



PAPER

Gainful Compression, Significance and the Formation of We

Michael Kontorovich · 7 September 2026 · 12 min read

IN BRIEF

What is expertise, and how could it reach more people without being diluted? An expert has compressed years of encounters into a few things worth noticing, and stays alert to the case that breaks the pattern. AI can compress thousands of audience encounters and bring the anomalies back to the expert, and trust lets what matters to one person start to matter to another. This is the learning loop behind a mindcast. It develops the idea of a “we” first set out in Semantic Field Dynamics, and applies the self-observation described in Learning the Next Mind to the expert.

Someone said that intelligence is compression.

There is something deeply right about this.

An intelligent system does not preserve every detail of every encounter. It discovers regularities. Observations become patterns; patterns become concepts; concepts allow the system to navigate situations it has never encountered before.

The expert is perhaps the clearest example.

Thirty years of expertise does not mean carrying thirty years of experience around as an enormous archive. Repeated encounters have changed the structures through which the expert perceives the next one.

The novice sees a hundred things.

The expert sees three.

Those three things may contain thirty years of experience.

This is gainful compression: information has been discarded, yet capability has increased.

But compression alone is not intelligence.

A compressed representation may preserve the wrong things. Random information may be difficult to compress without being meaningful. And a beautifully compressed theory can still be wrong.

What determines what should be preserved?

Information, surprise and significance

Shannon gives us part of the answer.

The information carried by an event is related to its improbability:

$$I(x) = -\log P(x)$$

The expected carries relatively little new information. The unexpected carries more.

Once a system has discovered a regularity, it has an expectation. What conforms to that expectation requires little additional representation. What violates it becomes visible.

Good compression makes surprise detectable.

But surprise alone is insufficient.

A random fluctuation can be extremely improbable and completely irrelevant.

Something else is required: significance.

Why does one difference matter while another does not?

A living system gives us a clue because a living system has something at stake.

It must continually maintain the conditions that allow it to remain the kind of system it is. For a cell, concentrations, gradients, energy availability and molecular organisation must remain within viable ranges.

Environmental differences therefore cease to be equivalent. Some differences affect the system's ability to maintain itself. Others do not.

Before life: differences cause differences.

With life: some differences matter.

This may be the primitive origin of significance.

significance = significance for a system

Nothing is simply significant. It is significant to something, with respect to something.

Meaning is compressed significance

A living system encounters enormous variation. It cannot preserve every detail.

Through repeated encounters, it discovers which differences have proved significant and which can safely be ignored.

$$X_1, X_2, X_3, \dots, X_n \rightarrow X$$

The representation X discards enormous amounts of information. But if it preserves the distinctions that matter, the compression is gainful.

Meaning is the gainful compression of significance.

Meaning is therefore not simply compressed information. It is compressed history concerning what has mattered.

This explains how a very small signal can carry enormous meaning.

A particular sound in the grass can activate an accumulated structure concerning predators. A particular number in a pathology report can activate decades of clinical experience. A particular hesitation from a client can tell an experienced therapist that something important has happened.

The representation is tiny compared with the history compressed within it.

We might call this meaning density:

$$\text{meaning density} \sim \text{significance preserved} / \text{representation required}$$

Expertise is extraordinarily high-density meaning.

Significant surprise

But every compression throws something away. That is not a defect. It is what makes compression useful.

The problem occurs when reality demonstrates that one of the distinctions we discarded actually mattered.

Suppose repeated encounters teach a system:

$$A_1 \sim A_2 \sim A_3 \sim A_4 \rightarrow A$$

Then it encounters:

$$A_5 = A + d$$

Its existing meaning treats d as irrelevant. But this time d changes the outcome. A previously discarded distinction has become significant.

This is significant surprise.

Your compression lost something you needed.

The system may now have to differentiate and construct a new compression:

$$A \rightarrow \{A^+, A^-\}$$

Learning therefore does not consist simply of progressively compressing experience. Sometimes learning requires decompression.

A novice sees thousands of differences because they do not know which ones matter. An intermediate learns the pattern and sees one thing.

The great expert says: Yes, normally this is X. But not this time.

Expertise therefore combines two capacities:

expertise = high-resolution significance detection + high-density meaning

This suggests that one of the most revealing questions we can ask about an expert is not simply: What do they know?

What surprises them?

The things that make an expert stop expose the boundary of their existing compression.

Intelligence as a loop

We can now describe intelligence as a recursive process:

*significance → gainful compression → meaning → expectation → significant
surprise → recompression*

Intelligence is not possession of a perfect model. It is the continual construction and repair of useful meaning through encounters with reality.

Intelligence is the gainful compression of significance into meaning, continuously revised through significant surprise.

The neuron as physical compression

Neurons provide an intriguing physical example.

A cortical neuron can receive thousands of synaptic inputs. Yet it does not appear to process these as thousands of independent symbolic messages requiring central inspection.

Its physical architecture performs successive integrations. Synapses interact locally. Dendritic branches can act as nonlinear computational subunits. Patterns of activity become local events. Those events contribute to the changing state of the neuron and ultimately to a comparatively small output: patterns of spikes.

At the same time, synaptic plasticity allows previous encounters to alter how future inputs will be interpreted.

The neuron does not need to retain a transcript of every previous event.

History becomes structure.

The vast complexity of previous interaction becomes a relatively compact physical configuration determining sensitivity to what happens next.

Compress the common. Remain sensitive to the significant exception.

Expertise as compressed relationship with reality

Human expertise appears to have a similar structure at a radically different scale.

An expert is not simply someone who possesses more information. Expertise consists partly in knowing what matters.

Two people can receive exactly the same information and perceive completely different situations.

The expert notices: That.

A surgeon sees something significant in a clinical trajectory that a novice overlooks. A therapist hears something in a sentence that changes their understanding of the interaction. An entrepreneur sees a number and thinks: That doesn't smell right.

Years of encounters with reality have changed what becomes significant to them.

This is why expertise cannot be understood merely as a database of facts. The expert carries a compressed history of their relationship with reality.

And crucially, that compression remains vulnerable to reality. Reality can still produce the significant surprise that makes the expert say: I was wrong.

That vulnerability is not a weakness of expertise. It is what keeps expertise alive.

The problem of many humans

Now put many intelligent humans together.

Each has their own loop:

significance → compression → meaning → attention

Connecting those people does not automatically create collective intelligence. It may simply create many individual intelligences competing for attention.

Imagine 100,000 people interacting with an expert. They generate millions of questions, observations, misunderstandings, experiences and anomalies. No human can inspect them individually.

The expert therefore loses the ability to perceive what may be significant in the experience of the whole audience.

The deeper scaling problem is significance discovery at scale.

Making the whole perceptible

This is one place where AI changes the equation.

AI can perform gainful compression across enormous numbers of human encounters:

100,000 encounters → clusters + patterns + anomalies + changes

It might reveal a new cluster of questions; an explanation that repeatedly generates confusion; an unexpected connection between subjects; a minority whose experience systematically differs from everyone else's; or observations outside an expert's existing conceptual manifold.

These are not automatically significant. They are candidates for significance.

The human encounters the compressed field and says: That's interesting.

At that moment something remarkable has happened. The experiences of thousands of humans have affected one human without requiring thousands of individual conversations.

The expert changes. Their teaching changes. That changed understanding propagates back to the audience. The audience responds. Those encounters are compressed again. New significance becomes perceptible.

*human → many humans → compression → significance → changed human →
many humans*

This is no longer simply publishing. It is a learning system.

But perceptibility is not enough

Now remove the privileged expert.

Suppose many humans encounter different parts of reality:

$H_1, H_2, \dots, H_n$

One encounters x , another y , another z . No individual can perceive $X = f(x, y, z, \dots)$ because no individual has access to all the local encounters.

AI may help make X perceptible. But that does not solve the deeper problem.

An individual can perfectly perceive the state of the whole and still ask: Why should that change what I do?

This is the point at which our information and compression account runs out.

What matters to me?

becoming.

What matters to us?

That transition is not compression. It is not communication. It is not consensus. And it is not simply cooperation.

It requires trust.

Trust allows significance to cross boundaries

To regulate myself with respect to the whole is risky.

My own experience may tell me A, while the experience of others suggests B.

If the larger system is to affect my behaviour, I must sometimes be willing to accept:

I cannot see everything from where I am. Something encountered by another person may be significant even though I have not encountered it myself.

This requires trust.

Trust is necessary precisely because verification is incomplete. If I could independently reproduce every experience, inspect every observation and derive every conclusion myself, I would not need trust.

Human bandwidth makes that impossible.

Trust is the willingness to allow information we cannot completely verify ourselves to participate in our self-regulation.

This is how significance crosses the boundary between autonomous agents.

Something happens to you. It matters to you. I did not experience it. But because I trust you, the fact that it matters to you can itself become significant to me.

significance_A --trust→ significance_B

This does not require agreement.

Trust does not mean: You say this matters, therefore I must agree that it matters.

It means: The fact that this matters to you is significant to me.

I remain free to interpret it differently. I remain responsible for my own judgement. But I have become permeable to significance originating outside my own experience.

We-formation

This may be the beginning of a we.

Initially:

A: I regulate myself according to what matters to me

B: I regulate myself according to what matters to me

The two individuals may communicate or cooperate, but there is no necessary collective intelligence.

Trust creates another possibility:

A ↔ B through trust and significance

A remains A. B remains B. There is no requirement for some new super-entity C to appear above them.

What changes is the relationship. Each becomes capable of regulating themselves partly with respect to realities encountered by the other.

The intelligence begins to exist between them.

We-formation is reciprocal permeability to significance.

This also explains why a we is not binary. There are degrees of permeability.

Something significant to a stranger may have little effect on my behaviour. Something significant to a colleague may have more. Something significant to a close friend, partner or child can immediately reorganise what matters to me.

The strength of a we may therefore be related to the degree to which significance can legitimately cross the boundaries between its members.

But this permeability must remain reciprocal and compatible with autonomy. Otherwise we becomes domination.

Trust is not surrender

This distinction is essential.

Collective intelligence cannot mean: The collective tells me what matters, and I obey.

That is centralized control, not intelligence between humans.

Nor should AI determine: This is what matters to everyone.

The purpose of a trustworthy collective system is not to manufacture significance. It is to make the distributed experience of others sufficiently perceptible that humans can exercise judgement about its significance.

So the requirement is not: Trust the whole.

It is: Trust the process by which the whole becomes perceptible.

That requires provenance.

If a system tells an expert: Thousands of people appear to be struggling with this idea, the expert should be able to ask: Show me.

The compression should permit descent:

field → cluster → encounters → human sources

Compression necessarily hides detail in order to make structure perceptible. Trust requires that the hidden detail has not simply disappeared.

trustworthy compression = compression + provenance + inspectability

We can abstract upward without severing the tether downward.

The whole becoming perceptible to its parts

We can now return to multicellular life.

The extraordinary transition is not simply that cells cooperate. At some point, an individual cell regulates its behaviour with respect to something larger than itself.

Its implicit question changes from:

What matters to me?

toward:

What am I looking after?

The cell does not need an intellectual representation of the organism. But its local environment allows it to become sensitive to states that only make sense with respect to the larger whole.

For humans, the analogous transition may begin when reading ourselves is no longer sufficient for orientation.

Where am I?

What is happening around me?

What are others encountering that I cannot see?

What appears to matter here?

What am I looking after?

Collective intelligence begins when the whole becomes perceptible to its parts, and the state of the whole becomes significant to how those parts orient themselves.

The first half is increasingly an engineering problem. The second half remains profoundly human.

Consensus is not collective intelligence

None of this requires consensus.

Indeed, consensus can be a dangerous form of lossy compression.

If 99% of a group sees X and 1% sees Y, majority aggregation gives X. But the minority observation may contain the most significant signal in the system.

One person may have encountered the thing that proves everyone else's model wrong.

A collectively intelligent system must. Compress the common. Illuminate the anomalous.

But anomaly is not automatically significance. The anomaly must encounter judgement: Does this matter?

This preserves the role of significant surprise at the collective level.

A healthy we is therefore not one in which differences disappear. It is one in which differences can cross boundaries without destroying trust.

From local significance to collective orientation

We can now describe the larger process:

local encounters → significance → gainful compression → perceptible field → trust → shared significance → collective reorientation

Compression makes the experience of many perceptible.

Trust allows significance to cross autonomous boundaries.

Shared significance makes we-formation possible.

Reorientation demonstrates that the collective has actually learned.

And significant surprise keeps the entire system open to reality.

Intelligence between us

This brings us somewhere rather different from the original proposition that intelligence is compression.

Compression remains fundamental. But intelligence requires something to determine what deserves preservation. That something is significance.

Significance compressed through experience becomes meaning. Meaning generates expectation. Reality produces significant surprise. Significant surprise forces meaning to change.

Between autonomous humans, trust allows significance discovered by one person to participate in the orientation of another.

When this becomes reciprocal, a we begins to form. And when that we can perceive and respond to significance unavailable to any individual member, collective intelligence becomes possible.

Meaning = gainful compression of significance

Intelligence = meaning continuously revised through significant surprise

Trust = what allows significance to cross autonomous boundaries

We-formation = reciprocal permeability to significance

Collective intelligence = the capacity of that we to discover and respond to significance that none of its members could perceive alone

The opportunity

AI creates an extraordinary new possibility here.

Not because AI should determine what matters to humans. And not because we should build an artificial collective intelligence above them.

AI can solve a different problem.

Human beings generate far more significant experience than other human beings can possibly encounter directly. Human bandwidth prevents locally discovered significance from becoming globally perceptible.

AI can perform gainful compression across those encounters. It can reveal patterns. Preserve anomalies. Expose significant surprise. Maintain provenance. Allow humans to move from the compressed field back toward the encounters from which it arose.

It can make the whole more perceptible to its parts.

But humans remain responsible for significance.

The model can say. This appears significant.

The human must still ask.

Significant for whom?

Significant with respect to what?

What should this change?

And the human remains answerable to reality.

The objective, then, is not to engineer a collective mind. It is to create the conditions in which humans can form a we without surrendering their autonomy.

A system in which what matters to you can become perceptible to me. In which I can allow your encounter with reality to modify my orientation without surrendering my judgement. In which significant differences are not crushed by consensus. In which the experience of many can become perceptible without becoming centralized authority.

And in which our individual relationships with reality can begin to contribute to an intelligence that exists not above us, but between us.

Perhaps that is the deeper promise of Human Intelligence:

not making the individual less important in service of the whole, but making the relationships between individuals capable of carrying significance at a scale they never could before.

Gainful Compression, Significance and the Formation of We · Michael Kontorovich, 7 September 2026. Published by Polyworld at polyworld.life/thinking/gainful-compression/