

Andrey Vlasov



PREDICTIVE
MINDS
UNDER
PRESSURE

Andrey Vlasov

Predictive Minds Under Pressure

«Издательские решения»

Vlasov A.

Predictive Minds Under Pressure / A. Vlasov — «Издательские решения»,

Predictive Minds Under Pressure is the map of the crumbling frontier. Our consciousness is not a mirror of reality, but a proactive engine of anticipation. But what happens to the self when the “noise” of modern stress overwhelms our internal models? Presented as a sharp philosophical dialogue, this book explores the fragile machinery of consciousness under fire. It moves beyond the mechanics of cognitive rigidity to introduce coping intelligence.

© Vlasov A.

© Издательские решения

Содержание

Predictive Minds Under Pressure: Dialogues In Noise	6
Editor's Preface	7
PREDICTIVE MINDS UNDER PRESSURE: DIALOGUES IN NOISE	8
Chapter I. What Is Consciousness?	8
Chapter II. Neural Foundations of Consciousness	11
Конец ознакомительного фрагмента.	15

Predictive Minds Under Pressure

Andrey Vlasov

Editor A. Vlasov

Cover designer A. Vlasov

© Andrey Vlasov, 2026

© A. Vlasov, cover design, 2026

ISBN 978-5-0069-8817-0

Created with Ridero smart publishing system

Predictive Minds Under Pressure: Dialogues In Noise
*A philosophical dialogue on consciousness,
stress, and coping intelligence.*

For those who think under pressure.

For those who teach under noise.

For those who refuse to stop asking

why there is something it is like to be here at all.

Editor's Preface

This book did not start as a book. It started as a problem.

The original plan was narrow and technical: to integrate three very different ways of talking about consciousness into one coherent framework. On one side stood contemporary neuroscience, with its neural correlates and global workspace models. On another, the Soviet and Russian school, with Vygotsky, Leontiev and Luria treating consciousness as socially shaped activity. On a third, computational theorists with predictive coding, the Free Energy Principle, and modern variants of Integrated Information Theory.

The working group brought these traditions into structured conversation. Over many iterations, they produced a shared ontology, a set of axioms, and a map of where the schools converge and where they disagree. They also produced something less formal, but just as important: a style of talking to each other that combined rigor, patience, and occasional sarcasm.

At some point, it became clear that the **dialogue** itself was the best educational tool. The deepest insights did not appear in tidy definitions. They appeared when someone said «Wait, that does not follow,» or «Let me try a different example,» or «I know what my equations say, but my own mind does not behave like that at 3 a.m.»

That is why this book adopts the form of a continuous dialogue between a Teacher and a Student. The Teacher is not always right. The Student is not always wrong. Both carry fragments of the three traditions, and both have to make sense of them while living in a noisy world.

Within this conversation, the Free Energy Principle (FEP) and Integrated Information Theory (IIT) are treated as complementary perspectives. FEP describes brains as systems that reduce surprise by updating predictive models. IIT provides a way to talk about how much information is integrated in a system. In the background, a more general idea appears: that conscious systems have a kind of *computational mass index* — a structured capacity to carry, integrate, and update patterns. FEP and IIT are viewed here as partial, high-level approximations to that deeper concept.

The intended readers are students and researchers of mind — broadly defined — between roughly 18 and 45, many with philosophical or technical training. They are comfortable with ideas, but tired of inflated jargon. For that reason, the prose is direct. Sentences are short. Equations stay off-stage. A bit of dry humor and self-irony is allowed, because minds under pressure rarely speak in a perfectly formal tone.

If this book works, it will not give you the final theory of consciousness. It will give you a sharper inner vocabulary for what your own mind is doing when the noise level rises.

A note on the book's structure. The book moves in three arcs. The first (Chapters I—III) asks what consciousness is and what the brain does to produce it. The second (Chapter IV) asks how consciousness is shaped by activity, body, and social life. The third (Chapters V—VII) asks what happens when this whole system is put under pressure — and what it means to cope intelligently. Each chapter is a set of dialogues. Each dialogue is self-contained enough to read on its own. The Appendices at the end provide a glossary, a convergence map, the shared axioms, and a self-assessment. You can read those first, last, or in parallel — depending on how your own mind likes to navigate.

PREDICTIVE MINDS UNDER PRESSURE: DIALOGUES IN NOISE

Chapter I. What Is Consciousness?

This chapter is your first map of the territory: what philosophers and neuroscientists mean when they say «**consciousness**», why there is a «hard problem», and how your own sunset experiences fit into this story.

Read it as a way to sharpen your inner vocabulary, not as a final answer — we are building tools you will reuse when we talk about stress and coping later.

#consciousness #qualia #hardproblem #access_vs_phenomenal #aboutness

Dialogue 1. The Sunset Question

Teacher: Before we start, I want to ask you something very simple. Not philosophical. Just concrete. What do you see when you look out of the window right now?

Student: The sky. An orange-pink sunset. Clouds. Something like... calm. It is hard to put into words.

Teacher: «Hard to put into words» is exactly where we begin — and it is a better starting point than most textbooks would admit. You named the objects easily: sky, sunset, clouds. But the moment you reached the **experience** — calm, beauty, that vague something — language started to resist. That gap is worth noticing.

Student: Maybe because experience is inside, and words are outside?

Teacher: That is one of the central distinctions in the philosophy of mind — and you landed on it without a map. Let us draw it more carefully.

When you look at the sunset, two different events happen at once.

First: your brain registers light in a certain wavelength band, your visual cortex classifies it as «orange», and this information becomes available to your cognitive systems. You can say «I see orange», decide to take a picture, describe what you see, store it in memory. This is what we call **access consciousness**: information that is globally available for report, reasoning, and control of action.

Student: So, there are two things happening: the information part, and the... feeling part?

Teacher: Exactly. And philosophers, characteristically, gave the second part a name that sounds more mysterious than it needs to: **qualia** — the raw qualitative feel of experience.

The redness of red. The sharpness of pain. The warmth of an orange sunset. Qualia are what phenomenal consciousness is made of, as far as we can tell.

Student: And what is the real difference? Isn't «knowing that you see orange» the same as actually seeing it?

Teacher: This is where things become tricky. Let me tell you a thought experiment by Frank Jackson.

Imagine a scientist named Mary. She is a world-class expert on the neuroscience of color. She knows, in principle, **everything** about color vision: which photoreceptors respond to which wavelengths, how signals travel through the thalamus and cortex, how the brain encodes «red». She can simulate the entire system.

But Mary has lived her whole life in a black-and-white room. She has never actually seen color.

Student: So, what happens when she leaves the room?

Teacher: She sees a red rose for the first time — and she learns something new. Even though she already knew every physical and neural fact about color, she did not know what it is like to see red.

That extra piece — which no equation gave her — is what we call qualia. The gap between complete objective knowledge and that missing subjective «what-it-is-like» is what David Chalmers named the hard problem of consciousness. Part of this book is about giving you more precise names for these pieces — not to kill the magic, but to let you see it more clearly when it happens to you.

Dialogue 2. The Hard Problem

Student: Can you state the hard problem in one sentence?

Teacher: I can try — though it is the kind of sentence that quietly expands in your head after you first read it: even with a complete description of all neural processes in the brain, we still have to explain why there is any subjective experience at all, instead of just silent computation.

This is often called the **explanatory gap**. On one side we have functional stories about what the brain does. On the other side we have the simple fact that there is an inner movie — or at least a faint inner radio — attached to those processes.

Student: That sounds like something science cannot quite touch yet — or at least not with its usual tools.

Teacher: It is a philosophical problem, not a standard empirical question. But that does not mean science has to stop and wait — it just means we use a slightly different set of tools on this particular page.

In practice, neuroscience follows a **correlational strategy**. It looks for **neural correlates of consciousness** (NCC): minimal sets of neural mechanisms that are necessary and sufficient for particular experiences. You can design experiments, record activity, and refine theories. This work does not close the explanatory gap, but it does not depend on closing it.

Student: Are there ways philosophers try to get around the gap?

Teacher: There are several — and none of them has won yet, which is itself informative. One is **illusionism** (Daniel Dennett, Keith Frankish). It says that qualia, as people usually talk about them, are a kind of introspective illusion. Our cognitive systems generate the *impression* of a special inner qualitative stuff. The hard problem looks hard because that impression is misleading.

Another is **panpsychism** (Philip Goff, Chalmers in some readings). Here, experience is treated as a **fundamental property** of reality, like charge or spin. Even simple systems have primitive proto-experiences. Human consciousness is a highly structured form of something that exists in a minimal way everywhere.

A third is **neutral monism** (Bertrand Russell, Ernst Mach). In this view, neither «mental» nor «physical» is basic. Both are different organizations of a deeper, neutral kind of stuff.

None of these views has become the consensus. The gap is stubborn, not merely a matter of taste. This is not a reason for despair. It is a sign that the question is genuine.

Dialogue 3. Subjectivity and Aboutness

Student: So, what are the basic features that all conscious experiences share?

Teacher: Two stand out — and once you see them, you will notice them everywhere.

First is **subjectivity**. Each experience belongs to someone in the first person. *My* pain is mine, not yours, and not «pain in general.» This «mine-ness» is hard to capture from a third-person point of view.

Second is **intentionality** — in the technical sense from Franz Brentano, not just «doing something on purpose». Conscious states are always **about** something. You do not simply see; you see *a sunset*. You do not simply fear; you fear *something*. Mental states have built-in **aboutness**.

Student: So, consciousness never just floats in a void. It is always tied to a world.

Teacher: Exactly — and this is one of those moments where neurobiology, phenomenology, and psychology stop arguing and quietly point in the same direction.

The Soviet psychologist **Alexei Leontiev** introduced the idea of an «**image of the world**». For him, a person's consciousness is not a collection of isolated snapshots. It is a **structured system of meanings and personal senses** in which the subject lives and acts. In this picture, qualia are not little atoms of experience. They are **nodes in a semantic network** that connects perception, action, and personal meaning.

An orange sunset is not only a wavelength. It is «the sunset after a hard day», «the memory of a specific place», «the sign that it is finally over». Meaning is not added on top of experience. It is woven into it from the start.

Dialogue 4. The Tip of the Iceberg

Student: Is everything in the brain conscious in this rich sense?

Teacher: Not even close. The brain is running an enormous amount of processing that never shows up in consciousness at all.

Autonomic regulation of heart rate, low-level sensory analysis, habitual motor sequences, implicit learning, most emotional pre-processing — all of this runs silently and competently **without** access consciousness. The iceberg metaphor is apt: conscious experience is a small and special subset of what the brain does. Most of the iceberg is doing just fine below the surface.

But consciousness is not just «the visible tip». It is a **distinct mode of processing** with several features:

- Information is globally broadcast across distant brain networks.
- Behavior can break out of rigid habits and become flexible.
- The system can form explicit verbal reports and reflect on its own states.
- The processing is accompanied by a subjective «what-it-is-like.»

Unconscious processes can be highly sophisticated. What they lack is not complexity, but this special mode of global availability and subjective presence.

Teacher: In the next chapter we ask a simple, hard question: if consciousness is not only neural activity, it is still **not less than** what the brain does. So what exactly does the brain do when a state becomes conscious?

Student: We dive into the physics of the inner movie.

Teacher: Yes — with one agreement: we use the equations as tools, not as decorations or as gods.

Chapter II. Neural Foundations of Consciousness

Here we ask what the brain is actually doing when an experience becomes conscious: where the «inner movie» is implemented, how signals become significant, and why timing matters.

If you ever wanted to know what NCC, gamma synchrony and predictive coding have to do with your everyday awareness, this is where the pieces start to line up.

#NCC #gamma #thalamocortical_loop #functional_systems #significance_detector
#predictive_coding

Dialogue 5. Neural Correlates: Where in the Brain Is the «Movie»?

Teacher: Let us move from sunsets to neurons — which is a bigger conceptual leap than it sounds. If consciousness is not less than what the brain does, we should ask the rudest possible question: where in the brain is the «inner movie» implemented?

Student: You mean like «this spot in the cortex = consciousness»?

Teacher: That is the tempting version. It is also wrong — which, to be fair, is a useful thing to know early. Francis Crick — one of the co-discoverers of DNA — spent the last decades of his life trying to make this question precise. Together with Christof Koch, he proposed a more careful program: look for **neural correlates of consciousness**, the so-called NCC.

Student: What exactly counts as an NCC?

Teacher: A standard definition is this: an NCC is the **minimal set of neural events** that are **jointly sufficient** for a specific conscious experience, given that the rest of the brain is functioning normally.

«Minimal» matters. Every experience involves the whole brain in some way, but that is not informative. We do not want a vague statement like «consciousness happens when the brain is active». We want to know which specific patterns make the difference between a perception being conscious and being unconscious.

Student: How do we even isolate such patterns?

Teacher: By finding situations where the **stimulus is the same**, but **experience differs** — which is a cleaner experiment than it sounds. Think of ambiguous figures, like the Necker cube. The image on your retina does not change, but your experience alternates between two interpretations. When that switch happens, something in the brain also switches.

Crick and Koch looked at such phenomena in vision. They found that simple activation in early visual cortex is often **not enough**. The crucial change appears in patterns of **synchrony** and long-range communication between areas. This led to a hypothesis that **gamma-band synchrony** — oscillations around 40 Hz — plays a role as a temporal «glue» that binds distributed neural assemblies into a single conscious content.

Student: So, consciousness is a kind of synchronized party in the brain?

Teacher: That is a picture worth keeping. If the party is too local — if the rooms are not talking to each other — you get rich processing, but no unified conscious scene. When local groups start to synchronize and broadcast information widely, a content enters the spotlight of consciousness.

This brings us to another key player: the **thalamocortical loop**. The thalamus is not just a relay station. It is part of a recurrent loop with the cortex, and this loop seems necessary for maintaining conscious states. Under general anesthesia or in deep coma, this loop is disrupted. Local cortical activity may continue, but the global dynamics changes and conscious experience fades.

Student: So the problem is not that the brain is quiet, but that the guests are all talking in separate kitchens.

Teacher: Exactly. Consciousness is what happens when someone opens the doors — and the conversation becomes one.

Dialogue 6. The Soviet and Russian Contribution: Functional Systems and Significance Detectors

Student: You mentioned earlier that the Soviet tradition has something important to add here. How does that fit into the picture?

Teacher: More naturally than you might expect — and more historically honestly than most neuroscience textbooks acknowledge. Let us take two examples: **functional systems** and **significance detectors**.

First, **Pyotr Anokhin** and his **Theory of Functional Systems**. Starting from the 1930s and through the 1960s, he described behavior not as a simple reflex chain, but as a dynamic system centered around an «**acceptor of the result of action**».

Before an action is executed, the brain forms an **anticipatory model** of the expected outcome. Sensory feedback is then compared to this model. If they match, the system confirms the action and moves on. If they do not, it triggers corrections. In modern language, this looks very much like a **predictive model plus error correction loop**.

Student: That sounds almost like predictive coding.

Teacher: Almost is an understatement. In our ontology, we treat Anokhin's functional system as **structurally** homologous to Friston's Free Energy framework. The languages differ — one is physiological, the other is mathematical — but the core move is identical: anticipate, compare, correct. The nervous system constantly anticipates future states and uses mismatches to drive learning and control.

Now, the second example: **Natalia Bekhtereva** and the «**detector of significance**.» She worked on what we would now call the neural basis of relevance and salience. Her experiments showed that certain distributed networks — involving limbic structures, basal ganglia, and frontal cortex — act as a system that marks incoming signals as **subjectively significant** or **insignificant**.

Student: So, this detector decides what gets into consciousness?

Teacher: In a way, yes — though it is less like a single bouncer at the door and more like a distributed committee. No single neuron shouts «important!»; a whole network biases which signals are allowed to reach the **global workspace**.

The modern language of **salience maps** and **attentional priority** is a direct descendant of this idea. When we ask «what gets on the stage of consciousness,» we are really asking how the brain's significance detectors and salience networks work.

In honest historical terms, Bekhtereva's lab work anticipated some of the core insights of current computational salience models — decades before the computational vocabulary existed. Our ontology simply names that continuity rather than pretending it is a coincidence.

Dialogue 7. Neural Dynamics: Consciousness as a Process, not a Snapshot

Student: So far, we have talked in terms of structures: cortex, thalamus, networks. But earlier you said consciousness is not a static pattern, but a process. Can you unpack that?

Teacher: This is where **neural dynamics** comes in — and this is one more place where Russian science got there earlier than the international field quite noticed.

Alexander Ivanitsky and his collaborators studied the **time course** of conscious perception. They showed that when a conscious experience forms, activity does not just «appear» at one moment. It evolves across time in a characteristic sequence.

A simple sketch looks like this:

- Early sensory areas respond to physical features of the stimulus.
- Associative areas integrate features into more complex patterns.
- Frontal and parietal areas integrate these patterns with context, goals, and meaning.

Conscious experience, in this view, is tied to the **final integration phase** — when sensation and meaning are woven together into a coherent, reportable state.

Student: Is this related to the idea of a «window of the present»?

Teacher: Yes — and this is one of those moments where a phenomenologist from 1905 and a neuroscientist from 2005 arrive at the same description from completely opposite directions. Empirically, this window seems to be on the order of a few seconds.

Neural dynamics gives this an implementation story: it takes roughly that long for the brain to complete one cycle of integration — from raw input, through intermediate processing, to a globally integrated, meaningful state. Consciousness is not a single frame. It is more like a short clip.

This time-based view will matter a great deal when we talk about stress. Under pressure, the integration window can be rushed, distorted, or flooded. But that is the next arc of the book — let us not get there before we have the tools.

Dialogue 8. Free Energy and Predictive Coding

Teacher: Now we can turn to a theory that has become genuinely central in computational neuroscience — and also, it turns out, a surprisingly good description of what it feels like to be wrong in a stressful situation. This is the **Free Energy Principle** (FEP) and its implementation through **predictive coding**.

The core claim is simple and radical at once: a brain can be modeled as a system that **minimizes a quantity related to surprise** about its sensory inputs.

Student: So, the brain is basically a Bayesian nerd?

Teacher: That is one good way to say it — the brain might object to the word «nerd», but the equations would quietly agree. A Bayesian reasoner constantly updates beliefs based on new evidence, always asking: given what I am sensing, what is most probably causing it, and what should I predict next? In predictive coding schemes, higher levels of the cortex send **predictions** downward; lower levels send back **prediction errors** — the gap between predicted and actual input. The system continuously adjusts its internal model to reduce these errors.

In this picture, perception is not passive reception. It is **active inference**: the brain is always guessing what is out there and checking those guesses against incoming data.

Student: And conscious experience appears when... what exactly happens?

Teacher: One intuitive proposal is that a state becomes conscious when **prediction errors at relatively high levels** of the hierarchy cannot be quickly suppressed — the model has to reorganize itself significantly. During that reorganization, the content becomes globally available and experienced.

Notice the thread back to Anokhin: his «acceptor of the result of action» is exactly a predecessor of the generative model. The mismatch he described is what we now call prediction error. The terminology changed; the insight stayed.

Correction of behavior is model updating. The Free Energy framework generalizes this logic to perception, action, and learning in a unified way.

In the ontology of this book, FEP is one of the main tools for describing how a **predictive mind** stays coherent in a noisy world — or fails to.

Dialogue 9. Integrated Information and Computational Mass

Student: Where does Integrated Information Theory fit into this? You said we would not treat it as a rival dogma.

Teacher: IIT comes at the problem from the opposite direction — which is exactly what makes it a useful companion to FEP, not a rival. Instead of starting with neurons and computation, it starts with the **structure of experience itself** and asks: what must any system have if it is going to have experience at all?

Giulio Tononi proposed a set of axioms: experience exists, it is structured, specific, integrated, and exclusive. From these, he constructed a formal quantity, Φ (**phi**), which measures how much information a system carries **as a whole**, over and above the sum of its parts. A system with high Φ is said to have a rich, integrated state.

In this book, we do not take the slogan «consciousness is Φ » literally — it is a tempting slogan but an overcommitment. Instead, we treat Φ as a candidate for a **computational mass index**: a way to ask how much structured, integrated «stuff» a system can carry in principle.

Student: And how does that relate to FEP?

Teacher: FEP tells us how systems maintain themselves over time by updating models and reducing surprise. IIT tells us something about **how densely structured** their internal states can be. They point to different, but potentially complementary aspects of conscious systems.

You could say, playfully, that FEP describes how a system **moves** through its state space, while IIT hints at how «heavy» or «thick» any given state is — how much is packed inside the moment. One tracks the journey; the other tracks the density of the traveller.

The deeper theory — if it ever arrives — might unify these. For now, we use both as tools. Neither breaks if you drop the other.

Dialogue 10. A Brief Truce

Student: So, we have gamma synchrony, thalamocortical loops, functional systems, significance detectors, predictive coding, and integrated information. That is a lot of moving parts.

Teacher: It is — and if your head is slightly overloaded right now, that is a sign the material has actually been landing. But notice a pattern before we move on.

— Neural correlates (NCC) tell us **where** and **when** conscious contents appear in the brain.

— Functional systems and predictive coding tell us **how** the brain anticipates and corrects its own states.

— Significance detectors and global workspaces tell us **which** contents make it into the spotlight.

— IIT and related measures hint at **how structured and unified** those states can be.

We do not need to collapse this into one slogan — and we should probably be suspicious of anyone who tries. For our purposes, it is enough to treat these as parts of a single **conceptual toolbox**, each with its own job.

In the next chapter, we will leave the brain for a moment and look at consciousness from another angle: as **activity** — as something that is produced in doing, in social interaction, and in meaning. This is where the Soviet tradition and modern enactive approaches have a lot to say.

Конец ознакомительного фрагмента.

Текст предоставлен ООО «Литрес».

Прочитайте эту книгу целиком, [купив полную легальную версию](#) на Литрес.

Безопасно оплатить книгу можно банковской картой Visa, MasterCard, Maestro, со счета мобильного телефона, с платежного терминала, в салоне МТС или Связной, через PayPal, WebMoney, Яндекс.Деньги, QIWI Кошелек, бонусными картами или другим удобным Вам способом.