AlpiType · Anton Lytvynenko · 2026

Stack Question 1: Does every AI system have an explicit EU AI Act classification? Documented and signed by a responsible person? 12 Anthropology of AI Systems AlpiType · Anton Lytvynenko Question 2: Does your cyber insurance cover AI decision damages? In most policies from 2023 and earlier — no. Question 3: What is your 72-hour action sequence in the event of an AI incident? If you have no clear answer — you have no compliance architecture. 13 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 5 Economic Anthropology Companies pay cloud providers six-figure sums and AI operators five-figure ones. The math doesn't add up. Market Observation: Salary vs. Cloud Costs Industrial plant 200–500 employees: cloud spending on AI 80,000–250,000 euros/year. AI operator salary: 45,000–75,000 euros gross. Senior: up to 95,000 euros. For comparison: Senior Software Architect 95,000–130,000 euros, ML Engineer 85,000–110,000 euros, SAP Architect 100,000–140,000 euros. The AI operator performs work that requires the competencies of at least three of these specialist fields. And is paid like junior IT support. What an AI Operator Actually Does

- Prompt architecture — deep understanding of model architecture, avoidance of hallucinations.

- Data curation — first feedback line between real-world deployment and training.

- Escalation mediation — normal variation or genuine anomaly? Both error directions carry costs.

- Compliance translation — translating EU AI Act, GDPR into concrete operating procedures.

Consequences: Turnover, Knowledge Loss, Stagnation The most common pattern: 12–18 months of enthusiasm, then stagnation. The system freezes on one version. Two years after go-live the model is outdated — but continues running on the initial momentum, consuming cloud budget. Three Questions for Your Own Stack Question 1: What is your cloud cost / operator salary ratio? Total annual cloud spending on AI vs. salaries of the people maintaining these systems. Question 2: What is lost if your AI operator resigns? Can you describe this concretely? Is this knowledge documented? Question 3: What career path do you offer a good AI operator over 3–5 years? If the answer is unclear — you have a retention problem. 14 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 6 Escalation Risk with GPU Four anonymized cases in which unclear roles caused more damage than technical failures. Introduction: Where Systems Really Break The four cases described are anonymized, but real. They share one thing: the technical part functioned within its specifications. The problem lay in the anthropological layer. Case A: Mechanical Engineering, Predictive Maintenance Upper Bavaria, 320 employees. Friday evening 9:30 PM: system reports "Critical" — predicted spindle malfunction within 12–36 hours, probability 91%. Potential damage: 45,000 euros. After one hour of discussion: decision against shutdown. The machine stopped by itself the following morning. Damage: 78,000 euros plus downtime. The model responded correctly. The problem: nobody had defined what "Critical" means in terms of action, who decides, according to which procedure. Case B: Logistics, Route Optimization Rhine-Main, 85 trucks. Systematically late deliveries — 22 minutes on average. Cause: a regional transport strike day was unmarked in the training data. 9 months of systematic error, contractual penalties, complete retraining. Case C: Chemicals, Quality Control Baden-Württemberg, AI system accuracy 94%. External ISO 9001 audit: auditor cannot obtain interpretation justification — cancels the approval. Lab re-measurement confirms the system. 4-day delay, 12,000 euros contractual penalty. Case D: Defense, Sensor Fusion Bavaria, "EU Sovereign Cloud" of an American provider. Geopolitical tension: provider restricts inference API due to ITAR. System: from 200 ms to 8–12 seconds — effectively non-functional. Pattern Summary

- Case A: Missing mandate for nighttime decision — control role vacant.

- Case B: No data curator — training role failed to fulfill editorial duty.

- Case C: No compliance translation mechanism for external auditors.

- Case D: No sovereignty audit — liability role failed to identify vendor risk.

15 Anthropology of AI Systems AlpiType · Anton Lytvynenko Three Questions for Your Own Stack Question 1: Describe the worst-case escalation path. Who decides, with what mandate, within what timeframe. If you cannot put this in writing within 5 minutes — the path is not defined. Question 2: Has anyone explicitly searched for anomalous events in your training data? Who made the last three significant dataset updates? Question 3: Under what conditions can any external provider restrict your access? For each critical provider: what is the emergency operation? 16 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 7 70/30 as an Architectural Principle How an AI system is built in which the 70% machine / 30% human split represents not a compromise but a planned architecture. The Layer Model: Technical and Social Layer The 70/30 architecture requires two diagrams: one for the technical layer (model, pipeline, infrastructure, logging, monitoring) and one for the social layer (roles, escalation paths, mandates, thresholds). Both layers are equal system components. One without the other is not an AI system. EXPECTED vs. REAL CONTROL ▲ Expectation layer Expert reads report · follows the reasoning · recognizes errors Assumption: human control detects AI errors ▼ GAP ▼ Summary replaces trace · error remains invisible ▼ Reality layer Summary instead of trace · plausible text conceals errors Reality: control only works with transparent output Fig. 2: Expected vs. Real Control — the gap between assumption and reality The diagram shows the central problem of the social layer: the expectation that a human "reads and understands" an AI report meets the reality that the model outputs a summary rather than a reasoning trace. The gap is not a misunderstanding — it is a design flaw. The Eight-Step Process Step 1: Process Mapping The first step consists of capturing the process, not the model. Every step, every decision, and every transition is recorded. Step 2: Identify Decision Points For each point, four questions are examined: error costs, regulatory burden, frequency, and variability. Step 3: Define the 70% High frequency, low variability, acceptable error costs, no requirement for human signature. Step 4: Define the 30% The second half of the system is defined via named individuals with documented substitutes for absence, vacation, or changed working-time arrangements — not functional roles. 17 Anthropology of AI Systems AlpiType · Anton Lytvynenko The concrete name, position, and the person who steps in as substitute are recorded. Step 5: Set Thresholds Thresholds are defined specifically and measurably and recorded both in the code and in the operating documentation. Step 6: Logging Design Every decision, every escalation, and every override is logged with name, timestamp, and justification. For high-risk systems, the EU AI Act prescribes a retention period of five years. Step 7: Operating Documentation Prior to go-live, the system description, operating procedures, and disaster recovery plan are available in signed form. Step 8: Escalation Drill Prior to go-live, a simulated incident is run through in real time. If the escalation path takes longer than the maximum time specified in the operating documentation, it must be revised. Typical Errors

- Mandate concentration. All roles with one person — single point of failure.

- Technical thresholds without human context. Confidence values without explanation of what they mean in practice.

- Documentation nobody reads. In the SharePoint folder, never opened since go-live.

- Governance without schedule. "Reviewed as needed" — means: not reviewed.

- No handover plan during personnel change. Knowledge transfer must be a planned system component.

Three Questions for Your Own Stack Question 1: Draw your AI system in two layers. If you cannot draw the social layer — it does not exist in a designed state. Question 2: When were the operating thresholds last reviewed? Who was present? What changed? Question 3: Conduct an escalation drill now. Choose a scenario, measure the time, record what took the longest. 18 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 8 Sovereignty as an Anthropological Question Cloud or local — that is not a technical decision. It is the question of who shuts down the system in an emergency — and what happens afterwards. The Wrong Choice: Cloud vs. On-Premises Reconsidered The "cloud vs. on-premises" debate is usually conducted in technical terms: latency, scalability, total cost of ownership, security. These dimensions are real — but not central. The central question is operational and strategic in nature: On whom does the continuation of your systems depend? Who can — at any time, for any reason — restrict or terminate access to a critical component? That is a sovereignty question. And sovereignty is anthropological, because the answer always points to concrete people and organizations: who holds the key — and who makes the decision to use it? The Kill-Switch Risk: Official Sources These risks are not theoretical. Three institutions have documented them:

- BSI C3A (April 2026): The Federal Office for Information Security establishes that most cloud solutions marketed as "sovereign" do not meet the criteria of genuine operational independence. BSI defines sovereignty through three vectors: technical portability, legal independence, and operational control.

- EP A10-0107/2025: The European Parliament notes over 80% digital dependency on non-European providers — for processors, operating systems, cloud platforms, and AI models.

- NATO Parliamentary Assembly (November 2024): Rapporteur Sven Clement documents risks of cloud-dependent AI systems in scenarios where providers are subject to legal or commercial restrictions — particularly critical in conflict situations.

Sovereign-Washing: Why "European Cloud" Rarely Equals Genuine Sovereignty "EU Data Residency" means that data is physically stored in the EU — not that the administering company is not subject to the law of a third state. American providers are subject to the Cloud Act of 2018, which grants US authorities data access regardless of the physical server location. "GDPR-certified" means data protection standards — not the absence of operational dependencies. The practical sovereignty test: Could you, tomorrow, if the provider blocks access, operate functionally — regardless of marketing labels? Local Does Not Mean Offline 19 Anthropology of AI Systems AlpiType · Anton Lytvynenko On-premises in the modern context does not mean an isolated server in the basement. It means that the control points of the system lie within your organizational jurisdiction: the model is with you, inference takes place on your hardware, your data does not leave your perimeter, model updates occur according to your decision and your schedule. In industrial environments in Bavaria or Austria — where network connectivity on the production floor may be restricted for security reasons — this is not a comfort option. It is a reliability requirement. A predictive maintenance system that cannot tolerate a network outage requires local inference as a foundational architecture decision. Three Questions for Your Own Stack Question 1: Kill-Switch Clause Does every contract with an AI provider contain a clause describing under what conditions the provider can restrict access — and what documented emergency plan your organization has? Question 2: Jurisdiction In which jurisdiction is the company registered that manages your critical AI infrastructure? Are there parent companies outside the EU? What legal mechanisms (FISA, Cloud Act) could affect your operational availability? Question 3: Degraded Mode Describe concretely what happens to your operations if the most critical AI system is unavailable for 4 hours, 24 hours, 72 hours. Is this mode documented? Is the personnel trained on it? 20 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 9 What Changes When AI Operators Receive Senior Salaries A thought experiment with consequences: adequate compensation of the 30% layer changes org charts, vendor selection, and time-to-value. The Thought Experiment Imagine: the role of AI operator is elevated to the salary level of a Senior Architect. Not because the HR department is suddenly generous, but because an analysis has shown: this role actually requires Senior Architect competencies — and the consequences of inadequate staffing already cost more than the salary difference. What changes? More than expected. A change in salary structure is catalytic — it influences recruiting, candidate quality, internal negotiating position, and the way vendors interact with the role. Position Architecture: AI Operator as Horizontal Senior Track AI operator fits into no existing vertical career line. The person who wants to grow in this role has no clear path: should they become a data scientist? ML engineer? Move into management? That is part of the problem. The solution: a horizontal senior track running parallel to existing career paths. Concretely this means: a clear job title with seniority level (e.g. "Senior AI Operator" or "Principal AI Systems Architect"), reporting line close to operational leadership (COO or CTO direct), compensation band 95,000–125,000 euros gross in DACH, and performance metrics measuring system accuracy over time, escalation quality, and compliance status — not document volume. Recruiting: Where These People Come From Anyone seeking an AI operator through the standard recruitment process will receive the wrong candidates — because the role rarely bears that name. The right profiles sit under other titles: ML engineers with an operations bias, senior DevOps with AI experience, business analysts with technical backgrounds, operations managers who have experienced AI rollouts. In the interview: no abstract "tell me about your ML experience." Instead, a concrete scenario: "Our quality classification model has been showing more false positives for a week. What do you do?" Vendor Relationships: A Senior Operator Negotiates Differently AI vendors negotiate differently when there is someone with subject-matter depth on the other side. A competent AI operator can formulate SLAs specifically — not "99.9% uptime," but "inference latency below 200 ms for batch sizes up to 500 units at our hardware configuration." 21 Anthropology of AI Systems AlpiType · Anton Lytvynenko They can negotiate model versioning terms (90 days

advance notice before updates) and demand exit rights (data portability, architecture documentation). ROI Effect: Calculating Total Cost of Ownership Fully The real TCO contains hidden costs that are rarely calculated: degradation costs (without active monitoring, model accuracy predictably declines), turnover costs (a complete turnover cycle costs 1.5–2 annual salaries of the role), compliance costs (EU AI Act fines for high-risk systems start at 300,000 euros), and negotiation value (a competent negotiator saves 10–20% of annual contract costs). When these costs appear in the TCO, the difference between "cheap AI operator" and "properly paid AI operator" disappears — or reverses. Three Questions for Your Own Stack Question 1: Real TCO Calculate the actual TCO of your largest AI system — including degradation costs, turnover risk, and compliance exposure. Compare with the salary difference between current AI operator and Senior Architect level. Question 2: Negotiation Representation Who represented your organization in your last vendor negotiation cycle? Did this person have the technical competence for model-specific SLA and versioning negotiations? Question 3: Job Description Write a job description for "Senior AI Operator" right now. If you find this difficult — you do not yet have a sufficiently clear picture of the role to staff it effectively. 22 Anthropology of AI Systems AlpiType · Anton Lytvynenko Chapter 10 Anthropology of the Next Decade What the DACH industry and European defense must build over the next ten years — and why the anthropological question comes at the beginning, not at the end. Three Scenarios for 2030–2035 Scenario 1 — Cloud Dependency: The DACH industry continues moving toward cloud-based AI. By 2030, critical operating systems depend on two to three major cloud providers, none of them European. Efficient in geopolitically calm times. In times of crisis: the same vulnerabilities as in Case D from Chapter 6 — but at industrial scale. Scenario 2 — National Fragmentation: In response to dependency risks, each country or sector builds its own isolated AI infrastructure. The problem: fragmentation without coordination means cost duplication, lack of scale, and — paradoxically — greater vulnerability. Because smaller isolated systems have lower operational maturity and can offer less resistance to attacks that do not stop at national borders. Scenario 3 — Anthropological Maturity: DACH industry and the European defense sector develop a shared architectural language for the human layer. Not a unified technical platform — but shared standards for roles, mandates, escalation paths, documentation, and sovereignty requirements. AI operator as a recognized professional role in the labor market, analogous to the data protection officer under GDPR. What Sovereignty Means Institutionally in the AI Era Sovereignty in the AI era is not the question of where servers are physically located. It concerns who makes decisions — at three levels:

- First order: which model, on which data, for which purposes.

- Second order: how the model is updated, who reviews updates, who can reject them.

- Third order: who shuts down the system when it behaves unacceptably — and who is liable for it.

An organization that has no control over third-order decisions is not sovereign in this system — regardless of who signed the contract. For a plant manager, sovereignty means: you can shut down any AI component in the production process without approval from an external provider — and the process continues in a degraded but functional mode. Call to Action: Three Decisions That Must Be Made Now

- Plant managers and operations managers: 30-day audit of the human layer (Appendix A) for every AI system in the

Production environment. Result: a list of gaps in training, control, and accountability. That is your action plan.

CTOs and technical directors: Include the social layer as a mandatory component in every new AI project plan. No go-live without signed role documentation and escalation paths.

23 Anthropology of AI Systems AlpiType · Anton Lytvynenko

Compliance officers and executive management: Review every active vendor contract against the three sovereignty questions from Chapter 8. Contracts with unsatisfactory answers are candidates for renegotiation.

Conclusion

This book has not described a single ML algorithm. Made no concrete model or platform recommendations. Given no Python tutorial. All of that exists — and is accessible. There are more qualified data scientists in the DACH market than ever before. What is missing — and what this book attempts to fill — is a language for the conversation about the human layer. Build the human layer first. The technical layer will find its place.

Chapter 11

Harari Was Wrong — No Neuralink Required

Cognitive augmentation of humans does not happen through physical implants, but through the quality of thinking that flows into the prompt. The interface already exists — it is called language.

Where Harari Failed to Think It Through

Yuval Noah Harari wrote Homo Deus in 2015. His central thesis: to remain relevant in an era where algorithms outperform humans at cognitive tasks, humans must connect with technology through physical interfaces — neural implants, a direct channel between cortex and processor. Elon Musk launched Neuralink. Transhumanists reached the intellectual mainstream.

But look at what actually happened between 2022 and 2026: The cognitive capabilities of an average technician at a Bavarian manufacturing company — no transhumanist, no researcher — have grown radically. He can analyze a problem in one hour that previously required a consulting team three weeks. He can formulate a legal argument without a law degree. Without Neuralink. Without surgical intervention. Only with the ability to formulate a request.

The Prompt as a Neural Implant Without Surgery

A neural implant is a hardware interface: an electrode reads a neural signal, converts it into a digital data stream, and passes it to an external processor. Goal: to reduce the latency between thought and machine response.

A prompt is the same thing at the software level: the human formulates a thought in language, passes it to the language model, and receives a computed, structured response. Latency: a few seconds. Operation: none.

The hardware level requires surgical procedures, regulatory approvals, decades of clinical trials. The software level requires: the ability to think clearly enough to articulate what one wants. This makes augmentation democratic rather than elitist.

What Has Changed in Humans — and What Has Not

What has not changed: the computational power of the human brain. Neurons do not fire faster. Working memory has not grown larger.

What has changed: access to an external computational resource that can be activated through natural language. Large language models have removed the filter that always stood between request and response — whether time, format, or the need to first learn the language of the system. The model meets the human where they are: at the level of human language, not machine language.

Why Harari Thought a Chip Was Necessary

Harari thinks in categories of evolutionary biology and cultural anthropology — both disciplines operate in millennia. From this perspective, natural language as a technology had long stagnated: it was no better as a

human-machine interface in 1970, 1990, or 2010 than before. So the next level lay at the hardware layer.

25 Anthropology of AI Systems AlpiType · Anton Lytvynenko

He did not account for the possibility that a language model could learn to meet the human where they are — at the level of human language — instead of the human having to learn the language of the machine. That is exactly what happened.

Consequences: What This Means for the AI Operator

If intellectual augmentation happens through the quality of the request, then the ability to formulate precise requests is a new form of cognitive capital. Not a replacement for intelligence — but its lever.

The AI operator within this framework is the person who masters this lever in their domain of expertise. The difference between an operator and a regular user does not lie in access to the tool — everyone has that. It lies in the quality of thinking that precedes the request. And that quality does not come from a neural implant. It comes from domain knowledge, practice, and the willingness to keep refining until the answer is actually useful.

26 Anthropology of AI Systems AlpiType · Anton Lytvynenko

Chapter 12

The Prompt as an Act of Thinking

A prompt is not a command to the machine. It is an act of thinking: the attempt to understand what one actually wants — and to articulate it precisely enough that the answer is useful. The quality of the prompt depends on the quality of problem understanding, not on knowing the right syntax.

Coder and Architect: What Has Happened to Developers

In 2022, a software developer spent 40–60% of their working time writing code. By 2025, a significant portion of that is handled automatically. This does not mean developers have disappeared — what has changed is what they are valued for.

The developer who thought like a coder — "my value lies in knowing syntax and writing functions quickly" — suddenly found themselves in a situation where that knowledge is no longer a bottleneck. The tool writes syntax no worse, often better.

The developer who thought like an architect — "my value lies in understanding the problem, defining the solution structure, and knowing where the system might break" — found that demand for their skills had not decreased but increased.

The logic is not new: when a tool automates execution, value shifts to definition. From craftsman to planner. From translator to author. From calculator to analyst. From coder to architect.

The Prompt Follows the Same Logic

When an AI operator formulates a task for a language model, they perform the same action as an architect before a developer: they do not write code — they define what is to be built and why. They do not generate text — they pose a task with sufficient precision so that the generation is meaningful.

The difference between coder and architect is the difference between syntax and semantics. Between "how to write" and "what and why to build." That same difference separates users from AI operators.

The Nature of the Event: Syntax to the Machine, Semantics Stays with Humans

Every skill consists of two layers:

The syntactic layer — rules, formats, sequences that ensure technically correct execution.

The semantic layer — the understanding of what is to be executed and why.

Human work in the industrial era combined both layers in a single role, because no tool could take over the syntactic level. Now one can. This is not the destruction of work. It is its layer separation. The syntactic layer goes to the machine. The semantic layer stays with the human — and becomes more visible than ever before, because the syntactic noise no longer obscures it.

27 Anthropology of AI Systems AlpiType · Anton Lytvynenko

Misconceptions Around Prompt Engineering

Since 2022, an entire industry has emerged around "prompt engineering": courses, certificates, handbooks with hundreds of pages, lists of "magic words." Some techniques do deliver better results — but not because they are "keys" to the model. Rather, because they force the person writing the prompt to think more clearly.

"Think step by step" is useful not because it is a command to the AI. But because it raises the question: what steps are there in this task in the first place?

A prompt is a mirror of thinking. An effective prompt is the visible form of clear thinking. A poor prompt is the visible form of unclear thinking.

The Prompt as a Process, Not a Command

An effective prompt is rarely the first one. It is the result of iteration:

Formulate the task: Not the symptom — the actual task. Not "something is wrong with our process," but "we want to reduce the time from order receipt to specification confirmation from 48 to 8 hours, while maintaining the same verification quality."

Provide context: What does the model need to know that it does not know by default? Industry, constraints, previous attempts, success criteria.

Determine the format: What will happen with the answer? An analysis for executive management differs from a technical solution for implementation.

Verify and refine: Does the answer solve the task? If not — where is the gap between the request and what was actually needed?

The Other Side: When the Prompt Becomes an Illusion of Thinking

If the prompt is a mirror of thinking, then a poor prompt is a mirror of poor thinking. And unlike without AI — where poor thinking produces a poor result — with AI it produces a convincingly formulated poor result: text that looks like an analysis but is not. A recommendation that appears well-reasoned but rests on an imprecisely posed task. A conclusion that sounds confident but answers the wrong question.

The architect who does not understand the real customer problem receives from the developer perfectly written code for the wrong solution. The AI operator who has not thoroughly understood the subject task receives from the model a perfectly formulated answer to the wrong question.

That is precisely why the 30% human layer is not a compliance requirement — it is an architectural decision. It preserves within the system a human who is responsible for semantic precision: for ensuring the question is right, not merely that the answer is grammatically correct.

28 Anthropology of AI Systems AlpiType · Anton Lytvynenko