groups of people create something new together that none could have managed alone? I call such groups Collaborative Innovation Networks — COINs for short, networks of shared renewal. From the inventors' team through the open-source community to the school class planning a trip: the best always arises not in the single head but between the heads.

The six signals of a group. Recall the six honest signals from Chapter 5 — there they applied to individual people, but they apply just as well to a whole group. Does leadership change with the topic, or does one person cling to power? Does everyone answer one another quickly, or do contributions seep away? Does everyone speak roughly equally, or do a few dominate? Is the mood in the language honest and warm? And does the strong, connecting leadership — the betweenness centrality — travel through the group? These measures predict astonishingly well whether a team will achieve great things, long before the result is there.

Who builds the swarm. That the travelling leadership really counts, we did not merely suspect but measured. Together with colleagues I observed sixteen medical online communities over seven years, in which around fourteen thousand professionals exchanged ideas. The question was simple: which of these communities grow, which shrink? Our study put the result in its very title — “It is rotating leaders who build the swarm.” The communities that grew most strongly were those in which leadership passed briskly from one to the next, in which there were clear hubs, and in which people wrote simply and understandably. Where, by contrast, the same people always had the say, the swarm ran dry.

Making the group visible. To measure this, you need not question the group — you need only watch it collaborate. From the patterns of communication (who writes to whom, when) we draw, with tools like our program Condor, a living network that changes over time. You see the swarm grow, see hubs arise and pass away, see when a group opens up and when it stiffens. It is the same map I once drew from my Deloitte inbox (Chapter 8) — only in motion.

The mirror changes. And now the most astonishing thing: as soon as you show a group this picture of itself, it changes. We call this — as with the Happimeter in Chapter 5 — virtual mirroring. In a large service company we measured teams' email patterns

for two years and regularly mirrored back to their leaders how they communicate: how quickly they answer, how understandably they write, how much a single person pulls everything to themselves. This mirror alone was enough. The teams that saw it became more attentive and more open — and their customers' satisfaction rose measurably, while in the comparison teams without a mirror it even fell. Measuring here is not cold surveillance but an invitation to self-knowledge: whoever sees themselves honestly becomes better.

Reading the mood of the crowd. Added to this is language itself. A machine can grasp the tone of a thousand messages at once — confident or anxious, open or irritable — and read from it the collective mood, the emotion curve of a group. This is the sentiment analysis from Part II, scaled up from one person to many. And it is exactly the signal that predicted the stock market in Chapter 14.

With these tools, the reading of single people becomes the reading of whole communities. And with that the last, greatest possibility lies at hand — and the last, greatest danger: to build from all of it a complete model of a person or a group. A digital twin.

Chapter 16 · The Digital Twin — and where the journey ends

WHAT THIS CHAPTER IS ABOUT

The last chapter of this part brings everything together — and poses the most serious question of the whole book. If you read a person’s face, voice, language and network, you can build from it a model that imitates and predicts them: a digital twin. This chapter shows what is good about it, what is dangerous — and why in the end the unpredictable person remains the most valuable thing.

Imagine everything this book has taught you to read flowing together: your face, your voice, your words, your place in the network of your friendships. From all of it a model can be built that predicts astonishingly well how you will react to something — a digital twin. Not science fiction: in rudimentary form it has long existed.

Tribes — how your words betray where you belong. Recall the Tribefinder from Chapter 14, which assigns people to a “tribe” from their word choice alone — from Treehugger to Fatherlander. With the digital twin its second side appears. What helps a company understand its customers can also put people into drawers they never chose for themselves. You need ask no one a question; language reveals belonging by itself. That is powerful — and uncanny at once.

A twin from five relationships — the case of Mia

How does such a twin arise? Let us build one, small and transparent. Take Mia, a schoolgirl. We collect — only with her consent — her WhatsApp messages to five people: her sister, her teacher, her best friend, her mother, and her boss at her holiday job. On each message we set loose the symbiont tool from Chapter 7 and count which of the five roles — Bee, Ant, Butterfly, Capybara, Leech — Mia is playing at the time.

MIA’S MESSAGES — FIVE RELATIONSHIPS (SYNTHETIC)

To her sister: *“Omg I have the best idea for Dad’s birthday — let’s build a scavenger hunt through the whole flat! 🐜”*

To her teacher: *“Dear Mr Berger, please find attached the outline of my presentation; the sources are numbered, submission on Friday as agreed.”*

To her best friend: *“Hey, you were so quiet today ☐ Do you want to talk? No pressure — I’m just here.”*

To her mother: *“Mum, will you drive me to the station at 8? And lend me 20 francs? I’ll pay it back, promise 🙏”*

To her employer: *“Good morning, I’ll take the early shift as agreed. Delivery is stocked, till cashed up.”*

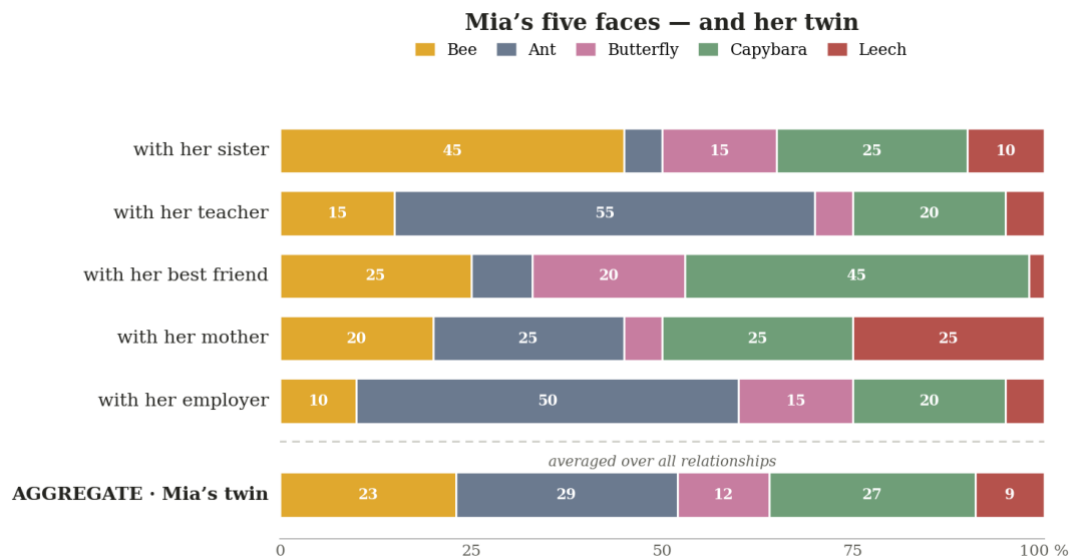


Figure 16: *Mia's five faces.* For each relationship the bar shows which symbiont roles Mia plays there. With her sister the Bee predominates, with teacher and boss the Ant, with her friend the Capybara — and with her mother a good dose of Leech comes through. The bottom bar is the aggregate: the average over all five relationships — Mia's “overall twin”.

Five faces, one person. With her sister Mia is above all Bee — playful, full of ideas. With the teacher she becomes the Ant: orderly, dutiful, punctual. With her friend she is a Capybara that listens and comforts. With her boss, the Ant again. And with her mother? There, quite honestly, a good dose of Leech comes through — Mia takes more here than she gives: “Can you drive me?”, “Will you lend me twenty francs?” That is often how it is between children and parents, and it makes Mia no bad person. It only shows: there is not the one Mia. There are as many Mias as she has relationships.

The twin depends on the counterpart. That is the first great lesson of this example: a digital twin is never the picture of a person alone — it is the picture of a relationship. Whoever knows Mia only from the chat with her mother takes her for a little Leech. Whoever sees her only with her friend, for a gentle Capybara. Both are right, and both are wrong, because they take a part for the whole.

A good twin is an aggregate. So what is Mia's “true” twin? Most likely the bottom bar: the mean over all her relationships. Only this aggregate smooths the outliers and shows a balanced person — much Ant and Capybara, a fair amount of Bee, little Leech. But beware: the aggregate too is only a summary. It conceals that Mia sparkles with her sister and nags with her mother. The more relationships a twin knows, the fairer it becomes — but the full truth about a person never sits in a single number.

Build it yourself. The whole calculation — from the five dialogues to these bars — you can retrace as a small, commented Python program on hidden-signals.swarmcreativity.com/kapitel16-zwilling-en.html and try out with your own, voluntarily shared chats. It is the same two-step as in Chapter 7: first use the finished

version, then rebuild it yourself and dispel the magic. The simple word-counting version there comes out sharper than the fine mixtures above — a good feel for how much more the “real” tool, with its language model, sees on top.

The twin as helper. At its best a digital twin is a gift. A model of your team can show you where the collaboration is about to snag before it breaks. A twin of yourself can be a patient mirror that helps you understand yourself better — like the face mirror, only deeper. Doctors simulate a treatment on the digital twin of an organ before they risk it on the person. Understand, practise, protect — that is what the tool is made for.

The twin as fetter. But the same twin, in someone else’s hands, becomes a fetter. Whoever can predict you accurately enough can also steer you: show you exactly the message that grabs you, exactly in the moment of your weakness. This is no distant invention — it is the business model of half of whole industries. And the better your twin, the more easily you are steered without noticing. Here the honest limit of the whole book becomes the most pressing question: a tool that understands can also manipulate.

The golden rule — and the Mule. Against both, the same thing helps that already stood in Chapter 9: the golden rule. What is found out about you belongs to you — not to whoever measures it. Your digital twin should serve you, not dispose of you. And then there is still the Mule from Chapter 14, that unpredictable one who bursts every prediction. He is the last word of this part: no model, however good, captures the whole person. You are always also the one who breaks ranks, who decides differently, who disappoints the statistics. This freedom no machine takes from you — as long as you know you have it.

With this the great arc closes: from the single honest signal through face, voice, plant and animal to the whole of society. What connects all these threads and where they lead together, the last part brings together — the Convergence.

TO DO YOURSELF — THE ACTIVITIES FOR THIS CHAPTER

Everything runs together in the digital twin (interactive version and instructions online: hidden-signals.swarmcreativity.com/kapitel16-zwilling-en.html).

16.1 The twin as aggregate — the case of Mia. You compute the symbiont proportions that Mia shows with five different people, and see in black and white: the twin depends on the counterpart — it is a picture of the relationship, not of the person alone.

16.2 My digital twin (Master). For advanced readers: fine-tune a mini language model on your own, voluntarily shared texts — and experience how much and how little of it is “you”.

Part V

The Convergence

Paracelsus · Rosicrucians · Adams · Asimov

“To listen across the kingdoms — that is not a crown, but a responsibility.”

— the guiding idea of this book

The same machine — now in all worlds at once. And it reads you too.

Chapter 17 • A Shared Timeline

WHAT THIS CHAPTER IS ABOUT

For four parts you have learned to read individual beings — humans, animals, plants, whole groups. This chapter finally lays them side by side, on a shared timeline, and shows the astonishing thing: it was always the same method, and always the same signals. The kingdoms of nature, once made audible, speak a surprisingly similar language.

At the beginning of this book stood a promise: that in the end plant, animal and human would lie side by side on a shared timeline — and that you would see what connects them. That moment has now come. Look back on the path, and a pattern emerges that previously lay hidden behind the individual chapters.

A single method. Look closely, and you recognise it: it was always the same move. A voice, a plant signal, a dog's body, the mood's ups and downs of a whole group — however different they are, with each we did the same thing. We recorded the signal over time and turned it into an image: a spectrogram, a point cloud, a network map. And then we let the same image-recognising machinery look at it — the trick from Chapter 3. A single method, from the tomato to the stock market. That is the first convergence: that of method.

The same honest signals. The second is deeper. Across all the kingdoms the same honest signals appear that Pentland described in the human being (Chapter 1). The rhythm in which something answers. The rise and fall of tension and calm. The recurrence of fixed patterns. The slow beat of the heart mirrored in the plant's leaf (Chapter 13); the time windows of the Venus flytrap recurring in the minute-band; the mood in language, whether in one person or in millions at once. It is as if life everywhere spoke the same basic dialect — and we have only just learned to hear it.

Viewed on a shared timeline, human, animal and plant thereby move surprisingly close together. Not because they are the same — but because they are akin in the way they answer their world. And that turns a very old order on its head.

TO DO YOURSELF — THE ACTIVITY FOR THIS CHAPTER

Bring the three kingdoms together (instructions and data online: hidden-signals.swarmcreativity.com/aktivitaeten-en.html).

17.1 Three kingdoms on one timeline. You lay human, animal and plant from the Phänomena data side by side on a shared timeline and look for what connects them across the kingdoms.

Chapter 18 • The Inverted Ladder

WHAT THIS CHAPTER IS ABOUT

For two thousand years we have thought of nature as a ladder: the human at the top, below it animal, plant, stone. This chapter shows why the tools of this book invert that ladder — and what takes its place: not a rank, but a net of mutual reading, in which every being is perfect in its own way.

Aristotle arranged living things into a ladder — the *scala naturae*. Right at the top the human, below it the animals, then the plants, at the very bottom the inanimate. For two thousand years this image shaped our thinking: the “higher”, the more valuable, the closer to reason and mind. The hidden signals turn this ladder on its head.

Who reads whom here? For who, in this book, actually reads whom? The dog reads our mood often more accurately than we read its (Chapter 12). The horse senses our tension before we notice it ourselves. And in the end it was the plant — the being right at the bottom of the old ladder — that read the human (Chapter 13). The rank order that placed the human at the top and the plant at the bottom falls apart as soon as one listens more closely. Top and bottom lose their meaning.

Incomparable perfections. In its place comes another thought: each kingdom is perfect in its own, incomparable way. The plant can do what no human can — make life from light and read its environment with a patience we lack. The swarm creates what no individual can (Part IV). Such abilities cannot be set off against one another; there is no common scale on which one would stand “higher” than the other. They are different and equal. That is the inverted ladder: no longer a ladder, but a net.

The one distinction of the human. One single distinction of the human remains — but it is not the one we long thought. The human is not the crown of creation. It is the only being that can build the tools to hear the others at last. That is no crown. That is a responsibility. Whoever alone can listen bears the duty to listen well — and not to misuse what is heard.

With this the guiding idea is fulfilled that stood at the very beginning of this book: to listen across the kingdoms. One last question remains — what this means for you, who now hold these tools in your hands.

Chapter 19 • The Open Vault — a Word to You

WHAT THIS CHAPTER IS ABOUT

The last chapter. It returns to the beginning — to the vault of the prologue — and addresses you. You have learned to read the hidden signals. What now? This chapter draws the book's five perspectives together into a stance and gives you the most important thing a responsible person needs in a readable world.

At the beginning of this book stood a vault: the fabled crypt of the Rosicrucians, in which an artificial light burned and a knowledge lay locked away — with a wager of one hundred and twenty years that one day the instruments and the minds would come to read it. Four centuries later, so the prologue told, a second vault stands in a hall in Dietikon, made of glass, cable and code. You have now passed through it. The vault stands open, and you now know how to read what shines within it.

Five perspectives, one stance. Four pioneers have accompanied you, and each leaves you more than a technique — he leaves a stance. From Paracelsus: read nature itself, not the old books about it. From the Rosicrucians: share your knowledge openly, across all walls. From Douglas Adams's *Babel fish*: cross the species boundary, hold no being to be mute. From Asimov's psychohistory: see the whole, the patterns of the crowd — and honour the individual who breaks ranks. And from their convergence: listen across the kingdoms, in humility. Together they are not a toolbox but a way of seeing the world.

You read — and you are read. But do not forget the most important thing: in a readable world you do not only read, you are also read. Your face, your voice, your words, your traces — others read them already, and not always for your good. Responsibility means knowing both: that you have this power, and that others have it over you. So ask, with everything you read: do I have the right to this? And with everything that is read about you: does it happen openly, and does the result belong to me? Remember the golden rule — what is found out about you belongs to you. And remember the Mule: no prediction captures you entirely. Your freedom to be other than any statistic expects, no one can take from you.

The old Rosicrucians dreamed of a fellowship that cracks the code of nature and shares its knowledge for the good of all. You now hold part of this code in your hand — earlier than they dared to hope, and more powerful. What you do with it is not in this book. It rests with you.

And do you remember the girl from the prologue, at the table with the *Kalanchoe*? We had promised to return. Now you know what happened there: no miracle, but an honest signal, read by a tool you can by now operate yourself — and place in context, with wonder and with scepticism. Perhaps it has long dawned on you who this girl is. She bore no name and no face, because the place in front of the plant was kept free for you.

Step up. Bend over the plant. See how the line rises.

The vault stands open. The light within it is bioelectric.

Go through — and listen well.

TO DO YOURSELF — THE CLOSING ACTIVITY

Not a building task but a thinking task to close.

19.1 AI literacy — where are the limits? Together, collect what the tools of this book can do — and what they cannot. Whoever has built an emotion recognition themselves no longer confuses it with mind-reading.

Sources and Further Reading

This book combines the ideas of astute pioneers with real research — much of it from the author’s laboratory, often together with students and colleagues from an international network. Anyone who wishes to dig deeper will find here the most important sources, ordered by theme. This list grows with the book.

Books and Inspirations

- Adams, D.: *The Hitchhiker’s Guide to the Galaxy*. London: Pan Books, 1979.
- Asimov, I.: *The Foundation Trilogy*. 1951–1953.
- Pentland, A.: *Honest Signals: How They Shape Our World*. Cambridge, MA: MIT Press, 2008.
- Pennebaker, J. W.: *The Secret Life of Pronouns: What Our Words Say About Us*. New York: Bloomsbury Press, 2011.
- Haidt, J.: *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Pantheon Books, 2012.

The Research Behind This Book

A large part of this book rests on the author’s own research and that of his collaborators.

- Gloor, P. A.: *Swarm Creativity: Competitive Advantage through Collaborative Innovation Networks*. New York: Oxford University Press, 2006.
- Gloor, P.; Fischbach, K.; Fuehres, H. et al.: *Towards ‘Honest Signals’ of Creativity – Identifying Personality Characteristics Through Microscopic Social Network Analysis*. *Procedia – Social and Behavioral Sciences* 26, 2011 (COINs 2010).
- Roessler, J.; Gloor, P. A.: *Measuring Happiness Increases Happiness (Happimeter)*. *Journal of Computational Social Science* 4, 123–146, 2021.
- Antonacci, G.; Fronzetti Colladon, A.; Stefanini, A.; Gloor, P. A.: *It Is Rotating Leaders Who Build the Swarm: Social Network Determinants of Growth for Healthcare Virtual Communities of Practice*. *Journal of Knowledge Management* 21 (5), 1218–1239, 2017.
- Gloor, P. A.; Fronzetti Colladon, A.; Giacomelli, G.; Saran, T.; Grippa, F.: *The Impact of Virtual Mirroring on Customer Satisfaction*. *Journal of Business Research* 75, 67–76, 2017.
- Zhang, X.; Fuehres, H.; Gloor, P. A.: *Predicting Stock Market Indicators Through Twitter – ‘I Hope It Is Not as Bad as I Fear’*. *Procedia – Social and Behavioral Sciences* 26, 55–62, 2011 (COINs 2010).
- Gloor, P. A.: *Predicting Stock Market Indicators Through Twitter (blog post)*. *Swarm Creativity Blog*, 2010 — early results on stock-market, Oscar and box-office prediction; swarmcreativity.blogspot.com.
- Gloor, P. A.; Fronzetti Colladon, A.; de Oliveira, J. M.; Rovelli, P.: *Put Your Money Where Your Mouth Is: Using Deep Learning to Identify Consumer Tribes from Word Usage (Tribefinder)*. *International Journal of Information Management* 51, 101924, 2020.

- Jha, A.; Gloor, P. A.: *Do What You Say — Computing Personal Values Associated with Professions Based on the Words They Use* (Fatherlanders, Nerds, Spiritualists, Treehuggers). *Algorithms* 18 (2), 72, 2025.
- Gloor, P. A.; Fronzetti Colladon, A.; Altuntas, E. et al.: *Your Face Mirrors Your Deepest Beliefs — Predicting Personality and Morals through Facial Emotion Recognition*. *Future Internet* 14 (1), 5, 2022.
- Corujo, L. A.; Kieson, E.; Schloesser, T.; Gloor, P. A.: *Emotion Recognition in Horses with Convolutional Neural Networks*. *Future Internet* 13 (10), 250, 2021.
- Ferres, K.; Schloesser, T.; Gloor, P. A.: *Predicting Dog Emotions Based on Posture Analysis Using DeepLabCut*. *Future Internet* 14 (4), 97, 2022.
- Bhave, A.; Hafner, A.; Bhave, A.; Gloor, P. A.: *Unsupervised Canine Emotion Recognition Using Momentum Contrast*. *Sensors* 24 (22), 7324, 2024.
- Bhave, A.; Kieson, E.; Hafner, A.; Gloor, P. A.: *Identifying Novel Emotions and Wellbeing of Horses from Videos Through Unsupervised Learning*. *Sensors* 25 (3), 859, 2025.
- Arora, S.; Kieson, E.; Rudd, C.; Gloor, P. A.: *Multi-Modal Feature Fusion and Hierarchical Classification for Automated Equine–Human Interaction Behavior Recognition*. *Sensors* 26 (7), 2202, 2026.
- Longhitano, D.; Sauchuk, Y.; Asadov, A.; Gloor, P. A.: *Leveraging Deep Learning to Predict Cat Emotions Using Audio*. In: *Artificial Intelligence and Networks for a Sustainable Future*. Springer, Cham, pp. 95–117, 2026.
- Oezkaya, B.; Gloor, P. A.: *Recognizing Individuals and Their Emotions Using Plants as Bio-Sensors through Electro-Static Discharge*. arXiv:2005.04591, 2020.
- de la Cal, L.; Gloor, P. A.; Weinbeer, M.: *Can Plants Sense Humans? Using Plants as Biosensors to Detect the Presence of Eurythmic Gestures*. *Sensors* 23 (15), 6971, 2023.
- Kruse, J. A.; Ciechanowski, L.; Dupuis, A. et al.: *Leveraging the Sensitivity of Plants with Deep Learning to Recognize Human Emotions*. *Sensors* 24 (6), 1917, 2024.
- Gloor, P. A.: *Plant Bioelectric Early Warning Systems: A Five-Year Investigation into Human–Plant Electromagnetic Communication*. arXiv:2506.04132, 2025.
- Gloor, P. A.: *Plant Bioelectrical Signals for Environmental and Emotional State Classification*. *Biosensors* 15 (11), 744, 2025.
- Gloor, P. A.: *Silent Signals: Correlating Plant Bioelectric Activity with Human Emotional States via Wearable Sensing*. *Biomimetics* 11 (4), 236, 2026.
- Gloor, P. A.: *Mechanosensory Coupling Between Human Cardiac Activity and Plant Bioelectric Potentials: A Naturalistic Self-Study with Spectral Validation*. Preprints 202605.0055, 2026.

Technical Foundations and Further Sources

- He, K.; Zhang, X.; Ren, S.; Sun, J.: *Deep Residual Learning for Image Recognition* (ResNet). *IEEE CVPR*, 2016.
- Mathis, A. et al.: *DeepLabCut: Markerless Pose Estimation of User-Defined Body Parts with Deep Learning*. *Nature Neuroscience* 21, 1281, 2018.
- Radford, A. et al.: *Robust Speech Recognition via Large-Scale Weak Supervision* (Whisper). OpenAI, arXiv:2212.04356, 2022.
- Khait, I. et al.: *Sounds Emitted by Plants under Stress Are Airborne and Informative*. *Cell* 186 (7), 2023.

About the Author

Peter A. Gloor has spent over twenty years researching how something can be read from the hidden signals of people, teams — and by now of animals and plants too. He is Honorary Professor of Information Management at the University of Cologne, teaches Data Science at the Lucerne University of Applied Sciences, and is a Research Affiliate at the Massachusetts Institute of Technology (MIT) in Cambridge, where he researched for over two decades. In 2024 he returned to Switzerland.

In the 1990s he worked at MIT on one of the first “maps” of the then still young World Wide Web. From this early fascination with networks grew his life’s theme: Collaborative Innovation Networks (COINs) — the idea that creative breakthroughs arise in open groups that communicate honestly with one another. His book “Swarm Creativity” coined the term.

Between his early years at MIT and his long return there, he gathered experience in industry: as head of the software-development section at UBS, as partner and head of eBusiness at PwC, and finally as partner and European head of eBusiness at Deloitte. These years at the interface of technology and large organisations shaped the questions he then pursued for twenty-two years as a research scientist at the MIT Center for Collective Intelligence: how do people really work together, and how can that be read?

Today he leads an international research network that reads the electrical signals of plants with home-built sensors and artificial intelligence. The experiments in this book — from the dog’s face through the cat’s meow to the Kalanchoe that seems to sense a person’s heartbeat — come largely from his own laboratory, often together with students from all over the world. He lives in Aarau.