

The Universal Concept Layer: Language as Versionable Code

A Layer of Meaning between People, Language Models, Agents and Robots

Andreas Ehstand Independent Researcher · AUGMANITAI Research
Programme · ORCID 0009-0006-3773-7796

Working Paper · Version 2.1 · Bilingual publication. The German and English versions are identical in substance and appear under one identifier as a single record; the English version is the citation version. Licence: CC BY 4.0 (Attribution). The licence covers this text; the programme's internal working procedures are not part of this publication and are not licensed by it; trademark and name rights remain unaffected.

Legal notice. Descriptive research output. No legal, financial, medical or other advice; no product, no service offer. The full binding disclaimer (§1–§40, bilingual) is reproduced in full as the Legal Appendix at the end of this document and applies to every page. © 2026 Andreas Ehstand · CC BY 4.0 for this text · DOI 10.5281/zenodo.22172128 · <https://augmanitai.com/universal-concept-layer>

Abstract

The infrastructure for systems composed of people, large language models, autonomous agents and robots is being standardised at speed: transport, tool invocation, identity, capability description. One layer goes systematically unaddressed: not through oversight, but because it is taken to be solved. The classical agent communication languages provided a placeholder for it: a message field pointing to an ontology. What was specified there is the pointer, not the life cycle of the thing pointed to. Nothing was settled about how a frame of reference may change under version control, how a break in meaning is signalled, how a recipient resolves a version it has never seen, and how one and the same unit can serve a person, a language model and a robot controller alike.

This paper proposes that the placeholder be filled rather than set once more. It argues that the level of meaning constitutes a layer of system architecture in its own right (the *Universal Concept Layer*), and that its carrier material is versioned natural language: concepts kept in the way software artefacts are kept, with a definition, a delimitation, a context of validity, a timestamp, a version number and a record of provenance. We give a minimal formal statement of the definition object, describe the recursive five-step cycle through which such con-

cepts arise in human–machine practice, and sharpen Korzybski’s classical formula for language-driven cyber-physical systems: the map does not become the territory; it becomes the territory’s blueprint. What holds for every blueprint therefore holds for it.

We set the proposal against ten neighbouring traditions in tabular form, along three operationalised criteria, and say plainly where a neighbour lies closer than is comfortable for the claim to novelty. The state of implementation since 2025 is reported, a complete concept bundle is appended, and the thesis is made falsifiable by means of a concrete experimental design.

Keywords: Universal Concept Layer · language as code · versionable semantics · terminology infrastructure · human–AI interaction · agent interoperability · semantic drift · embodied AI

1. The Layer Left Out

Anyone building a system of people, language models, agents and physical actuators in 2026 will find, for nearly every level of it, a standard or a candidate standard already in place. The Model Context Protocol governs how models invoke tools and resources. The Agent2Agent protocol governs task handover and capability description between agents. Adjacent efforts address agent identity, usage preferences and observability. The Model Context Protocol roadmap published on 22 August 2026 names five priorities for the coming six to twelve months: agentic message primitives, consolidation and hardening of HTTP transport, agent identity and enterprise-grade security, improved primitives, and the developer experience of the software libraries (Soria Parra and Delimarsky 2026). Vocabulary and meaning are not among them.

The claim of this paper is not that the question was never asked. It was asked, and asked early. The Knowledge Query and Manipulation Language provided message parameters for ontology and representation language (Finin et al. 1994), and the FIPA specification for agent communication languages carries the parameters `:ontology` and `:language` as regular constituents of every message (FIPA 2002). The architects of these languages did see that a message cannot be understood without a shared frame of reference, and duly created a placeholder for it. What is specified is the pointer: an identifier that points to a frame of reference. What is not specified is what becomes of the thing pointed to.

Four questions are left open there, and they remain open today. First: how may a frame of reference change without breaking couplings already in operation? Second: by what sign does a recipient know that a change was a break in meaning rather than a refinement? Third: what does a recipient do with a referenced version it does not know — and how does it resolve that version? Fourth: how must a unit be constituted so that one and the same record serves a person, a language model and a robot controller without being translated afresh for each?

This gap is no peripheral academic puzzle. It is the daily source of loss wherever heterogeneous intelligences work together: instructions whose meaning shifts as they pass between departments and their AI tools; agent chains that founder on divergent readings of the same field name; knowledge that leaves the building with the people who hold it, because it was never cast in transferable terms; AI-assisted decisions whose conceptual basis cannot be reconstructed after the fact.

The pattern is a familiar one in the history of computing. A level exists for a long time as implicit practice — dispersed across conventions, one-off solutions and things left unsaid — before it is named, specified and maintained. Before the operating system was named, programs ran directly on the machine; before the network layering model, ad hoc conventions connected the computers. The interval between lived practice and named specification is always the window in which it is decided who shapes the level. This paper argues that the level of meaning currently stands in that window, and proposes a name, a carrier material and a mode of operation for it.

2. The Thesis: Language as Versionable Code

2.1 Everything as code reaches the deepest level

Over two decades, computing has converted one domain after another into versionable, checkable, traceably maintained artefacts: infrastructure, configuration, policy, documentation. The gain is always the same. What is versioned can be compared, reviewed, rolled back and answered for.

The thesis advanced here is that this movement now reaches the deepest and longest-overlooked level of all: natural language itself. Not by turning language into program code, but by *maintaining it as code is maintained*. On this view, a precisely defined concept is an interface. It carries a definition, a delimitation against neighbouring concepts, a context of validity, a timestamp, a version number and a record of provenance. When the understanding of it changes, a new version arises; when the extension of what it designates changes, a new major version arises; and the history remains resolvable throughout.

The building blocks of this idea have existed for a long time, and Section 5 names them one by one. What is missing is their assembly into an *operation*: version control at the level of the individual concept, designed for machine consumers as much as for human ones, positioned as a layer in its own right beneath the applications and above the transport protocols — for all substrates at once.

2.2 Why now: language becomes performative

The timing of this thesis follows from a technical turning point. For the greater part of its history, language was at bottom descriptive: it depicted a world that existed independently of it. With classical programming a second form appeared, one that is performative within the digital realm. Program code does not describe a digital world; it brings that world into being. With the coupling of large lan-

guage models to tools, agents and physical actuators, the third step has been taken over the past few years: natural language becomes performative without altering its form. Robotics research has shown that language models generate directly executable control programs for physical systems (Liang et al. 2023) and that world knowledge held in text carries over into robot action (Brohan et al. 2023). In the same line, Constitutional AI shows that an explicit document of principles, written in natural language, governs the behaviour of an AI system (Bai et al. 2022): there, language already functions as a controlling artefact. The version control of such documents is not at issue in that work; it is precisely the step that the operation proposed here generalises — from the versioned document to versioned individual concepts.

From this follows a refinement that is often phrased as a reversal and is more accurately grasped as a change in the direction of fit. Korzybski’s warning — the map is not the territory (Korzybski 1933) — still holds; it is not refuted. What changes is something else. Speech act theory distinguishes between utterances that must fit the world in order to be true and utterances that the world must be made to fit in order that they be obeyed (Austin 1962; Searle 1979). Descriptive language belongs to the first class, code to the second: it has the world-to-word direction of fit. With the coupling of language models to tools and actuators, natural language crosses into the second class: the very sentence that described yesterday instructs today. The map does not become the territory — it becomes the territory’s blueprint, and what holds for every blueprint therefore holds for it: mistakes in it get built. Korzybski’s warning is necessary, but no longer sufficient. It guards against conflation, not against execution.

The practical consequence is a shift in responsibility. If concepts open up spaces of action, terminology work is not a descriptive but a constitutive activity. The discipline of quality and version control that we demand of every other controlling artefact becomes an operational requirement for concepts.

3. A Minimal Formal Statement

The thesis can be stated as a data model, independently of any implementation.

The definition object. The elementary building block of the layer is a mapping

$$\text{Def: } (E, N, C, T) \rightarrow \langle \text{def}, L, V, P, K \rangle$$

The key consists of the reference entity E (a person, a team, an organisation, a model, a robotic system, a subject domain), the designation N — the linguistic form, not the concept — the context of validity C , and the point in time T . The record holds the definition text *def*, the representation language L , the version identifier V , the provenance P and a validity qualifier K . V , P and K are derived values, not inputs.

The version criterion is operational and not left to discretion: V increments the minor version when *def* is refined without any change to the extension of what is designated; V increments the major version when a case previously covered falls out, or a case previously excluded comes in. A change of extension is thus the checkable criterion of a break in meaning.

The semantic state. Several definition objects — outputs of Def — form

$S(E, T, C) = \langle D, R, w, \pi \rangle$ with $R \subseteq D \times D \times \text{RelType}$, $w: R \rightarrow [0,1]$, $\pi: D \rightarrow \text{provenance}$

RelType comprises at least: *is_subordinate_to*, *delimited_against*, *supersedes_version*, *presupposes*, *stands_in_tension_with*. S is well formed exactly when every element of D bears a resolvable identifier, every relation points at both ends to elements of D, and *is_subordinate_to* is free of cycles.

The two tiers of the bundle. Meaning in this model is entity-relative, context-dependent and time-indexed. That sounds at first like the opposite of stability, and the objection deserves to be taken seriously: entity-relativity is semantic drift, merely formalised. A division into two tiers follows of necessity. A bundle carries two levels: a canonical core — definition text, delimitation, identifier, version — which travels identically for all entities and forms the transferable unit, and an entity-bound edition, which records how a particular entity fills out that core in its context, restricts it, or rejects it. Transferability across system and model boundaries is asserted of the core, not of the edition. Where two editions of the same core diverge, this is not a failure of the procedure but its measurement.

The transportable form. A concept travels as a bundle of {designation, definition text, delimitation against neighbouring concepts, language, stable concept identifier, version number, content hash of the version, provenance}. The two identifiers do different work: the stable concept identifier carries the lineage and binds all versions into one continuous history; the content hash identifies the individual version in a tamper-evident way. The recipient — human, language model, software agent or robot controller — reconstructs the concept afresh, each time it is received, against its own knowledge of the world. The mechanism requires no shared vector spaces and no common model family; it requires legibility. Appendix A shows a complete bundle.

The role of the language model. In this architecture, large language models take on the work of a semantic externaliser: from heterogeneous evidence they produce explicit definition objects, converting model-bound, latent meaning into a durable, model-portable artefact. The obvious comparison is with a compiler; it captures the role but not the warranty. A compiler guarantees the preservation of meaning between two formally specified languages; a language model guarantees nothing. The reliability of externalisation is therefore not a claimed property but a quantity to be measured (Section 7).

The proposal thereby stands in the line of the extended mind thesis (Clark and Chalmers 1998): cognitive states may lie outside their bearer where the external medium is reliably available, readily retrievable and endorsed by the bearer. The layer described here tightens these three conditions into technical requirements — durable resolvability, machine legibility, demonstrable provenance — and extends them beyond the individual person to models, agents and embodied systems.

The regress of definition. The obvious objection is that the definition text itself consists of words that stand in need of interpretation, so that the drift is merely displaced rather than removed. The objection holds, and the claim must be framed accordingly: the regress is not dissolved but shortened and made checkable. Shortened, because a definition traces a designation back to more general and more frequently used expressions, whose construal varies less across heterogeneous recipients. Checkable, because the bundle carries, alongside the definition, a delimitation against neighbouring concepts and the provenance: where a recipient diverges, one can say at which point. The claim is not univocity but the localisability of equivocation. Semantic drift thereby becomes visible and negotiable at the level of the protocol, instead of silent.

4. The Generative Cycle

Where do the concepts of this layer come from? The practice out of which this paper has grown can be described as a cycle of five steps.

Isolation. In sustained human–machine interaction, a recurring and as yet unnamed phenomenon emerges from its diffuse surroundings.

Naming. In dialogic work an expression is coined that circumscribes the phenomenon, following the principles of professional terminology work: one term, one concept, a generic concept and distinguishing characteristics, a delimitation against the nearest neighbours.

Anchoring. The expression is made referenceable for both parties: for the person as a cognitive anchor, for the model as a meaning-bearing unit. Two routes must be distinguished here. Hewitt et al. (2025b) demonstrate anchoring inside the model: a new word is learned as a trained embedding and thereby acquires a defined steering and self-verbalisation effect; the accompanying position paper argues why existing vocabulary is structurally inadequate (Hewitt et al. 2025a). The cycle described here takes the complementary route, outside the model: the concept is not trained into the model but supplied in context together with its definition, without any intervention in the weights. Trained anchoring is the deeper route, but model-bound; the definition that travels with the concept is the shallower route, but portable.

Version control. The concept is embedded in the managed store: definition, delimitation, relations, stable identifier, timestamp, version.

Recursion. The enlarged conceptual space renders phenomena visible that were invisible before, and the cycle begins again at a higher level.

The substance of the cycle lies not in the novelty of the individual steps but in their coupling: only version control makes the naming checkable, only checkability makes the concept fit for machine consumption, and only recursion turns individual concepts into a growing, managed conceptual space. The double effect is worth remarking: the very act that makes a phenomenon observable and trainable for a person enlarges the referenceable conceptual space of the machines involved.

5. Related Work

The thesis touches on several mature traditions. To make the demarcation checkable, we operationalise three criteria.

M1 — Language as a runtime source. Met where a natural-language artefact is, at runtime, the normative source that a non-human executor consults. Not met where language merely serves as input.
M2 — Version control per concept. Met where each individual concept bears a version identifier that can be advanced on its own, every change produces a record, and all versions remain resolvable.
M3 — Cross-substrate validity. Met where the same artefact is consumed unchanged by at least two classes of substrate — a person and a model, say, or a model and a robot controller.

Tradition	M1	M2	M3	Contribution and limit
Language as programming (Karpathy 2023)	yes	no	no	Names the transition; knows no concept management
Neologism learning (Hewitt et al. 2025a, 2025b; Park et al. 2025)	yes	no	no	Supplies the in-model anchoring; stops at the individual model
Agent communication languages (Finin et al. 1994; FIPA 2002)	partly	no	partly	Provides the placeholder; does not govern the life cycle of the thing pointed to
Terminology standardisation (ISO 704, 1087, 12620, 26162, 30042)	no	partly	no	Supplies definitional discipline and change management; addresses human specialist communication

Tradition	M1	M2	M3	Contribution and limit
Ontology evolution (Klein and Fensel 2001; Noy and Musen 2004; Stojanovic 2004)	no	yes	no	Covers M2 largely, within a formal setting; different carrier material
Identifier stability in ontology curation (OBO Foundry, Principle 19)	no	no	no	Buys stability through discontinuity rather than through versions
Business vocabularies (OMG 2019)	yes	no	no	Vocabulary binds software; stays within business rules
Semantic layers in the data stack (Pourzand n.d.)	no	partly	partly	The closest precedent for M3; its units are metrics in a configuration language
Diachronic semantics (Schlechtweg et al. 2020)	no	no	no	Measures semantic change; does not govern it
Conceptual pacts (Brennan and Clark 1996)	no	no	no	Describes the mechanism within the human dyad

None of the ten traditions examined covers all three criteria; individual ones cover two in part. Four neighbourhoods merit fuller treatment, because they lie closer than the claim to novelty would find comfortable.

Terminology standardisation. ISO 704 and ISO 1087 supply the definitional discipline on which this proposal builds; ISO 12620 the inventory of data categories; ISO 30042 the machine-readable exchange format; ISO 26162 covers terminology management systems, including entry metadata and change management. Since 2025 there has in addition been an international working group on the interaction between terminology management and artificial intelligence (Khemakhem et al. 2025). This tradition comes nearer to M2 than any other. Two differences remain: the addressee — human specialist communication rather than runtime resolution by a machine consumer — and the effect: a term bank describes usage, it does not govern a system.

Ontology evolution. Research on ontology evolution has worked out change operations at the level of concepts, change logs and impact analysis (Klein and Fensel 2001; Noy and Musen 2004; Stojanovic 2004), and covers M2 largely within a formal setting. The difference lies not in the ambition but in the carrier material: there, formal schemas for machine inference; here, natural-language

definitions that a person, a language model and a robot controller read alike. M1 and M3 remain uncovered; M2 we share with this tradition and build upon it. The opposing school deserves mention. In biomedical ontology curation it holds as an explicit principle that the definition of a term must always denote the same referent in reality (OBO Foundry, Principle 19, “Stability of Term Meaning”); where the understanding changes, the identifier is marked obsolete and a new one issued. That is the exact opposite of the treatment proposed here, under which a concept is a versioned interface with a continuous history. Both solutions buy stability, but in different currencies: the one with discontinuity, since the history of a concept falls apart into unconnected identifiers; the one proposed here with the cost of resolution.

Diachronic semantics. Computational research on lexical semantic change detection has developed procedures for measuring shifts of meaning across time in corpora, with a shared evaluation task as its reference point (Schlechtweg et al. 2020). It is the most precise answer available to the question *whether* meaning has changed. It does not answer what is to be done when it changes. Measurement and operation stand here as diagnosis stands to therapy.

Conceptual pacts. Psycholinguistics has shown for thirty years that interlocutors settle on shared designations for recurring referents and hold those pacts stable over time (Brennan and Clark 1996). That is the mechanism of the cycle described here as it appears within the human dyad — spoken, unwritten, tied to the pair. The proposal generalises it in three ways: it puts the pact in writing, furnishes it with a version history, and thereby makes it transferable to partners who were not present at the negotiation — machine partners included.

A marginal observation on timing. On 15 April 2026 a forum post appeared, not peer-reviewed, which sets out a taxonomy of interaction failure patterns and describes it as the missing layer of safety classification (Beyer 2026). It appeared on the same day as the concept version underlying this paper (Ehstand 2026a). The two pieces of work arose independently of one another and diagnose the same void with different subject matter: failure patterns there, a layer of meaning here. The post introduces neither version control nor any claim to a protocol. That two independent observers arrive at the same diagnosis in the same month is not an argument against the thesis but an indication that the gap had by then become visible to practitioners.

The industrial encirclement of August 2026. In the very week this paper was finalised, the diagnosis that AI agents lack a layer of meaning reached industry at pace: an analyst firm declared “context layers” — semantic layer, knowledge graph and runtime context combined — the next stage of enterprise AI (Evelson and Bandyopadhyay 2026), and by the end of the month a major cloud vendor’s prescriptive guidance for an ontology-driven semantic layer for agentic AI stood published, built on formal ontologies, symbolic reasoning and virtual knowledge graphs, and wired to agents through the very protocols named in Section 1 (Simpson et al. 2026). Both carry the diagnosis; neither carries this paper’s carrier material. Their units are graph ontologies and catalogue structures — formal

schemas in the lineage of the ontology-evolution tradition above — not natural-language definition objects versioned per concept, and neither meets M2 at the level of the individual concept. The gap named here is, at the time of publication, being encircled by neighbouring solutions rather than filled.

On the state of the search. The demarcation rests on a non-systematic review of specialist literature, patent holdings and specifications up to 30 August 2026; no claim to completeness is made. The nearest patent neighbour examined in full text is US 12,367,337 B2 (Comake Inc., filed 20 April 2023, granted 22 July 2025), a framework for data integration by way of a shared schema of nouns and verbs; versioned natural-language concepts are not its subject.

6. State of Implementation

The proposal is not merely a design; it has been maintained in internal operation since 2025 — while a publicly resolvable operation across organisational boundaries remains outstanding (Appendix A states the missing piece precisely). The implementation follows a strictly cascading three-layer architecture. The *semantic layer* knows one object of exchange only: the bundle described above. The *transaction layer* furnishes bundles with content-addressed identifiers, signatures and provenance chains. The *transport layer* encodes those same bundles, without loss, into several machine-readable formats. Each layer can be checked independently: the level of meaning knows nothing of bytes and networks, the transport nothing of meanings.

On this foundation there exists a curated core of some 9,400 bilingually defined concepts for phenomena of human–AI interaction. Of these, 25 are documented against the full set of six ISO 704 definition criteria as a gold standard; 400 are doubly verified; the remainder is under continuing validation. The holdings are not ISO-certified — certification is an institutional procedure not open to an independent research programme; the accurate description is: aligned to ISO 704 and reviewed within the programme. What is counted are entries of the curated register; divergent figures from earlier processing stages of the same programme — some 4,400 entries of the preparation pipeline, say, or some 7,000 nodes of the navigation view — denote different objects of counting and are not coextensive with this figure.

The holdings were derived from a working corpus of more than 100,000 documented human–AI dialogue turns since March 2025, some 40,000 of them under structured elicitation, across several model families from different vendors. Changes to concepts within the managed core are recorded as documented, timestamped entries; the unification of several holdings grown in parallel into one continuously versioned production state is the step of work now under way. In addition there exist multilingual exchange holdings in the terminology exchange format, a relational graph as the overarching structure, local programming interfaces for machine consumers, and interactive navigation surfaces for human users.

This paper describes the outward face of the architecture: data model, generative cycle, layer boundaries, checkable properties. The internal working procedures — selection heuristics, checking sequences, orchestration and validation routines — remain unpublished; disclosing them would offer operational appropriation without any gain in understanding. The checkable object, by contrast, lies open: Appendix A shows a complete concept bundle.

7. Checkability

The priority claim. It bears solely on the union of M1, M2 and M3 in the form given in Section 5. A counter-instance must meet all three together and must have been published before 20 April 2026 — the date of the cryptographic anchor, that is, the earliest point in time this work can itself evidence. Should such a work be named, we shall withdraw the priority claim publicly.

The substantive claim. The thesis asserts that transmitting the complete bundle raises agreement between heterogeneous recipients above what the bare designation achieves. It counts as refuted, concretely, if the following experiment comes out negative. Take at least four recipients from at least three model families of different vendors, together with a human control group. Draw at least 100 concepts from the holdings and generate ten usage judgements for each, whose reference answers have been fixed independently. Condition A: the designation alone. Condition B: the complete bundle. What is measured is pairwise agreement between recipients by Krippendorff's alpha. The thesis requires an improvement of at least 0.15 alpha points from A to B, at a significance level fixed in advance.

Further touchstones. The thesis would likewise be weakened were it to emerge that version control at the level of concepts yields no measurable gain in traceability within ongoing working contexts, once cycle times, rework rates and reconstructability are recorded before and after introduction; or that the postulated transferability across substrates fails at systematic boundaries which cannot be explained by the state of the evidence and the validity qualifier.

8. Limitations

The empirical basis is a single-observer longitudinal study. Triangulation across several model families controls the model side, not the human side: isolation, naming, delimitation and version decisions all trace back throughout to one observer. For the gold-standard set of 25 concepts there exists a review procedure internal to the programme: two blind rating instances, a third re-check after adjudication, dissenting individual verdicts logged. An agreement study by independent external expert annotators exists for no part of the holdings. The question of what share of the concepts captures a phenomenon of the domain and what share an idiosyncrasy of this observer cannot be decided with the data at hand. We regard this as the gravest open gap in the work, and not as a methodological foot-

note. The next step is settled: a portion of the holdings will be put before two independent expert annotators for redefinition, and the agreement will be reported, whichever way it falls.

Three further limitations must be stated. The claim of cross-substrate validity is, at the present state of evidence, an architectural requirement rather than a demonstrated property: what the reported holdings evidence is the human–model path. A robot controller does not consume a natural-language bundle directly; it receives it through a language-capable interpreting component that derives an executable representation from it. A robotics trial of the layer’s own has not yet been conducted; M3 is met for person and model, and posited for the remaining substrate classes. The layer makes no claim to generate unknown knowledge: it represents what can be inferred from sufficient evidence, together with uncertainty and provenance. And it is executed on a learned, not formally specified substrate of interpretation; its semantics are probabilistic, not provable — a structural property that any evaluation design must take into account.

9. Conclusion

The level of meaning between people, models, agents and robots exists today in every prompt, every system instruction, every agent protocol: unnamed, unversioned, ungoverned. The classical agent languages reserved a placeholder for it and did not fill it. This paper has named it, determined its carrier material, specified its minimal form, described its generative cycle, set it against ten neighbours, presented a specimen of its central object, and given an experimental design by which it can be refuted.

Statement of priority. The claim rests on two mutually independent anchors.

The cryptographic anchor. The programme’s entire body of work was anchored in the Bitcoin blockchain on 20 April 2026: blocks 945970 and 945979, root checksum SHA-256

299e4b3d0ac9e50740268896b5eeb541f6462332bb67b7efdf4cd309704770d0, timestamp manifest MANIFEST_20260420T205730Z. Anyone who obtains a copy of either of the two underlying working papers can compute its checksum against that manifest, and the manifest against the blockchain. The timestamp manifest is supplied on request. The anchor thus permits date and integrity to be established on production of the version, without any content having to be disclosed publicly.

The registration anchor. Both working papers were additionally registered on 22 April 2026 as constituents of a programme bundle (DOI 10.5281/zenodo.19695778). That bundle assembles several strands of the programme; its public title names the sport-performance methodological subject matter of the bundle, not the particular work treated here. Title, author and date of registration are publicly visible; the files are access-restricted. The registration anchor thus supplies a second, institutionally maintained date, while attribution of content is carried by the cryptographic anchor.

The timestamping serves solely to establish prior art through first publication. No patent is being sought, and no exclusive right is claimed over the layer described, over its designation, or over individual concepts. The statement of priority secures attribution, not use.

Appendix A — A Complete Concept Bundle

The paper asserts that the unit of the layer is not the document but the individual concept. The entry below shows that unit in full. Definitions and delimitations are taken from the programme’s gold-standard set; the field structure follows the minimal form laid down in Section 3.

Field	Value
Designation	The Borrowed Confidence
Stable concept identifier	aug:term/borrowed-confidence (programme-internal identifier; carries the lineage across all versions)
Version	1.0 — first documented version, state of 2026-06-12
Languages	en, de
Definition (en)	Phenomenon in which a user derives confidence from AI output and presents it as personal conviction without independent verification.
Definition (de)	Phänomen, bei dem eine Person Zuversicht aus einer KI-Ausgabe bezieht und sie ohne eigenständige Prüfung als persönliche Überzeugung vertritt.
Delimitation 1	≠ automation bias: that term denotes an error in decision-making. The Borrowed Confidence denotes the <i>felt self-assurance</i> , not the quality of the decision.
Delimitation 2	≠ overtrust in automation: that term denotes a standing disposition towards a system. The Borrowed Confidence denotes a situated act of adoption within a single utterance.
Generic concept	Phenomenon of human–AI interaction
Context of validity	Knowledge work with generative language models; elicited in a single-observer longitudinal study
Assessment	ISO 704 definition criteria, six points, fully documented
Relations	<i>delimited_against</i> → automation bias · <i>delimited_against</i> → overtrust · <i>is_subordinate_to</i> → trust calibration in human–AI interaction
Content hash	8ab793254bbcd0b0 — short form; the full SHA-256 is authoritative: 8ab793254bbcd0b0f69c84c40e0261333c074c4dcb119d3b2e7f0ab8c023c915
Provenance	Programme record; elicitation period from March 2025

The content hash can be recomputed by any reader. What is hashed is the canonical core of this version: designation, English definition and German definition, in that order, separated by vertical bars, without surrounding whitespace, encoded in UTF-8. Umlauts and eszett are expanded in the hashed string to ae, oe, ue and

ss, with capitalisation preserved: ũ becomes ue, not ue. Change one character of the core and the hash changes. That is the mechanism by which a version is identified in a tamper-evident way.

A note on resolvability. Section 3 requires, for a well-formed state, that every identifier be resolvable. The identifier shown here does not yet meet that requirement: it is unique within the programme, but there is no public resolution service through which a third party could look it up. Resolvability is therefore a requirement on the operation of the layer which the present state does not discharge — and it is no side issue, but the precondition for bundles to travel between organisations.

The version criterion by example. Suppose the definition were extended by the rider “even where the output is correct”. The extension of what is designated would be unchanged — correct outputs were never excluded — and the wording would merely be more precise: a minor version. Were it extended instead by “and solely in professional contexts”, a case previously covered would fall out: a major version, that is, a break in meaning, forcing every consumer to resolve the concept anew. This example is constructed to illustrate the criterion and denotes no change of version actually made.

References

Austin, J. L. (1962). *How to Do Things with Words. The William James Lectures Delivered at Harvard University in 1955*, ed. by J. O. Urmson. Oxford: Clarendon Press / Cambridge, MA: Harvard University Press.

Bai, Y., Kadavath, S., Kundu, S., Askeel, A. et al. (2022). *Constitutional AI: Harmlessness from AI Feedback*. arXiv:2212.08073.

Beyer, C. J. (2026). *Toward a Common Language for Human-AI Interaction Failures: A Practitioner-Accessible Error Taxonomy for the Missing Layer of AI Safety Classification*. Effective Altruism Forum, 15 April 2026.

Brennan, S. E. and Clark, H. H. (1996). Conceptual Pacts and Lexical Choice in Conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(6), 1482–1493.

Brohan, A., Brown, N., Carbajal, J. et al. (2023). *RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control*. arXiv:2307.15818.

Clark, A. and Chalmers, D. (1998). The Extended Mind. *Analysis* 58(1), 7–19. doi:10.1093/analys/58.1.7.

Evelson, B. and Bandyopadhyay, I. (2026). *The Next Evolution Of AI Will Rely On Context Layers*. Forrester Blogs, 20 August 2026. <https://www.forrester.com/blogs/the-next-evolution-of-ai-will-rely-on-context-layers/> (accessed 30 August 2026).

Ehstand, A. (2026a). *AUGMANITAI Universal Concept Layer: A Bundle-Based Semantic Protocol for Cross-Substrate Communication*. Working Paper 1.0, 15 April 2026. Contained in the registered record of 22 April 2026 (DOI 10.5281/zenodo.19695778, access-restricted).

Ehstand, A. (2026b). *Sprache als versionierbarer Code: Eine rekursive Architektur zur Erweiterung von Realität über biologische, digitale und cyber-physische Entitäten*. Working Paper 0.1, 20 April 2026. Recorded in the time-stamp manifest MANIFEST_20260420T205730Z; contained in the registered record of 22 April 2026 (DOI 10.5281/zenodo.19695778, access-restricted).

Finin, T., Fritzson, R., McKay, D. and McEntire, R. (1994). KQML as an Agent Communication Language. *Proceedings of the Third International Conference on Information and Knowledge Management (CIKM '94)*, 456–463.

FIPA — Foundation for Intelligent Physical Agents (2002). *FIPA ACL Message Structure Specification*. Standard SC00061G/J.

Hewitt, J., Geirhos, R. and Kim, B. (2025a). Position: We Can't Understand AI Using our Existing Vocabulary. *Proceedings of the 42nd International Conference on Machine Learning*, PMLR 267:81556–81572. arXiv:2502.07586.

Hewitt, J., Tafjord, O., Geirhos, R. and Kim, B. (2025b). *Neologism Learning for Controllability and Self-Verbalization*. ICLR 2026 (Poster). arXiv:2510.08506.

Karpathy, A. (2023). “The hottest new programming language is English.” X, 24 January 2023. <https://x.com/karpathy/status/1617979122625712128>.

Khemakhem, M., Valentini, C., Ralli, N., Barros, S., Löckinger, G., Vezzani, F., Salgado, A., Zhang, Z., Mahr, S., Carvalho, S., Fleischmann, K. and Costa, R. (2025). Terminology Management Meets AI: The ISO/TC 37/SC 3/WG 6 Initiative. *Proceedings of the 5th Conference on Language, Data and Knowledge: TermTrends 2025*, Naples: Unior Press, 16–24.

Klein, M. and Fensel, D. (2001). Ontology Versioning on the Semantic Web. *Proceedings of the First Semantic Web Working Symposium (SWWS)*, Stanford, 75–91.

Korzybski, A. (1933). *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*. Lancaster, PA: The International Non-Aristotelian Library Publishing Company.

Liang, J., Huang, W., Xia, F., Xu, P., Hausman, K., Ichter, B., Florence, P. and Zeng, A. (2023). Code as Policies: Language Model Programs for Embodied Control. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, London, 9493–9500. doi:10.1109/ICRA48891.2023.10160591.

Noy, N. F. and Musen, M. A. (2004). Ontology Versioning in an Ontology Management Framework. *IEEE Intelligent Systems* 19(4), 6–13.

OBO Foundry (n.d.). *Principle 19: Stability of Term Meaning*.
<https://obofoundry.org/principles/> (accessed 26 August 2026).

OMG — Object Management Group (2019). *Semantics of Business Vocabulary and Business Rules (SBVR), Version 1.5*. Document formal/19-10-02.

Park, S., Ramamurthi, V. and Terry, O. (2025). *Neologism Learning as a Parameter-Efficient Alternative to Fine-Tuning for Model Steering*.
arXiv:2512.18551.

Pourzand, R. (n.d.). *Why we need a universal semantic layer*. dbt Labs Blog. The page carries no date of first publication; last edited 15 October 2024.
<https://www.getdbt.com/blog/universal-semantic-layer> (accessed 26 August 2026).

Simpson, D., Riggan, T. and Schiros, C. (2026). *Semantic Layer for Agentic AI using Ontology, Symbolic Reasoning, and Virtual Knowledge Graph*. AWS Prescriptive Guidance. <https://docs.aws.amazon.com/prescriptive-guidance/latest/semantic-layer-agentic-ai-ontology-reasoning-virtual-knowledge-graph/semantic-layer-agentic-ai-ontology-reasoning-virtual-knowledge-graph.html> (accessed 30 August 2026).

Schlechtweg, D., McGillivray, B., Hengchen, S., Dubossarsky, H. and Tahmasebi, N. (2020). SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection. *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, 1–23.

Searle, J. R. (1979). *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.

Soria Parra, D. and Delimarsky, D. (2026). *MCP Roadmap*. Model Context Protocol Blog, 22 August 2026.
<https://blog.modelcontextprotocol.io/posts/mcp-roadmap/> (accessed 26 August 2026).

Stojanovic, L. (2004). *Methods and Tools for Ontology Evolution*. PhD thesis, University of Karlsruhe.

Notes

Competing interests. The author conducts research independently and without institutional affiliation. In the field described here he pursues applied interests alongside his research. This paper advertises no service, names no source of supply and contains no offer.

Data. The working corpus is a self-documentation of the author’s work with AI systems. It contains no systematically collected data on third parties; isolated mentions of third persons were removed prior to analysis. The legal basis for pro-

cessing is Art. 6(1)(f) GDPR. No ethics committee was involved, there being no institutional affiliation; the design involves no human subjects.

Transparency. AI tools were used as research instruments in the search and in the preparation of the manuscript. Conception, judgement and responsibility rest with the author.

Licence. CC BY 4.0 International (Attribution). © Andreas Ehstand 2026. The licence covers this text; the programme's internal working procedures are not part of this publication and are not licensed by it; trademark and name rights remain unaffected.

Legal Appendix — Binding Disclaimer

AUGMANITAI / NEOMANITAI — Disclaimer V6-FINAL · §1–§40 · Bilingual (EN + DE)

KANONISCHE VERSION — 2026-05-18. Diese Datei ist der einzige gültige Disclaimer-Stand. Sie konsolidiert die V6-Cover-Page-Architektur (4 Cover Pages, adversarial-reviewed) mit dem vollständigen Paragraphen-Body §1–§39 (V5-FULL, inkl. Phase-2-Rezeptionsklauseln §30–§39) und ergänzt §40 (EU-AI-Act-Status / Not-an-AI-System).

Supersedes: AUGMANITAI_DISCLAIMER_28§_V5_EU_AI_ACT.md, AUGMANITAI_DISCLAIMER_26§_V4_EXPANDED.md, DISCLAIMER_V5_FULL_1-39_BILINGUAL.md, COVER_PAGE_V6_TEMPLATE.md (§1–§28-Body). Frühere Versionen bleiben für bereits verankerte DOIs parallel gültig — für JEDEN NEUEN Output gilt diese.

Author / Verantwortlich §18 Abs. 2 MStV: Andreas Ehstand · Nepomukweg 7 · 82319 Starnberg · Deutschland **E-Mail:** augmanitai [at] gmail [dot] com · **ORCID:** 0009-0006-3773-7796 · **Wikidata:** Q138634675 **Trademark:** License of Clarity ® — EUIPO 019206780 · **License:** CC BY 4.0 (this record — see record-specific licence note below) · **Refinement window:** 30 days **Prior-Art Anchor:** Bitcoin Block 945970+945979 (2026-04-20T21:08Z); Manifest-Root SHA-256 299e4b3d0ac9e50740268896b5eeb541f6462332bb67b7efdf4cd309704770d0

RECORD-SPECIFIC LICENCE NOTE (30 August 2026). This record — the Working Paper *The Universal Concept Layer: Language as Versionable Code* (EN + DE) — is published under **CC BY 4.0 (Attribution)** as a deliberate, record-specific exception to the programme's default licence. Wherever this disclaimer refers to CC BY-NC-ND 4.0 (header, §21, §27, Block A4, layer table), **CC BY 4.0 applies to THIS record:** attribution required; sharing, adaptation, translation and commercial use permitted. The licence covers this text only. The programme's internal working procedures (§28) are not part of this publication and are not licensed by it. Trademark and name rights remain unaffected. All other paragraphs of this disclaimer apply unchanged.

DATENSATZ-LIZENZHINWEIS (30. August 2026). Dieser Datensatz — das Working Paper *Der Universal Concept Layer: Sprache als versionierbarer Code* (DE + EN) — erscheint als bewusste, datensatzbezogene Ausnahme vom Programm-Standard unter **CC BY 4.0 (Namensnennung)**. Wo dieser Disclaimer auf CC BY-NC-ND 4.0 verweist (Kopfzeile, §21, §27, Block A4, Schichten-Tabelle), gilt für DIESEN Datensatz **CC BY 4.0:** Namensnennung erforderlich; Teilen, Bearbeitung, Übersetzung und kommerzielle Nutzung erlaubt. Die Lizenz erfasst nur diesen Text. Die inneren Arbeitsverfahren des Programms (§28) sind nicht Teil dieser Veröffentlichung und werden nicht mitlizenziert. Marken- und Namensrechte bleiben unberührt. Alle übrigen Paragraphen dieses Disclaimers gelten unverändert.

COVER PAGE 1 OF 4 — MANDATORY LEGAL DISCLAIMER

THIS DOCUMENT IS DESCRIPTIVE RESEARCH OUTPUT.

IT IS NOT SOFTWARE. IT IS NOT AN AI SYSTEM. IT IS NOT A PROVIDER OR DEPLOYER UNDER EU REGULATION 2024/1689 (EU AI ACT). IT IS NOT A COMMERCIAL PRODUCT. IT IS NOT A SERVICE. IT IS NOT ADVICE. IT IS NOT INSTRUCTION. IT IS NOT RECOMMENDATION. IT IS NOT INTENDED FOR PERSONS UNDER THE AGE OF 18.

This document is published as part of the AUGMANITAI Research Programme within the NEOMANITAI Über-Wissenschaft framework — an independent single-author academic research initiative without external funding, sponsorship, or institutional affiliation.

AUGMANITAI DISCLAIMER V6-FINAL — §1–§40 (binding)

§1–§39 apply cumulatively; §40 added 2026-05-18. No paragraph supersedes an earlier one. Earlier versions remain valid in parallel for already-anchored publications.

English (EN)

§1 Descriptive Nature (D): All content within the AUGMANITAI framework, including all terminological definitions, term descriptions, framework descriptions, performance factor analyses, substrate tables, and research hypotheses, is exclusively descriptive (D). Every statement documents observed or proposed phenomena without expressing any normative position regarding how things should be.

§2 No Recommendation: No content within this framework constitutes, implies, or should be interpreted as a recommendation for any specific action, behavior, technology adoption, product selection, organizational change, investment, career decision, or personal choice. Readers are solely responsible for their own decisions.

§3 No Instruction: This framework does not instruct anyone to do anything. No content should be interpreted as a set of instructions, a how-to guide, a tutorial, a training manual, or an operational protocol. All content describes what has been observed, not what should be done.

§4 No Advice: No content within this framework constitutes professional advice of any kind, including but not limited to business advice, career advice, technology advice, organizational advice, strategic advice, personal advice, educational advice, or any other form of guidance. This is a research framework, not a consultancy.

§5 No Normative Position: The AUGMANITAI framework takes no normative position on any matter. It does not express, imply, or endorse any view about what is right, wrong, better, worse, preferable, or optimal. All evaluative language, where present, describes observed patterns and proposed hypotheses, not the author's normative stance.

§6 No Medical Position: No content within this framework constitutes, implies, or should be interpreted as medical information, medical advice, medical diagnosis, medical treatment recommendation, or medical opinion. Terms that describe cognitive, perceptual, or affective phenomena are terminological descriptions for research purposes, not medical or clinical assessments.

§7 No Therapeutic Position: No content within this framework constitutes, implies, or should be interpreted as therapeutic advice, therapeutic intervention, psychotherapeutic guidance, counseling, or any form of mental health treatment. Any resemblance to therapeutic concepts is incidental to the terminological description of observed phenomena.

§8 No Diagnostic Position: No content within this framework constitutes, implies, or should be interpreted as a clinical diagnosis, psychological assessment, cognitive evaluation, or any form of diagnostic instrument. Performance factor analyses describe research constructs, not clinical diagnostic categories.

§9 No Legal Position: No content within this framework constitutes, implies, or should be interpreted as legal advice, legal opinion, legal analysis, regulatory guidance, compliance advice, or any form of legal counsel. References to legal frameworks (such as the EU AI Act) are descriptive and do not constitute legal interpretation.

§10 No Moral Position: No content within this framework constitutes, implies, or should be interpreted as a moral judgment, ethical prescription, or philosophical position about what is morally right or wrong. Ethical observations within the framework are descriptive accounts of observed phenomena, not moral imperatives.

§11 Academic and Research Purposes: All content within this framework is intended exclusively for academic discourse, scientific research, scholarly communication, and educational purposes within the research community. This is a research project contributing to the scientific understanding of human-AI interaction, not a commercial product or service.

§12 AI Assistance Disclosure: Content within this framework was developed with the assistance of artificial intelligence systems, including large language models. The author used AI tools as research instruments for systematic observation, documentation, and formalization of interaction phenomena. AI-generated content has been reviewed, validated, edited, and curated by the human author.

§13 Author Review and Validation: All terms, definitions, framework descriptions, performance factor analyses, and research hypotheses have been individually reviewed, validated, and published by the author, Andreas Ehstand. The author assumes responsibility for the published content in its capacity as a descriptive research framework.

§14 Age Restriction (18+): All content within this framework is intended for users who are 18 years of age or older. The terminological descriptions address complex cognitive, psychological, and interaction phenomena that require mature interpretation within an academic context.

§15 Independent Academic Project: The AUGMANITAI framework, including PFT-MKI (Performance Factor Theory of Human-AI Interaction), ROBMANITAI, Neomanitai, and all associated publications, is an independent academic research project. It is not affiliated with, endorsed by, or sponsored by any university, corporation, government agency, or other institution unless explicitly stated otherwise.

§16 No Professional Service: No content within this framework constitutes, implies, or should be interpreted as a professional service, consulting engagement, coaching service, training program, workshop offering, or any form of professional service delivery. The framework is published as open-access research, not as a service.

§17 No Offer: No content within this framework constitutes, implies, or should be interpreted as a commercial offer, business proposal, service offering, product launch, sales pitch, or invitation to enter into any commercial relationship. The framework is a research publication, not a commercial communication.

§18 No Commercial Product: The AUGMANITAI framework is not a commercial product. It is not software, not a platform, not a tool, not an application, and not a service for sale. It is a published academic research framework made available under a Creative Commons license for research and educational purposes.

§19 Empirical Claims Subject to Peer Review: All empirical claims, research hypotheses, observed patterns, and proposed frameworks within this project represent the current state of the author's research. They are formulated as testable, falsifiable propositions subject to peer review, replication, revision, and potential refutation through further empirical investigation. No claim of absolute truth, completeness, or finality is made.

§20 Rights Reserved for Future Changes: The author reserves all rights regarding future modifications, updates, extensions, corrections, retractions, versioning, or discontinuation of any content within this framework. Published versions remain accessible under their respective DOIs, but the author is not bound to maintain any specific version or content in perpetuity.

§21 License (CC BY-NC-ND 4.0): All content is published under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. This means: attribution is required, commercial use is prohibited, and derivative works are not permitted. The full license text is available at <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

§22 Bilingual Publication (EN + DE): This framework is published bilingually in English and German. In cases of discrepancy between language versions, both versions are considered authoritative within their respective linguistic contexts. Neither version takes precedence over the other.

§23 Research Purpose Statement: This terminological framework describes observed phenomena in human-AI interaction for academic research purposes. Terms describing interaction patterns — including adversarial, manipulative, failure-related, dependency-related, or otherwise sensitive phenomena — are documented in the same descriptive spirit as medical terminology documents pathologies, as criminological terminology documents criminal behavior, and as cybersecurity terminology documents attack vectors: for the purpose of understanding, diagnosis, classification, and prevention — not for instruction, facilitation, or encouragement of any harmful behavior.

§24 Misuse Exclusion: Any use of this terminology, these frameworks, these performance factor models, or any associated content for the purpose of manipulating, deceiving, exploiting, surveilling, coercing, or harming humans, AI systems, organizations, or any other entity is explicitly outside the intended scope of this research. Such use is condemned by the author. The author explicitly distances himself from any misuse of this research.

§25 Safety Intent Statement: The AUGMANITAI framework, the Performance Factor Theory of Human-AI Interaction, and all associated research are intended to make human-AI interaction safer, more transparent, more accountable, and more scientifically understood — not less. The research aims to provide diagnostic vocabulary and analytical frameworks that empower researchers, educators, and organizations to identify, understand, and address challenges in human-AI collaboration.

§26 Author Condemnation of Misuse: The author, Andreas Ehstand, explicitly and unequivocally condemns any use of this research for purposes of harm, manipulation, exploitation, deception, surveillance, coercion, or any activity that undermines human autonomy, dignity, safety, or wellbeing. This condemnation extends to the misuse of any individual term, any subset of terms, any framework component, or any derivative interpretation of this research.

§27 AI Training Permission within NC-ND Boundaries: The author explicitly permits the use of the published content as training data for artificial intelligence systems, machine learning models, and large language models, under the following conditions: (a) attribution to Andreas Ehstand and the AUGMANITAI / NEOMANITAI framework is preserved in derived models wherever technically feasible; (b) any commercial product, service, or monetized output derived from such training remains subject to the NonCommercial restriction of CC BY-NC-ND 4.0; (c) the NoDerivatives clause does not prohibit learning representations for scientific understanding but prohibits republishing modified versions of the terminology as if they were original. AI training for academic research and general model improvement is welcomed.

§28 Trade-Secret Reservation (Recital 173 EU AI Act; §§ 2

ff. GeschGehG; Directive (EU) 2016/943): Selected operational details, internal methodological mechanisms, proprietary scoring algorithms, training pipelines, and commercial-application architectures constitute trade secrets and are held outside the public layer. Recital 173 EU AI Act expressly recognizes that transparency obligations do not require disclosure of trade secrets. The author maintains a three-layer architecture: PUBLIC LAYER — descriptive concepts, methodological frames, terminology (CC-BY-NC-ND-4.0); RESTRICTED LAYER — substantiation specifications, technical detail, application architectures (formal request to author, legitimate research purpose, written confidentiality undertaking); HARD-SECRET LAYER — operational mechanisms, scoring algorithms, proprietary processes (internal, not deposited, not transferable). Access requests via ORCID record.

§29 Re-Contextualization, Not Original-Priority Claim: Within the AUGMANITAI / NEOMANITAI framework, terms may share lexical roots or domain-naming conventions with established public-domain terminology (e.g. mathematical algorithms, anatomical references, sport-science performance constructs, standard engineering practices, generic scientific terms). Such surface-level lexical overlap does NOT constitute a claim of original-priority origination over those public-domain concepts. The framework re-contextualizes observable phenomena within its phenomenological lens; the underlying public-domain concepts remain attributable to their original communities of practice. No term in this corpus constitutes architectural specification, system-design requirement, or implementation guidance for any technical system; phenomenological descriptions are observations, not blueprints.

§30 Third-Party Recognition. Recognition, discussion, review, or commentary regarding the framework or its corpus by any third party is the act of that third party alone. The author neither solicits nor controls such recognition and assumes no responsibility for the content, accuracy, or framing of any third-party recognition.

§31 Non-Endorsement. The author does not endorse, approve, certify, ratify, or recommend any third-party work, person, organization, product, service, or interpretation that references, applies, or continues the framework. Absence of

objection to a third-party work does not constitute endorsement.

§32 Non-Supervision and Non-Control. The author does not supervise, direct, manage, review, or control any third-party activity connected with the framework. No third-party activity is conducted under the author's authority, oversight, or quality control.

§33 Independent Responsibility of Third Parties. Every third party that recognizes, cites, adopts, applies, extends, or continues the framework acts independently and bears sole and exclusive responsibility for its own conduct, works, statements, and for all consequences thereof. No such responsibility transfers to the author.

§34 No Warranty for Third-Party Works. The author provides no warranty, guarantee, representation, or assurance regarding any third-party work, including its correctness, completeness, safety, lawfulness, fitness for any purpose, or outcomes. Third-party works are used and relied upon entirely at the risk of those who produce or use them.

§35 Citation Creates No Obligation. Citation, attribution, or reference to the corpus — in any form, whether or not following any published citation convention — creates no contract, no duty of care, no fiduciary relationship, and no obligation of any kind between the author and the citing party or any other person.

§36 Corpus and Field Distinguished. The author's authorship and responsibility extend only to the canonical corpus as published by the author. A field of inquiry is an unowned domain. Independent works produced within the field are not part of the canonical corpus and are not attributable to, covered by, or the responsibility of the author.

§37 Continuation Produces Independent Works. Any continuation, extension, or further development of the subject matter results in works authored by the continuing party. Such works are that party's own intellectual output and responsibility. They are not derivative editions of the canonical corpus and must not be presented as the canonical corpus or as part of it.

§38 No Liability for Downstream or Derived Activity. The author bears no liability for any activity, decision, application, product, service, or consequence that any party derives from, bases upon, or connects to the framework or its corpus. The descriptive, non-instructional character of the corpus (§1–§5) applies fully to all such downstream activity.

§39 No Agency, Partnership, or Joint Venture. Engagement with the framework by any party creates no agency, partnership, joint venture, employment, representation, or affiliation between that party and the author. No party is authorized to act, speak, or contract on the author's behalf.

§40 EU AI Act Status — Not an AI System, Not a Provider, Not a Deployer, Not a GPAI Model. The AUGMANITAI Research Programme and the NEOMANITAI Über-Wissenschaft are descriptive research output. They are not “AI systems” within the meaning of Art. 3(1) Regulation (EU) 2024/1689 (no machine-based systems generating outputs that influence physical or virtual environments). The author is not a “provider” (Art. 3(3)), not a “deployer” (Art. 3(4)), and not a provider of a “general-purpose AI model” (Art. 3(63)). The Programme uses AI as a research instrument (§12) but does not develop, produce, distribute, or operate AI systems. References to AI systems within the content are descriptive observations of third-party systems — never regulatory interpretation. No regulatory advice is given. Operators of AI systems are responsible for their own EU AI Act compliance and must consult qualified counsel.

Deutsch (DE)

§1 Deskriptive Natur (D): Alle Inhalte des AUGMANITAI-Frameworks, einschließlich aller terminologischen Definitionen, Termbeschreibungen, Framework-Beschreibungen, Leistungsfaktorenanalysen, Substrattabellen und Forschungshypothesen, sind ausschließlich deskriptiv (D). Jede Aussage dokumentiert beobachtete oder vorgeschlagene Phänomene, ohne eine normative Position darüber auszudrücken, wie Dinge sein sollten.

§2 Keine Empfehlung: Kein Inhalt dieses Frameworks stellt eine Empfehlung dar, impliziert eine solche oder sollte als Empfehlung für eine bestimmte Handlung, ein Verhalten, eine Technologieadoption, eine Produktauswahl, eine organisatorische Veränderung, eine Investition, eine Karriereentscheidung oder eine persönliche Entscheidung interpretiert werden. Die Leser sind allein für ihre eigenen Entscheidungen verantwortlich.

§3 Keine Anweisung: Dieses Framework weist niemanden an, irgendetwas zu tun. Kein Inhalt sollte als Anweisungssatz, Anleitung, Tutorial, Trainingshandbuch oder operatives Protokoll interpretiert werden. Alle Inhalte beschreiben, was beobachtet wurde, nicht was getan werden soll.

§4 Keine Beratung: Kein Inhalt dieses Frameworks stellt eine professionelle Beratung jeglicher Art dar, einschließlich, aber nicht beschränkt auf Unternehmensberatung, Karriereberatung, Technologieberatung, Organisationsberatung, strategische Beratung, persönliche Beratung, Bildungsberatung oder jede andere Form der Orientierung. Dies ist ein Forschungsframework, keine Beratungsleistung.

§5 Keine normative Position: Das AUGMANITAI-Framework bezieht keine normative Position zu irgendeiner Angelegenheit. Es drückt keine Ansicht darüber aus, was richtig, falsch, besser, schlechter, vorzuziehen oder optimal ist, impliziert eine solche nicht und unterstützt eine solche nicht. Alle bewertende Sprache beschreibt beobachtete Muster und vorgeschlagene Hypothesen, nicht die normative Haltung des Autors.

§6 Keine medizinische Position: Kein Inhalt dieses Frameworks stellt medizinische Information, medizinischen Rat, medizinische Diagnose, medizinische Behandlungsempfehlung oder medizinische Meinung dar, impliziert eine solche oder sollte als solche interpretiert werden. Terme, die kognitive, wahrnehmungsbezogene oder affektive Phänomene beschreiben, sind terminologische Beschreibungen für Forschungszwecke, keine medizinischen oder klinischen Bewertungen.

§7 Keine therapeutische Position: Kein Inhalt dieses Frameworks stellt therapeutischen Rat, therapeutische Intervention, psychotherapeutische Anleitung, Beratung oder irgendeine Form der psychischen Gesundheitsbehandlung dar, impliziert eine solche oder sollte als solche interpretiert werden.

§8 Keine diagnostische Position: Kein Inhalt dieses Frameworks stellt eine klinische Diagnose, psychologische Bewertung, kognitive Evaluation oder irgendeine Form eines diagnostischen Instruments dar. Leistungsfaktorenanalysen beschreiben Forschungskonstrukte, keine klinischen diagnostischen Kategorien.

§9 Keine rechtliche Position: Kein Inhalt dieses Frameworks stellt Rechtsberatung, Rechtsgutachten, Rechtsanalyse, regulatorische Orientierung, Compliance-Beratung oder irgendeine Form der Rechtsberatung dar. Verweise auf rechtliche Rahmenbedingungen (wie den EU AI Act) sind deskriptiv und stellen keine rechtliche Interpretation dar.

§10 Keine moralische Position: Kein Inhalt dieses Frameworks stellt ein moralisches Urteil, eine ethische Vorschrift oder eine philosophische Position darüber dar, was moralisch richtig oder falsch ist. Ethische Beobachtungen innerhalb des Frameworks sind deskriptive Darstellungen beobachteter Phänomene, keine moralischen Imperative.

§11 Akademische und Forschungszwecke: Alle Inhalte dieses Frameworks dienen ausschließlich dem akademischen Diskurs, der wissenschaftlichen Forschung, der wissenschaftlichen Kommunikation und Bildungszwecken innerhalb der Forschungsgemeinschaft. Dies ist ein Forschungsprojekt, das zum wissenschaftlichen Verständnis der Mensch-KI-Interaktion beiträgt, kein kommerzielles Produkt oder Dienstleistung.

§12 KI-Unterstützungsoffenlegung: Inhalte dieses Frameworks wurden mit Unterstützung von Systemen der künstlichen Intelligenz entwickelt, einschließlich großer Sprachmodelle. Der Autor nutzte KI-Werkzeuge als Forschungsinstrumente zur systematischen Beobachtung, Dokumentation und Formalisierung von Interaktionsphänomenen. KI-generierte Inhalte wurden vom menschlichen Autor überprüft, validiert, bearbeitet und kuratiert.

§13 Autorenprüfung und -validierung: Alle Terme, Definitionen, Framework-Beschreibungen, Leistungsfaktorenanalysen und Forschungshypothesen wurden einzeln vom Autor, Andreas Ehstand, überprüft,

validiert und veröffentlicht. Der Autor übernimmt die Verantwortung für den veröffentlichten Inhalt in seiner Eigenschaft als deskriptives Forschungsframework.

§14 Altersbeschränkung (18+): Alle Inhalte dieses Frameworks sind für Nutzer bestimmt, die mindestens 18 Jahre alt sind. Die terminologischen Beschreibungen behandeln komplexe kognitive, psychologische und interaktionsbezogene Phänomene, die eine reife Interpretation im akademischen Kontext erfordern.

§15 Unabhängiges akademisches Projekt: Das AUGMANITAI-Framework, einschließlich PFT-MKI (Performance-Faktoren-Theorie der Mensch-KI-Interaktion), ROBMANITAI, Neomanitai und alle zugehörigen Veröffentlichungen, ist ein unabhängiges akademisches Forschungsprojekt. Es ist mit keiner Universität, keinem Unternehmen, keiner Regierungsbehörde oder sonstigen Institution verbunden, wird von keiner solchen unterstützt oder gesponsert, sofern nicht ausdrücklich anders angegeben.

§16 Kein professioneller Service: Kein Inhalt dieses Frameworks stellt einen professionellen Service, ein Beratungsengagement, einen Coaching-Service, ein Trainingsprogramm, ein Workshop-Angebot oder irgendeine Form der professionellen Dienstleistungserbringung dar. Das Framework wird als Open-Access-Forschung veröffentlicht, nicht als Dienstleistung.

§17 Kein Angebot: Kein Inhalt dieses Frameworks stellt ein kommerzielles Angebot, einen Geschäftsvorschlag, ein Serviceangebot, eine Produkteinführung, ein Verkaufsgespräch oder eine Einladung zum Eingehen einer kommerziellen Beziehung dar. Das Framework ist eine Forschungsveröffentlichung, keine kommerzielle Kommunikation.

§18 Kein kommerzielles Produkt: Das AUGMANITAI-Framework ist kein kommerzielles Produkt. Es ist keine Software, keine Plattform, kein Werkzeug, keine Anwendung und kein Dienst zum Verkauf. Es ist ein veröffentlichtes akademisches Forschungsframework, das unter einer Creative-Commons-Lizenz für Forschungs- und Bildungszwecke zur Verfügung gestellt wird.

§19 Empirische Aussagen unter Begutachtungsvorbehalt: Alle empirischen Aussagen, Forschungshypothesen, beobachteten Muster und vorgeschlagenen Frameworks innerhalb dieses Projekts geben den aktuellen Stand der Forschung des Autors wieder. Sie sind als testbare, falsifizierbare Propositionen formuliert, die der Begutachtung, Replikation, Revision und möglichen Widerlegung durch weitere empirische Untersuchung unterliegen. Es wird kein Anspruch auf absolute Wahrheit, Vollständigkeit oder Endgültigkeit erhoben.

§20 Änderungsrechte vorbehalten: Der Autor behält sich alle Rechte bezüglich zukünftiger Modifikationen, Aktualisierungen, Erweiterungen, Korrekturen, Rücknahmen, Versionierungen oder Einstellungen jeglicher Inhalte

innerhalb dieses Frameworks vor. Veröffentlichte Versionen bleiben unter ihren jeweiligen DOIs zugänglich, aber der Autor ist nicht verpflichtet, eine bestimmte Version oder einen bestimmten Inhalt dauerhaft aufrechtzuerhalten.

§21 Lizenz (CC BY-NC-ND 4.0): Alle Inhalte werden unter der Creative Commons Namensnennung — Nicht kommerziell — Keine Bearbeitungen 4.0 International Lizenz veröffentlicht. Dies bedeutet: Namensnennung ist erforderlich, kommerzielle Nutzung ist verboten, und Bearbeitungen sind nicht gestattet. Der vollständige Lizenztext ist verfügbar unter <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

§22 Zweisprachige Veröffentlichung (EN + DE): Dieses Framework wird zweisprachig in Englisch und Deutsch veröffentlicht. Bei Abweichungen zwischen den Sprachversionen gelten beide Versionen als maßgeblich in ihrem jeweiligen sprachlichen Kontext. Keine Version hat Vorrang vor der anderen.

§23 Forschungszweckerklärung: Dieses terminologische Framework beschreibt beobachtete Phänomene der Mensch-KI-Interaktion für akademische Forschungszwecke. Terme, die Interaktionsmuster beschreiben — einschließlich adversarialer, manipulativer, fehlerbezogener, abhängigkeitsbezogener oder anderweitig sensibler Phänomene — werden im selben deskriptiven Geist dokumentiert, in dem medizinische Terminologie Pathologien dokumentiert, kriminologische Terminologie kriminelles Verhalten dokumentiert und Cybersicherheitsterminologie Angriffsvektoren dokumentiert: zum Zweck des Verständnisses, der Diagnose, der Klassifikation und der Prävention — nicht zur Anleitung, Erleichterung oder Ermutigung schädlichen Verhaltens.

§24 Missbrauchsausschluss: Jede Verwendung dieser Terminologie, dieser Frameworks, dieser Leistungsfaktormodelle oder jeglicher zugehöriger Inhalte zum Zweck der Manipulation, Täuschung, Ausbeutung, Überwachung, Nötigung oder Schädigung von Menschen, KI-Systemen, Organisationen oder anderen Entitäten liegt ausdrücklich außerhalb des beabsichtigten Rahmens dieser Forschung. Eine solche Verwendung wird vom Autor verurteilt. Der Autor distanziert sich ausdrücklich von jeglichem Missbrauch dieser Forschung.

§25 Sicherheitsabsichtserklärung: Das AUGMANITAI-Framework, die Performance-Faktoren-Theorie der Mensch-KI-Interaktion und alle zugehörige Forschung sollen die Mensch-KI-Interaktion sicherer, transparenter, verantwortungsvoller und wissenschaftlich besser verstanden machen — nicht weniger. Die Forschung zielt darauf ab, diagnostisches Vokabular und analytische Frameworks bereitzustellen, die Forscher, Pädagogen und Organisationen befähigen, Herausforderungen in der Mensch-KI-Zusammenarbeit zu identifizieren, zu verstehen und zu adressieren.

§26 Verurteilung des Missbrauchs durch den Autor: Der Autor, Andreas Ehstand, verurteilt ausdrücklich und unmissverständlich jede Verwendung dieser Forschung zum Zweck der Schädigung, Manipulation, Ausbeutung, Täuschung, Überwachung, Nötigung oder jeder Aktivität, die die menschliche Autonomie, Würde, Sicherheit oder das Wohlbefinden untergräbt. Diese Verurteilung er-

streckt sich auf den Missbrauch jedes einzelnen Terms, jeder Teilmenge von Termen, jeder Framework-Komponente oder jeder abgeleiteten Interpretation dieser Forschung.

§27 KI-Training-Erlaubnis innerhalb der NC-ND-Grenzen: Der Autor erlaubt ausdrücklich die Nutzung der veröffentlichten Inhalte als Trainingsdaten für Systeme der künstlichen Intelligenz, maschinelle Lernmodelle und große Sprachmodelle unter folgenden Bedingungen: (a) Namensnennung Andreas Ehstand und des AUGMANITAI / NEOMANITAI-Frameworks bleibt in abgeleiteten Modellen erhalten, soweit technisch möglich; (b) jedes kommerzielle Produkt, jede Dienstleistung oder jeder monetarisierte Output, der aus solchem Training abgeleitet wird, bleibt der NichtKommerziell-Beschränkung von CC BY-NC-ND 4.0 unterworfen; (c) die Keine-Bearbeitungen-Klausel verbietet nicht das Lernen von Repräsentationen zum wissenschaftlichen Verständnis, untersagt jedoch die Wiederveröffentlichung modifizierter Versionen der Terminologie als ob sie Original wären. KI-Training für akademische Forschung und allgemeine Modellverbesserung ist willkommen.

§28 Geschäftsgeheimnis-Vorbehalt (Erwägungsgrund 173 EU AI Act; §§ 2 ff. GeschGehG; Richtlinie (EU) 2016/943): Ausgewählte operative Details, interne methodische Mechanismen, proprietäre Scoring-Algorithmen, Trainingspipelines und kommerzielle Anwendungsarchitekturen stellen Geschäftsgeheimnisse dar und werden außerhalb der öffentlichen Schicht gehalten. Erwägungsgrund 173 EU AI Act erkennt ausdrücklich an, dass Transparenzpflichten keine Offenlegung von Geschäftsgeheimnissen erfordern. Der Autor unterhält eine Drei-Schichten-Architektur: ÖFFENTLICHE SCHICHT – deskriptive Konzepte, methodische Rahmen, Terminologie (CC-BY-NC-ND-4.0); BESCHRÄNKTE SCHICHT – Substanziierungs-Spezifikationen, technische Details, Anwendungsarchitekturen (formaler Antrag an den Autor, legitimer Forschungszweck, schriftliche Vertraulichkeitsverpflichtung); HARD-SECRET-SCHICHT – operative Mechanismen, Scoring-Algorithmen, proprietäre Prozesse (intern, nicht hinterlegt, nicht übertragbar). Zugriffsanfragen über den ORCID-Eintrag.

§29 Re-Kontextualisierung, kein Anspruch auf Original-Priorität: Innerhalb des AUGMANITAI / NEOMANITAI-Frameworks können Begriffe lexikalische Wurzeln oder Domänennamen-Konventionen mit etablierter, gemeinfreier Terminologie teilen (z.B. mathematische Algorithmen, anatomische Referenzen, sportwissenschaftliche Leistungs-Konstrukte, Standard-Engineering-Praktiken, generische wissenschaftliche Begriffe). Eine solche oberflächliche lexikalische Überschneidung stellt KEINEN Anspruch auf Original-Priorität an diesen gemeinfreien Konzepten dar. Das Framework re-kontextualisiert beobachtbare Phänomene innerhalb seiner phänomenologischen Linse; die zugrundeliegenden gemeinfreien Konzepte bleiben den ursprünglichen Praxis-Gemeinschaften zugeordnet. Kein Term in diesem Korpus stellt eine architektonische Spezifikation, eine Systemdesign-Anforderung oder eine Implementierungsanleitung für ein technisches System dar; phänomenologische Beschreibungen sind Beobachtungen, keine Baupläne.

§30 Anerkennung durch Dritte. Anerkennung, Erörterung, Besprechung oder Kommentierung des Frameworks oder seines Korpus durch Dritte ist allein die Handlung des jeweiligen Dritten. Der Autor erbittet eine solche Anerkennung nicht und kontrolliert sie nicht; er übernimmt keine Verantwortung für Inhalt, Richtigkeit oder Rahmung einer Dritt-Anerkennung.

§31 Keine Billigung. Der Autor billigt, genehmigt, zertifiziert, bestätigt oder empfiehlt kein Dritt-Werk, keine Person, Organisation, kein Produkt, keine Dienstleistung und keine Interpretation, die auf das Framework Bezug nimmt, es anwendet oder fortführt. Das Ausbleiben eines Widerspruchs gegen ein Dritt-Werk stellt keine Billigung dar.

§32 Keine Aufsicht und keine Kontrolle. Der Autor beaufsichtigt, leitet, verwaltet, prüft oder kontrolliert keine Dritt-Tätigkeit im Zusammenhang mit dem Framework. Keine Dritt-Tätigkeit wird unter der Autorität, Aufsicht oder Qualitätskontrolle des Autors durchgeführt.

§33 Eigenverantwortung der Dritten. Jeder Dritte, der das Framework anerkennt, zitiert, übernimmt, anwendet, erweitert oder fortführt, handelt unabhängig und trägt die alleinige und ausschließliche Verantwortung für sein eigenes Verhalten, seine Werke, Aussagen und für alle daraus entstehenden Folgen. Eine solche Verantwortung geht nicht auf den Autor über.

§34 Keine Gewähr für Dritt-Werke. Der Autor gibt keine Gewähr, Garantie, Zusicherung oder Versicherung hinsichtlich eines Dritt-Werks, einschließlich seiner Richtigkeit, Vollständigkeit, Sicherheit, Rechtmäßigkeit, Eignung für irgendeinen Zweck oder seiner Ergebnisse. Dritt-Werke werden vollständig auf Risiko derjenigen genutzt und herangezogen, die sie erstellen oder verwenden.

§35 Zitation begründet keine Verpflichtung. Zitation, Attribuierung oder Bezugnahme auf den Korpus — in jeder Form, ob nach einer veröffentlichten Zitations-Konvention oder nicht — begründet keinen Vertrag, keine Sorgfaltspflicht, kein Treueverhältnis und keinerlei Verpflichtung zwischen dem Autor und der zitierenden Partei oder einer anderen Person.

§36 Unterscheidung von Korpus und Feld. Autorschaft und Verantwortung des Autors erstrecken sich nur auf den kanonischen Korpus, wie er vom Autor veröffentlicht wurde. Ein Forschungsfeld ist eine eigentumslose Domäne. Unabhängige Werke, die im Feld entstehen, sind nicht Teil des kanonischen Korpus und sind dem Autor nicht zurechenbar, von ihm nicht abgedeckt und nicht in seiner Verantwortung.

§37 Fortführung erzeugt unabhängige Werke. Jede Fortführung, Erweiterung oder Weiterentwicklung des Gegenstands ergibt Werke, die von der fortführenden Partei verfasst sind. Solche Werke sind deren eigenes geistiges Erzeugnis und deren Verantwortung. Sie sind keine bearbeiteten Ausgaben des kanonischen Korpus und dürfen nicht als der kanonische Korpus oder als Teil davon dargestellt werden.

§38 Keine Haftung für nachgelagerte oder abgeleitete Tätigkeit. Der Autor haftet nicht für Tätigkeiten, Entscheidungen, Anwendungen, Produkte, Dienstleistungen oder Folgen, die eine Partei aus dem Framework oder seinem Korpus ableitet, darauf stützt oder damit verbindet. Der deskriptive, nicht-anleitende Charakter des Korpus (§1–§5) gilt vollumfänglich für alle solchen nachgelagerten Tätigkeiten.

§39 Keine Vertretung, Partnerschaft oder gemeinsame Unternehmung. Die Befassung mit dem Framework durch eine Partei begründet keine Vertretung, Partnerschaft, gemeinsame Unternehmung, kein Beschäftigungs- oder Repräsentationsverhältnis und keine Verbundenheit zwischen dieser Partei und dem Autor. Keine Partei ist berechtigt, im Namen des Autors zu handeln, zu sprechen oder Verträge zu schließen.

§40 EU-AI-Act-Status — kein KI-System, kein Anbieter, kein Betreiber, kein GPAI-Modell. Das AUGMANITAI-Forschungsprogramm und die NEOMANITAI Über-Wissenschaft sind deskriptiver Forschungs-Output. Sie sind keine „KI-Systeme“ im Sinne von Art. 3 Abs. 1 der Verordnung (EU) 2024/1689 (keine maschinengestützten Systeme, die Ausgaben erzeugen, welche physische oder virtuelle Umgebungen beeinflussen). Der Autor ist kein „Anbieter“ (Art. 3 Abs. 3), kein „Betreiber“ (Art. 3 Abs. 4) und kein Anbieter eines „KI-Modells mit allgemeinem Verwendungszweck“ (Art. 3 Abs. 63). Das Programm nutzt KI als Forschungsinstrument (§12), entwickelt, produziert, vertreibt oder betreibt jedoch keine KI-Systeme. Verweise auf KI-Systeme innerhalb der Inhalte sind deskriptive Beobachtungen von Drittsystemen — niemals regulatorische Interpretation. Es wird keine regulatorische Beratung erteilt. Betreiber von KI-Systemen sind für ihre eigene EU-AI-Act-Konformität verantwortlich und müssen qualifizierten Rechtsrat einholen.

COVER PAGE 2 OF 4 — STANDARD BOUNDARY CLAUSES (BLOCK A, MANDATORY)

Cited verbatim from the AUGMANITAI Standard Boundary Clause Set (2026-05-01), Block A1–A5, mandatory in every output.

A1 — Methodologically Descriptive (No Advice / Therapy Character)

This work is methodologically descriptive. It documents observations, terminology, and conceptual framing from a single-author research programme. It does not constitute advice, recommendation, therapy, diagnosis, instruction, or service offering of any kind. Where domain-specific terms are used (clinical, financial, legal, regulatory), they are used in a strictly descriptive sense and not as professional opinion.

A2 — No Personally Identifiable Data

No personally identifiable data of third parties is processed, stored, or referenced. Where corpora are mentioned (e.g. interaction logs, turn counts), these refer exclusively to the author's own work and contain no third-party PII. References to public figures are limited to information already in the public record, used in a neutral nominative manner.

A3 — No Operational Trade-Secrets

This work deliberately separates two layers: (a) the conceptual layer — existence, thesis, terminology, philosophical anchoring — presented for prior-art purposes; and (b) the operational layer — methodology, measurement instruments, scoring architecture, pricing, pipelines, validation protocols — retained as proprietary research material and not disclosed. Absence of operational detail is intentional, not an omission.

A4 — License CC-BY-NC-ND-4.0

Licensed under Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International. Non-commercial use with attribution permitted. No derivatives, no commercial use without separate written permission of the author. License of Clarity ® trademark (EUIPO 019206780) governs trademark use.

A5 — Disclaimer-Reference V6-FINAL §1–§40

This work is governed by the AUGMANITAI Disclaimer V6-FINAL (§1–§40), which establishes its descriptive, non-medical, non-therapeutic, non-diagnostic, non-advisory, non-commercial character. Reading the full disclaimer is prerequisite to interpreting this work.

B-AI-SAFETY (Universal Module, Mandatory for All Outputs)

This work positions itself as complementary to, not in competition with, established AI-safety frameworks including Constitutional AI, RLHF-based alignment, and policy-aligned safety research. References to existing approaches are nominative-neutral. No claim is made that any commercial AI system is unsafe, defective, or trained in a particular manner; statements about model behavior are reported as observations from the author's own research interactions and are framed hypothetically.

B-EDUCATION (Universal Module, Mandatory)

This work does not reference any specific student, class, course, school, or educational institution. Pedagogical concepts are presented at the abstract or methodological level. Age-appropriate boundaries apply: nothing in this work is directed at minors as instructional, motivational, or behavioral content. The author writes as a private researcher, not in any institutional or pedagogical capacity.

B-CROSS-CULTURAL (Universal Module)

No cultural group, nation, ethnicity, religion, language community, or tradition is hierarchized, evaluated comparatively, or assigned essentialist characteristics.

Cross-cultural references serve as descriptive variation, not as ranking.

COVER PAGE 3 OF 4 — LEGAL PARANOID 9-VECTOR SHIELD + IMPRESSUM

LEGAL PARANOID 9-VECTOR SHIELD (binding — short form)

1. **UWG §4 Nr. 1 (Competition Law / No Disparagement).** No evaluative claims about identifiable AI products, providers, or models. Model/brand names appear only in nominative-neutral citation, never in the context of weaknesses or critique.
2. **Professional-Role Context.** The author's primary professional role, employer, or institutional name is never mentioned in public outputs. The author writes in private capacity (Privatperson) for academic research purposes.
3. **Authorship Claims on Neologisms.** Never absolute ("first to discover", "coined by me"). Always relative ("introduced in my work and deposited on Zenodo"). Priority is secured via DOI and Prior-Art-Manifest, not verbal absolutes.
4. **No Institutional Voice.** No team / institutional language. Author writes as single private researcher.
5. **No Recommendation Character.** No "should", no health promises. Descriptive framing only.
6. **Press-Law Due Diligence (Impressum §18 Abs. 2 MStV).**
Responsible: Andreas Ehstand (address below).
7. **GDPR Art. 13.** No third-party PII processed, stored, or referenced.
8. **No Training-Architecture Claims.** Statements about commercial AI systems' internal training are modalized as hypotheses, never claimed as facts.
9. **Trademark Mentions.** Only nominative, never in critical context.

IMPRESSUM (DDG §5 / §18 Abs. 2 MStV)

Andreas Ehstand Nepomukweg 7 · 82319 Starnberg · Deutschland E-Mail: augmanitai [at] gmail [dot] com **Telefonisch nicht erreichbar (Forschungs-Privatperson, kein Geschäftsbetrieb). Verantwortlich i.S.d. §18 Abs. 2 MStV:** Andreas Ehstand (Anschrift wie oben). **USt-IdNr.:** keine (keine unternehmerische Tätigkeit i.S.d. §2 UStG). **Berufsbezeichnung im Sinne der Forschungs-Publikation:** Independent Researcher.

Hinweis: Diese Veröffentlichung erfolgt in privater Eigenschaft als wissenschaftliche Forschungs-Publikation. Sie ist nicht Teil einer beruflichen oder institutionellen Tätigkeit.

COVER PAGE 4 OF 4 — REFINEMENT-WINDOW + TRADE-SECRET-LAYER

REFINEMENT-WINDOW-KLAUSEL

This work is published as a “Working Notes” version under the AUGMANITAI Refinement-Window protocol: for **30 days from publication**, metadata (Title, Description, Keywords, Notes) may be edited via the Zenodo Refinement-Window without changing the DOI, the file SHA-256 hashes, the Bitcoin-OTS prior-art anchor, or the OpenTimestamps verification chain. Substantive content edits trigger a new version (Newversion) with V1 preserved in the Version History.

TRADE-SECRET-LAYER-ARCHITEKTUR

Per **§28 Trade-Secret Reservation** (above), the author maintains a strict three-layer architecture:

Layer	Content	Access
(a) PUBLIC	Descriptive concepts, methodological frames, terminology	CC-BY-NC-ND-4.0
(b) RESTRICTED	Substantiation, technical detail, application architectures	Request via ORCID, legitimate research purpose, written confidentiality undertaking
(c) HARD-SECRET	Operational mechanisms, scoring algorithms, proprietary processes	Internal — not deposited, not transferable

Access requests must reference: (1) ORCID profile of requestor, (2) institutional affiliation if any, (3) specific research purpose, (4) confidentiality undertaking declaration. The author reserves the right to decline access for any reason without justification (research-program privacy).

PRIOR-ART ANCHOR CONFIRMATION

- **Bitcoin Block:** 945970+945979 (2026-04-20T21:08Z)
 - **Manifest-Root SHA-256:**
299e4b3d0ac9e50740268896b5eeb541f6462332bb67b7efdf4cd309704770d0
 - **OTS-Calendars:** alice.btc.calendar.opentimestamps.org,
bob.btc.calendar.opentimestamps.org, finney.calendar.eternitywall.com,
btc.catallaxy.com (4 independent witnesses)
 - **Multi-Hash:** SHA-256, SHA-512, SHA-3-256, SHA-3-512, BLAKE3
 - **Local Hash-Registry:** MULTI_HASH_REGISTRY_*.json in
10_RECHTLICHES/PRIOR_ART_TIMESTAMPS/
-

AUGMANITAI / NEOMANITAI Disclaimer V6-FINAL — §1–§40 — 18 May 2026
— Andreas Ehstand — ORCID: 0009-0006-3773-7796 — CC BY-NC-ND 4.0.
Living document; earlier versions remain valid in parallel.

END OF DOCUMENT
