ULOGIC-MIND
Long-Term Plan and Vision

SECTION 1: What is a "Universal Linguistic Mind" (UMIND) and what is a "Universal Logic Language" (ULOGIC)?

We will call a **Universal Linguistic Mind** (abbreviated **UMIND**) a system that possesses the capabilities listed below. This enumeration of capabilities offers us a roadmap of what we must achieve (at a minimum!) to attempt to approach a supposed "AGI":

a. **Discursive Capability:** It is capable of handling human natural language to create stories, narratives, and tell tales, respecting grammatical correctness and the semantic coherence of the discourse. It is speaking as in novels, as in a speech, or as in a book. At this level, the capability for rigorous reasoning, whether mathematical or logical, is not achieved, nor is the ability to define and execute algorithms.

b. **Discursive-Reasoning Capability:** Informal arguments of plausibility, similar to discourses linking phrases like "if" and "then" and using basic and-or-implies connectors. However, at the discursive level, the premises are not clear, the rules are not explicit, and the reasoning is not truly rigorous. They are based on similarity, trust, and the power of conviction due to the content of what is said (repetition and emotionality have an influence). At this level, it is impossible to do mathematics, complex reasoning, or science.

c. **Logical-Reasoning Capability:** Utilization of a strict internal language (ULOGIC), with strict rules, where reasoning and derivations are verifiable and exact. Mathematics is a progressive invention over centuries where these rules have been created, implicitly and vaguely. The objective is to use an explicit and defined linguistic system (which we will call ULOGIC).

d. **Semantic Algorithmic Capability:** Within ULOGIC, algorithms and processes can be defined, and it is possible to execute them. And to reason about the execution, about why it is correct or incorrect. In reality, it must be an extension of logical capability. Algorithms should be seen as an extension of logical rules; they are "mathematics," and that is why one can reason about algorithms, because they are an extension of mathematics-logic. The term "semantic algorithm" means just that: the execution of an algorithm is actually "performing derivations and demonstrations" following the instructions of the algorithm.

e. **Metalinguistic Capability:** The ULOGIC language itself is capable of talking about expressions, grammatical rules, derivations, and algorithms. To be able to model itself as an internal mathematical structure of the language itself.

f. **Complete Logical-Mathematical Capability:** The ULOGIC language is sufficient to develop the complete arguments of advanced mathematics books, expressed in the internal structure, with a rigorous validator that checks the derivations and application of rules.

g. **Closed Semantic Capability:** The formal-internal language (ULOGIC) of UMIND is NOT interpreted in "external semantics" as occurs with FOL, SOL, HOL, or CIC. Nor does it follow the "model theory" paradigm. When we talk about mathematics in the real world, or create algorithms or processes, our words "are not interpreted in another language": we use a single, self-contained language, whose semantics are internal without reference to anything external (with the exception of the words and labels used in perception, but that is another field).

h. **Mathematical Foundational Capability:** Additionally, it must be able to reconstruct a set theory free of contradictions and close to natural intuition. Zermelo's set axioms will cease to be useful and will merely be "the axioms of a structure" as the structure of groups or rings might be. The system will have the capability to define what a set is. And above all: to resolve set-theoretic contradictions.

In conclusion: UMIND is a system capable of handling natural language in a "discursive" form and for "discursive-arguments" as LLMs already do, but it will additionally use a novel language, ULOGIC, with exact grammar and rules, which allows for exact mathematical derivations, reasoning, and algorithms, has metalinguistic capability, and will additionally offer a new foundation for mathematics, where set axioms are not necessary, where the notion of a set is defined, and the contradictions of intuitive set theory are resolved from a new perspective.

The implementation strategy consists of combining LLM techniques with the novel ULOGIC language to create a new foundational paradigm of artificial intelligence with complete, verifiable, and exact linguistic reasoning capability. On one hand, LLMs make heuristic proposals to generate proofs and algorithms, and on the other hand, a Hard-Kernel is a verifier of indubitable correctness and exact validation.

SECTION 2: State of the Art in LLMs and their Limitations in Logical Reasoning

Large Language Models (LLMs), such as GPT-4, Claude, Gemini, and Llama, have demonstrated extraordinary capabilities in the processing and generation of natural language. These models, based on Transformer architectures and trained on vast corpora of text and even multimodal data, excel at tasks that require linguistic fluency, contextual understanding, and coherent text generation. Specifically, they have reached a notable level in what we have defined here as (a) Discursive Capability and (b) Discursive-Reasoning Capability.

Discursive and Discursive-Reasoning Capabilities in LLMs:

Current LLMs can generate complex narratives, maintain coherent conversations on various topics, summarize extensive texts, and answer questions informatively. They are capable of imitating writing styles, from literary prose to technical reports, largely respecting grammatical correctness and local semantic coherence. Evaluations and benchmarks like COHESENTIA attempt to measure this coherence, distinguishing between the logical connection of adjacent sentences (local coherence) and the thematic unity of the complete text (global coherence). LLMs demonstrate strong local coherence, but maintaining global coherence in long texts remains a challenge.

Regarding discursive-reasoning, LLMs can construct informal arguments, chain ideas using basic logical connectors ("if... then," "and," "or"), and generate plausible explanations for various phenomena. They can follow complex instructions and participate in dialogues that simulate thought processes, as evidenced in techniques like Chain-of-Thought (CoT). These capabilities make them useful tools for tasks such as drafting, brainstorming, and writing assistance. However, this form of reasoning is based more on the recognition of statistical patterns and superficial plausibility than on rigorous logical inference.

Fundamental Limitations in Logical-Mathematical Reasoning:

Despite their linguistic successes, LLMs exhibit fundamental and persistent limitations in Logical-Reasoning and Semantic Algorithmic Capability as defined for UMIND. Numerous studies and benchmarks have documented these deficiencies.

- **Lack of Rigor and Verifiability:** The reasoning of LLMs often lacks formal rigor. Although they can generate steps that *seem* logical, they do not operate on a system of explicit and verifiable rules. Their inferences are based on correlations learned from training data, not on guaranteed logical deduction. This leads to subtle but critical errors, especially in complex problems that require multiple steps of deduction. The evaluation of LLMs is often limited to the correctness of the final answer (Accuracy), without guaranteeing the validity of the reasoning process (Rigorousness).

- **Logical Inconsistency:** LLMs are prone to generating contradictory answers, violating basic logical principles such as consistency of negation, transitivity, or factuality. They can affirm A and not-A in slightly different contexts, or fail in simple transitive inferences (if A implies B and B implies C, does A imply C?). This inconsistency undermines their reliability for tasks requiring logical precision.

- **Fragility to Perturbations:** The performance of LLMs on reasoning tasks is surprisingly fragile to small alterations in the input that do not affect the underlying logical structure. Changes in the wording of a problem, the order of premises, the inclusion of irrelevant information, or the modification of numerical values can significantly degrade their accuracy. This suggests that LLMs learn superficial heuristics and specific patterns from the training set instead of understanding the deep structure of the problem. A "curse of complexity" has been observed, where accuracy decreases dramatically as the complexity of the problem increases, even if the basic steps are the same.

- **Arithmetic and Strategic Planning Errors:** In mathematical problems, even at the high school level, LLMs make not only arithmetic errors (especially with large numbers or those outside the training distribution) but also failures in spatial reasoning, strategic planning, and the translation of physical intuition into formal mathematical steps. They can make unjustified assumptions or rely excessively on superficial numerical patterns. Accuracy in arithmetic calculations decreases significantly when they are embedded within word problems rather than presented as isolated tasks.

- **Difficulty with Abstraction and Generalization:** LLMs struggle to generalize abstract logical principles to new domains or problems with superficially different structures from those seen during training. Their dependence on probabilistic "pattern matching" over formal logical reasoning limits their capacity for genuine inference.

- **Hallucination Phenomenon:** LLMs can generate incorrect or nonsensical information with great confidence, a phenomenon known as "hallucination." This is particularly problematic in reasoning tasks where truthfulness and accuracy are crucial.

Current Efforts to Integrate Logic into LLMs:

Aware of these limitations, the research community is actively exploring various strategies to equip LLMs with more robust logical reasoning capabilities. Research teams at leading laboratories such as OpenAI, Google DeepMind, Meta AI, and Anthropic are working in this direction. The main strategies can be classified as:

- **Improving Prompting Techniques and Inference Scaling:**
    - **Chain-of-Thought (CoT) and its Variants:** CoT and its extensions like Tree-of-Thoughts (ToT), Graph-of-Thoughts (GoT), and structures like RATT aim to guide the LLM to generate explicit intermediate reasoning steps.
    - **Inference-Time Scaling:** Techniques like generating multiple reasoning paths and using voting (self-consistency) or beam search to select the best answer. This involves allocating more computational resources during response generation.
    - **Program-of-Thought (PoT):** Generating executable code (e.g., Python) as an intermediate step to perform precise calculations or verifications.

- **Use of Verifiers and Feedback:**
    - **External Verifiers:** Training separate models (verifiers) to evaluate the correctness of the reasoning steps or final answers generated by the LLM. These can be reward models (Reward Models - ORMs, PRMs) or preference-based models (DPO, SimPO). Math-Rev and Code-Rev are examples of verifiers trained on datasets of correct and incorrect solutions.
    - **Environment Feedback:** Using compilers or interpreters to verify the correctness of generated code (in PoT) or proof assistants to validate steps in formal proofs.
    - **Self-Critique / Self-Refinement:** Methods where the LLM critiques or refines its own outputs. Panel, for example, uses self-generated critiques in natural language as feedback to guide the search at the step level, instead of scalar reward signals. CoTnPoT combines the interpretability of CoT with the precision of PoT for more robust verification.

- **Integration with Formal and Symbolic Systems (Neuro-Symbolic):**
    - **LLM as an Interface for Solvers:** Using LLMs to translate problems from natural language to formal languages (like FOL, SMT, ASP) that can be processed by external symbolic solvers. The LLM interprets the solver's output back into natural language.
    - **LLM Guided by Symbolic Knowledge:** Incorporating knowledge bases or logical rules to guide or constrain the LLM's generation.
    - **Neuro-Symbolic Learning:** Developing hybrid architectures that integrate neural and symbolic modules more deeply, sometimes allowing for end-to-end training. This can include the use of symbolic feedback during training.

- **Automatic Formalization (Autoformalization):** Using LLMs to translate informal mathematical statements (in natural language) into formal languages used by proof assistants like Lean, Coq, Isabelle/HOL. This is a crucial step for applying LLMs to the formal proof of theorems.
- **Tactic Generation and Proof Search:** Training LLMs to suggest the next steps (tactics) in a formal proof within a proof assistant or to guide the search in the proof space. Large datasets of formal proofs (e.g., in Lean) are being created to train these models.
- **Improving Training (Pre-training and Fine-tuning):**
  - **Augmented Data:** Training LLMs with enriched datasets that include examples of step-by-step reasoning, logical proofs, or solved mathematical problems.
  - **Supervised Fine-tuning (SFT) and Reinforcement Learning (RL):** Adjusting pre-trained LLMs using SFT on high-quality reasoning data and then refining with RL using rewards based on the correctness or quality of the reasoning (as in DeepSeek-R1).

Despite these efforts, achieving robust, verifiable, and generalizable logical-mathematical reasoning comparable to human capability (or to the capabilities described for UMIND/ULOGIC) remains a fundamental challenge for the current LLM paradigm.

SECTION 3: Impact of UMIND/ULOGIC as a New Paradigm and Limitations of Current Approaches

The concept of a Universal Linguistic Mind (UMIND) powered by a Universal Logic Language (ULOGIC) represents a fundamental paradigmatic shift in Artificial Intelligence, with the potential to overcome the inherent limitations of current approaches, including pure LLMs and existing neuro-symbolic systems.

Differentiation and Potential Impact of UMIND/ULOGIC:

UMIND/ULOGIC differs from current paradigms in several crucial aspects:

1. **Advanced Neuro-Symbolic Architecture:** The main differentiation lies in its explicit neuro-symbolic architecture. It is not a monolithic system, but a **synergy between two specialized components: a Large Language Model (LLM) that acts as a natural language interface and heuristic proposal generator, and a Verifier-Kernel that implements ULOGIC.** The superiority of the approach does not lie in an "intrinsic

integration" that blurs capabilities, but in the power of the kernel's language. While current neuro-symbolic approaches couple LLMs with verifiers based on formal languages like Lean or FOL, the UMIND kernel uses ULOGIC, a language with far superior expressive power and foundational capabilities (meta-linguistic, closed semantics). It is this interaction between a flexible LLM and a logical kernel of unprecedented power that defines the new paradigm.

2. **Formal Language with Superior Capabilities (ULOGIC):** The core of the proposal is ULOGIC, a postulated formal language with capabilities that go beyond standard formal languages (FOL, HOL, Type Theory, etc.). These capabilities include:
   - **Semantic Algorithmic Capability:** To define, execute, and *reason about* algorithms as extensions of the internal logic.
   - **Metalinguistic Capability:** ULOGIC can describe and reason about itself (its expressions, rules, derivations).
   - **Closed Semantic Capability:** The semantics of ULOGIC are intrinsic, self-contained, without depending on interpretations in external models, reflecting how humans seem to use a single language to reason about abstract domains like mathematics.
   - **Mathematical Foundational Capability:** ULOGIC is capable of *defining* concepts like "set" and reconstructing mathematical foundations without depending on external axiomatic systems like ZFC, potentially resolving inherent paradoxes.

3. **Verifiability and Accuracy:** Unlike the often opaque and error-prone reasoning of LLMs, the derivations and algorithmic executions within ULOGIC would be strict, verifiable, and exact, thanks to its explicit rules and validation kernel.

4. **New Foundation for AI:** UMIND/ULOGIC is not just an incremental improvement but proposes a new foundation for AI centered on a unified and sufficiently powerful logical-linguistic language. This could lead to AI systems with a deeper understanding, capable of performing advanced science, solving complex mathematical problems, and reasoning about their own knowledge and processes in a way that current systems cannot.

The impact of such a system would be transformative. It could drastically accelerate scientific discovery, enable the creation of automatically verified complex software, solve large-scale engineering problems, and offer much more reliable and explainable AI systems.

Insufficiency of Current Formal Languages:

The central claim of the UMIND/ULOGIC proposal is that current attempts to endow LLMs with logic are fundamentally limited by the insufficiency of the formal languages they use.

- **Propositional Logic (PL):** It is too simple, incapable of representing objects, properties, or relations, nor of using quantifiers. It only handles true or false atomic propositions.
- **First-Order Logic (FOL):** Although more expressive than PL by allowing predicates, functions, and quantification over objects, FOL has significant limitations:
  - It cannot quantify over predicates or functions, which makes it difficult to express higher-order properties common in mathematics and natural language.
  - It struggles to naturally represent concepts like beliefs, knowledge, or modalities (necessity, possibility).
  - Its model-based semantics (interpretations in external structures) differs from ULOGIC's concept of closed semantics.
- **Higher-Order Logics (HOL):** They extend FOL by allowing quantification over predicates and functions, making them more expressive and suitable for formalizing much of mathematics. Systems like Coq, Isabelle/HOL, and Lean are based on variants of HOL or related type theories. However:
  - They can become extremely complex.
  - The semantics are still generally model-based.
  - They do not intrinsically possess the self-referential metalinguistic capabilities or the foundational capability (defining sets from scratch) postulated for ULOGIC.
  - They remain formal languages distinct from discursive natural language, requiring a translation (autoformalization) that is itself a challenge for LLMs.
- **Type Theory (TT):** Used in many proof assistants (Coq, Lean) and in formal semantics of natural language (Montague Grammar). It classifies expressions into types to avoid paradoxes and ensure correctness. Rich type theories (with dependent types, etc.) offer great expressiveness. Nevertheless:
  - Capturing the full flexibility and ambiguity of natural language remains a challenge. The rigidity of types can clash with the fluidity of language.
  - Representing gradual, plausibility-based, or context-based reasoning is difficult.
  - Like HOL, they lack the closed semantics and the self-referential foundational and metalinguistic capabilities of ULOGIC in the described form.
  - Learning these complex semantic structures from primary linguistic data is an open problem.

In summary, current formal languages, although powerful in their domains, do not possess the unique combination of complete logical-mathematical expressiveness, semantic algorithmic capability, self-referential metalinguistics, closed semantics, and foundational capability attributed to ULOGIC.

SECTION 4: Advanced Capabilities of UMIND/ULOGIC

A Universal Linguistic Mind (UMIND) equipped with the Universal Logic Language (ULOGIC) would transcend the capabilities of current AI systems, enabling an unprecedented level of reasoning, creation, and self-understanding. The advanced capabilities attributed to this combination are as follows:

- **Performing Complex Mathematics:** UMIND/ULOGIC could not only *follow* mathematical proofs but actively *develop* them. The Complete Logical-Mathematical Capability (f) postulates that ULOGIC is expressive enough to formalize the arguments found in advanced mathematics books. This implies the ability to handle complex definitions, intermediate lemmas, proofs by induction, reduction to absurdity, and any other technique, all within the internal structure of ULOGIC and rigorously validated by an internal kernel. This contrasts with current LLMs, which can generate mathematical text but often fail in the underlying logical validity or multi-step complexity.

- **Defining, Executing, and Reasoning about Semantic Algorithms:** The Semantic Algorithmic Capability (d) places algorithms as a natural extension of logic within ULOGIC. This means UMIND could:
    - **Define algorithms:** Express computational procedures precisely using ULOGIC's syntax.
    - **Execute algorithms:** Execution is understood as a form of step-by-step "derivation" or "proof" following the algorithm's rules, maintaining traceability and logical verifiability.
    - **Reason about algorithms:** Formally analyze the properties of algorithms (correctness, termination, complexity) using ULOGIC's logical capabilities. This includes being able to explain *why* an algorithm is correct or incorrect. This capability is fundamental for reliable software engineering and for an AI that can understand and optimize its own processes, overcoming the difficulty of LLMs in guaranteeing code correctness or reasoning formally about it.

- **Deep Metalinguistic Capability:** The Metalinguistic Capability (e) is one of the most distinctive features. ULOGIC can be used to talk *about itself*: about its own expressions, its grammar, its inference rules, and the algorithms defined in it. This allows UMIND to:
    - **Self-model:** Represent itself as an internal mathematical structure.
    - **Self-reflection and Self-improvement:** Reason about its own reasoning processes, identify potential errors or inefficiencies, and potentially modify or improve its strategies.
    - **Explainability:** Generate deep and formally grounded explanations of its own derivations and conclusions. This capability is crucial for the trust and debugging of complex systems and goes beyond the post-hoc or superficial explanations that current LLMs can offer.

- **Closed and Self-Contained Semantics:** The Closed Semantic Capability (g) postulates that ULOGIC does not need to be interpreted in external "models" or "worlds" to have meaning, unlike the standard semantics of FOL or HOL (model theory). Its meaning derives from the internal relationships between its expressions and rules. This more closely resembles how humans seem to operate with abstract concepts like mathematics, where the language used is self-contained. This property could be key to avoiding certain types of ambiguity or dependence on external context that affect current systems and allow for more robust and self-consistent reasoning.

- **Refounding Mathematics and Resolving Paradoxes:** The Mathematical Foundational Capability (h) is the most ambitious. UMIND/ULOGIC would not only use existing mathematics but could *reconstruct* it from a new foundation. Specifically:
    - **Emergent Set Theory:** The notion of a "set" would not be primitive or defined by external axioms (like those of Zermelo-Fraenkel, ZFC), but would *emerge* as a definable construction within the ULOGIC language itself.
    - **Overcoming External Axioms:** The axioms of ZFC, which are the standard basis of modern mathematics but present certain complexities and philosophical problems (Skolem's paradox, etc.) or limitations, would become unnecessary as an ultimate foundation. They would simply be the axioms that define a particular structure (the "ZFC set structure"), similar to the axioms of a group or ring.
    - **Resolution of Paradoxes:** ULOGIC offers a new perspective to resolve the paradoxes that plagued intuitive set theory (like Russell's paradox) and motivated the creation of ZFC. The key is to understand that definitions ceased to be "eliminable abbreviations" when set theory appeared, and that the necessary logical rules are many and more complex than what logic has brought to light so far. The history of logic is an attempt at reverse engineering to figure out the

structure of the languages we use (and this reverse engineering has not yet been solved).

Together, these capabilities would paint a picture of a fundamentally different AI: one that not only imitates linguistic patterns or performs specific calculations but possesses an integrated and verifiable capability for language, logic, mathematics, and self-reflection, laying the groundwork for a truly general and comprehensive artificial intelligence.

SECTION 5: TekDocs as a Universal Knowledge Store

Within the UMIND/ULOGIC paradigm, **TekDocs** (Transportable Encapsulated Knowledge Documents) constitute the fundamental mechanism for storing, organizing, and reusing the formal and algorithmic knowledge generated and used by the system. They are much more than simple data files; they are structured expressions within the ULOGIC language itself.

Nature and Content of TekDocs:

A TekDoc is a self-contained unit of knowledge expressed in ULOGIC. It can encapsulate a variety of interrelated elements:

- **Premises and Contexts:** The assumptions, definitions, and contextual framework necessary for a specific derivation or algorithm.
- **Logical Derivations:** Sequences of formally valid reasoning steps within ULOGIC, proving theorems or conclusions.
- **Algorithms:** Precise definitions of computational processes, expressed as logical/mathematical structures within ULOGIC.
- **Algorithm Executions:** Detailed and verifiable traces of an algorithm's execution on specific data, understood as a sequence of derivations.
- **Questions and Problems:** Precise formulations of issues to be solved or investigated.
- **Strategies:** Descriptions of plans or heuristics to achieve a goal, which can themselves be seen as high-level algorithms (e.g., a strategy for finding a proof or solving a particular type of problem).

Key Properties of TekDocs:

- **System Expressions:** Crucially, TekDocs are not external to the UMIND/ULOGIC system; they *are* valid expressions within ULOGIC. The entire content of a TekDoc, as well as the TekDoc as a whole, is a formal structure defined in the language. This property is the basis of UMIND's metalinguistic capability: since TekDocs (which represent the system's knowledge and processes) are expressions of ULOGIC, the system can use ULOGIC to "talk about" its own TekDocs, analyze them, modify them, combine them, or reason about their content and structure. This allows for much deeper self-reflection and knowledge management than current systems.

- **Universal and Reusable Store:** TekDocs form a dynamic and constantly growing knowledge base. They can be stored, indexed, and retrieved by UMIND. The formal structure of ULOGIC allows for precise and reliable reuse of knowledge:
    - **Referencing:** A TekDoc can explicitly reference other TekDocs, importing their definitions, premises, or results (theorems, algorithms). This allows for the modular construction of complex knowledge, similar to how it is done in software libraries or mathematical theories.
    - **Composition:** TekDocs can be combined to address more complex problems, ensuring logical coherence through ULOGIC's rules.
    - **Verification:** The validity of any TekDoc or combination of TekDocs can be rigorously verified by the ULOGIC kernel.

- **Granularity and Context:** TekDocs can vary in size and complexity, from the definition of a simple concept to the formalization of a complete mathematical theory or a complex software system. Each TekDoc encapsulates its own context (premises, imported definitions), ensuring that the knowledge is self-contained and its applicability is clearly defined.

Function in the UMIND/ULOGIC Ecosystem:
TekDocs would act as the formal and operational "memory" of UMIND. They would be the medium through which:

- Mathematical and algorithmic knowledge is stored.
- Derivations and execution results are recorded.
- Knowledge and software components are shared and reused.
- The system reasons about its own state and capabilities (by analyzing its own TekDocs).
- Structured knowledge bases, including those for common sense, are built (see Section

In essence, TekDocs provide the necessary knowledge infrastructure for the advanced capabilities of ULOGIC (logical, algorithmic, metalinguistic, foundational) to manifest in an organized, scalable, and verifiable way.

SECTION 6: Connecting ULOGIC with World Knowledge and Common Sense

While ULOGIC is postulated as an internally rigorous and self-contained formal language for logic, mathematics, and algorithms, its ultimate utility as a basis for a general AI (UMIND) depends on its ability to connect with the vast informal knowledge contained in natural language and to represent aspects of the real world. The UMIND/ULOGIC paradigm envisions specific mechanisms to achieve this connection.

1. Creation of Common Sense Bases using TekDocs:
Common sense knowledge – the vast set of implicit facts, causal relationships, social norms, and intuitive physical knowledge that humans constantly use – is notoriously difficult to formalize. The proposal is to use the structure of TekDocs (Section 5) to build common sense knowledge bases:

- **Extraction from Existing Sources:** Large volumes of text (books, articles, websites) could be processed using the discursive capabilities of UMIND's LLM component to identify conceptual relationships, facts, and heuristic common sense rules.
- **Formalization in ULOGIC:** These extracted relationships and facts would then be formalized as expressions within TekDocs, using ULOGIC's syntax. For example, a statement like "birds usually fly" could be represented using predicates and quantifiers (possibly non-standard or probabilistic if ULOGIC supports them) within a TekDoc dedicated to knowledge about animals.
- **Interconnected Structure:** Sets of interrelated TekDocs would be created that express relationships between concepts (e.g., a TekDoc about "water" could link to TekDocs about "liquid," "drinking," "ice," "vapor"). This would form structured knowledge graphs formalized in ULOGIC.
- **Reuse and Reasoning:** UMIND could then use these common sense TekDocs as premises or context in its logical reasoning processes, allowing it to perform inferences that go beyond purely mathematical logic and approach everyday human reasoning. The verifiable nature of ULOGIC would allow for more robust common sense reasoning that is less prone to the inconsistencies of current knowledge handling by LLMs.

2. Extension of ULOGIC to Represent All Natural Language:

A crucial step is the hypothesis that ULOGIC not only serves for mathematics and logic but is also *extensible* to be able to formally express the semantics of *all* natural language, not just statements with explicit logical or mathematical content.

- **Beyond Traditional Formal Logic:** This implies that ULOGIC (extension-v2) would need constructs capable of representing the enormous variety of nuances in natural language: modality, temporality, causality, beliefs, intentions, vagueness, metaphors, etc. Current formal languages like FOL or even HOL and Type Theories struggle to capture all this richness naturally and completely.
- **A Universal Language for Meaning:** If ULOGIC (extension-v2) could achieve this extended expressiveness, it would become a universal semantic *interlingua*. Any sentence in natural language could, in principle, be translated into a formal expression in ULOGIC that captures its meaning precisely and unambiguously.
- **Capability to "Talk about the World":** By being able to formally represent statements about objects, events, properties, and relationships in the world described in natural language, ULOGIC would endow UMIND with the ability to "talk about the world" in a structured and logically coherent way. This would allow UMIND not only to process text but to *understand* its meaning at a deeper level and use that knowledge in rigorous reasoning.

Implications:

The combination of common sense knowledge bases formalized in TekDocs and an extended ULOGIC (extension-v2) to encompass all of natural language would allow UMIND to overcome one of the biggest gaps in current AI: the connection between superficial linguistic processing and deep, knowledge-based reasoning. UMIND could:

- Read and understand texts with much greater semantic depth.
- Integrate knowledge from diverse sources coherently.
- Perform complex reasoning that combines formal knowledge (mathematical, algorithmic) with common sense knowledge and facts about the world.
- Generate natural language that is not only fluent but also logically grounded and consistent with an internal knowledge base.

This connection between the logical formalism of ULOGIC and the richness of natural language and world knowledge is essential for the vision of UMIND as a truly general and comprehensive artificial intelligence.

SECTION 7: Integration with Perceptual Systems and Reasoning about Experience

For an artificial intelligence to interact meaningfully with the real world, it is not enough to process language and abstract knowledge; it must be able to perceive its environment and connect its internal representations with sensory experience.

Internal Perceptual Representations:

The first necessary component is the existence of **perceptual systems** (visual, auditory, tactile, etc.) that capture information from the external world and transform it into **internal representations**. These representations would not be mere raw data, but **structured abstract spaces** (e.g., an internal visual space with topological, metric, or geometric properties; an auditory space with dimensions of pitch, timbre, time).

- **Abstraction and Structure:** These internal spaces would capture the regularities and relationships inherent in sensory data. For example, a visual space could represent objects, their spatial relationships (near, far, inside, outside), shapes, colors, and movements in an organized manner.
- **Topology:** Endowing these internal spaces with topology is crucial. It allows for defining concepts such as continuity, neighborhood, connection, and separation, which are fundamental for coherently segmenting and understanding sensory scenes and events.

Dynamic Labeling with ULOGIC:

The key connection between perception and reasoning is established through the use of **ULOGIC labels** to name or refer to elements within these internal perceptual representations.

- **Naming Continuous Regions:** ULOGIC labels would not necessarily apply to individual "pixels" or "samples," but to significant **"continuous-regions"** within the perceptual spaces. For example, a ULOGIC label like spherical_red_object_1 could be assigned to a specific region of the internal visual space corresponding to a perceived red ball. Another label sharp_short_sound_start could refer to a segment of the internal auditory space.
- **Dynamic Labeling:** This labeling process would be dynamic. As perception changes (the ball rolls, another sound is heard), labels would be assigned, updated, or removed accordingly, reflecting the current state of the perceived environment as represented internally.

- **Symbol-Perception Connection (Grounding):** This mechanism provides a way to "ground" the abstract symbols of ULOGIC in perceptual experience. The label spherical_red_object_1 is not just an abstract symbol but is directly linked to an internal representation generated from real-world sensory data.

Reasoning about Experience:

Once this connection is established, UMIND could use ULOGIC to **reason directly about the labeled perceptual experience**:

- **Talking about What is Perceived:** UMIND could generate descriptions in natural language (or in ULOGIC) about what it is perceiving, using the labels as reference points. For example, it could derive the statement "I see spherical_red_object_1 near cubic_blue_object_1." This is analogous to how humans talk about what we see or hear.
- **Perception-Based Inference:** It could perform logical inferences that combine perceptual information with prior knowledge stored in TekDocs. For example:
    - Perception: Detects spherical_red_object_1 moving towards gray_wall_1.
    - Knowledge (TekDoc): "Solid objects cannot pass through walls."
    - Inference (ULOGIC): "It is likely that spherical_red_object_1 will collide with gray_wall_1."
- **Perception-Based Planning:** It could use perceptual information to plan actions in the real world.

Implications:

This integration of ULOGIC with perceptual systems represents a crucial step towards an embodied AI situated in the real world. It would overcome the current disconnection of many AI systems that operate solely in the abstract domain of language or symbolic data. It would allow UMIND to:

- Understand the world more holistically, integrating abstract and sensory information.
- Interact with physical environments intelligently and adaptively.
- Learn from the real world through direct experience.
- Validate its abstract knowledge against perceptual evidence.

This ability to reason about experience would finally connect the logical and linguistic power of UMIND/ULOGIC with the physical world, an indispensable requirement for general artificial intelligence.

SECTION 8: Modeling Physical Systems and Advanced Science

Once UMIND/ULOGIC has established a robust connection between its internal language, abstract knowledge, and perception of the world (Sections 6 and 7), the next logical step is to apply these capabilities to the fundamental task of science: modeling physical systems and discovering new scientific knowledge.

Utilization of Mathematical Structures for Modeling:

Modern science relies heavily on the use of **mathematical structures** (differential equations, linear algebra, probability theory, differential geometry, etc.) to describe and predict the behavior of physical systems, from subatomic particles to the evolution of the universe. UMIND/ULOGIC would be exceptionally equipped for this task:

- **Complete Mathematical Formalization:** Thanks to its Complete Logical-Mathematical Capability (Section 1f), UMIND could formally represent within ULOGIC the complex mathematical structures used in physics, chemistry, biology, and other sciences. This includes not only the equations but also the underlying definitions, axioms, and theorems.
- **Integration with Perceptual/Experimental Data:** Through the perceptual connection (Section 7), UMIND could link these formal mathematical structures with data obtained from experiments or real-world observations. It could use ULOGIC to label and represent experimental data in a structured way.
- **Model Construction and Validation:** UMIND could actively participate in the scientific process:
    - **Hypothesis Formulation:** Propose mathematical models (expressed in ULOGIC) to explain observed phenomena.
    - **Derivation of Predictions:** Use ULOGIC's logical and algorithmic reasoning capabilities to derive testable predictions from the proposed models.
    - **Comparison with Evidence:** Compare the formal predictions with experimental data (also represented in ULOGIC through perceptual labeling).
    - **Model Refinement:** Use discrepancies between predictions and observations to refine or revise the mathematical models, leveraging its metalinguistic capability to reason about the model structures themselves.

Capabilities for Advanced Science:

The application of UMIND/ULOGIC to science could enable significant advances:

- **Automation of Scientific Reasoning:** Many complex cognitive tasks in science, such as deriving consequences from a theory, checking the consistency of hypotheses, or searching for patterns in large datasets, could be automated or assisted by UMIND, thanks to the combination of rigorous logical reasoning and algorithmic capabilities.
- **Accelerated Discovery:** By being able to explore vast spaces of mathematical hypotheses and quickly compare them with data, UMIND could accelerate the discovery of new physical laws, biological mechanisms, or novel materials. Its ability to handle mathematical complexity that surpasses human intuition would be a key advantage.
- **Verifiable and Reproducible Science:** The expression of theories, models, and data analyses in the formal and verifiable language of ULOGIC (stored in TekDocs) would dramatically increase the reproducibility and reliability of scientific research. Derivations and analyses would be explicit and checkable by the ULOGIC kernel.
- **Interdisciplinary Integration:** ULOGIC could serve as a common formal language to express knowledge and models from different scientific disciplines, facilitating the integration of knowledge and the tackling of complex problems that require multidisciplinary approaches (e.g., climate change, systems biology).
- **New Types of Scientific Theories:** The mathematical foundational capability of ULOGIC (Section 4) could even lead to the formulation of scientific theories based on new mathematical foundations, potentially offering novel perspectives on fundamental problems in physics or cosmology.

Partial Conclusion:

The ability to use ULOGIC to model physical systems and reason about them represents the culmination of UMIND's capabilities. It would transform AI from a tool primarily based on data and patterns into a true collaborator in the process of scientific discovery, capable of handling mathematical abstraction, rigorous logic, and the connection with the empirical world in an integrated and powerful way. This would not only validate the UMIND/ULOGIC paradigm but could revolutionize the very practice of science.

SECTION 9: UMIND/ULOGIC as a Necessary Foundational Paradigm

The argument developed in the previous sections converges on a fundamental conclusion: the UMIND/ULOGIC paradigm does not represent a mere incremental improvement over existing Artificial Intelligence systems, but a **new foundational paradigm** necessary to achieve an AI with advanced linguistic thinking capabilities, robust logical-mathematical reasoning, and a deep understanding of the world.

Recapping Current Deficiencies:

- **LLMs:** Despite their discursive fluency, LLMs lack logical rigor, are prone to inconsistency and hallucination, and their reasoning is fragile to minor perturbations. Their statistical basis fundamentally limits their capacity for verifiable deduction and deep abstract understanding. Current strategies to improve their reasoning (prompts, verifiers, superficial symbolic integration) act as patches on a foundation inadequate for advanced logic.
- **Traditional Symbolic Systems:** They offer logical rigor but are brittle in the face of the ambiguity and noise of the real world, lack the flexibility of natural language, and often face computational scalability problems.
- **Current Neuro-Symbolic Approaches:** They attempt to combine the best of both worlds but often result in shallow couplings or are limited by the expressiveness and capabilities of the standard formal languages (FOL, HOL, etc.) they use. These languages, as argued in Section 3, lack the required combination of capabilities: deep meta-linguistics, closed semantics, semantic-algorithmic capability, and mathematical foundational potential.

The Need for ULOGIC:

The central thesis is that the fundamental barrier to an AI with advanced linguistic thinking lies not solely in the neural architectures or learning algorithms, but in the **insufficiency of the underlying representation languages**. A language is needed that can:

- **Unify Language and Logic:** Serve as a basis for both fluid natural language processing and rigorous logical and mathematical derivation.
- **Be Self-Referential and Metalinguistic:** Allow the system to reason about its own knowledge and processes.
- **Have Intrinsic Semantics:** Avoid dependence on external interpretations and allow for self-consistent reasoning.
- **Be Foundationally Powerful:** Have the ability to define basic concepts (like 'set') and reconstruct complex structures (like mathematics) from within.

- **Be Algorithmically Complete:** Integrate the definition, execution, and reasoning about algorithms as an intrinsic part of the logic.

ULOGIC, as postulated in this report, is specifically designed to meet these requirements. It is the missing piece in the puzzle of advanced AI. Without a language with these capabilities, attempts to build truly comprehensive and rational artificial minds will continue to hit the inherent limitations of current formalisms.

UMIND/ULOGIC as a Superior Paradigm:

The UMIND/ULOGIC paradigm, based on the synergy between the heuristic and language generation capabilities of LLMs and the logical, algorithmic, and meta-cognitive rigor enabled by ULOGIC, offers a coherent vision to overcome these limitations. It proposes a path towards an AI that:

- **Combines Fluency and Rigor:** Handles natural language with the ease of an LLM but grounds its reasoning in the verifiable logic of ULOGIC.
- **Is Explainable and Verifiable:** Its reasoning and algorithmic execution processes are transparent and can be formally validated.
- **Has Deep Understanding Capability:** Can model complex mathematical structures, reason about itself, and potentially reconstruct foundations of knowledge.
- **Can Connect to the World:** Integrates perception and action, reasoning about sensory experience.
- **Is a Platform for Science:** Enables advanced modeling of physical systems and accelerates scientific discovery.

Therefore, UMIND/ULOGIC is not simply an alternative but is presented as a **necessary and potentially superior foundational paradigm** to current approaches. It recognizes that language and logic are not separate capabilities that must be forced to work together, but interconnected aspects of a single cognitive ability that requires a unified and sufficiently powerful underlying language like ULOGIC.

Only through the adoption of a language with these advanced features can an artificial mind be created that is capable of achieving the level of advanced linguistic thinking, complex reasoning, and self-understanding that defines general intelligence.

Referencias:

1. RATT: A Thought Structure for Coherent and Correct LLM Reasoning - arXiv, https://arxiv.org/html/2406.02746v5
2. Achieving >97% on GSM8K: Deeply Understanding the Problems Makes LLMs Better Solvers for Math Word Problems - arXiv, https://arxiv.org/html/2404.14963
3. The Ultimate Guide to LLM Reasoning (2025) - Kili Technology, https://kili-technology.com/large-language-models-llms/llm-reasoning-guide
4. Benchmarking Reasoning Robustness in Large Language Models - arXiv, https://arxiv.org/html/2503.04550v1
5. Large Language Models and Mathematical Reasoning Failures - arXiv, https://arxiv.org/html/2502.11574v1
6. Are Large Language Models Really Good Logical Reasoners? A Comprehensive Evaluation and Beyond - arXiv, http://arxiv.org/pdf/2306.09841
7. arxiv.org, https://arxiv.org/html/2502.15652
8. [2502.11574] Large Language Models and Mathematical Reasoning Failures - arXiv, https://arxiv.org/abs/2502.11574
9. ZebraLogic: On the Scaling Limits of LLMs for Logical Reasoning - arXiv, https://arxiv.org/html/2502.01100v1
10. Line Goes Up? Inherent Limitations of Benchmarks for Evaluating Large Language Models, https://arxiv.org/html/2502.14318
11. Mathematical Reasoning in Large Language Models: Assessing Logical and Arithmetic Errors across Wide Numerical Ranges - arXiv, https://arxiv.org/html/2502.08680v1
12. Benchmarking LLMs' Math Reasoning Abilities against Hard Perturbations - arXiv, https://arxiv.org/html/2502.06453v2
13. Generative AI's Act o1: The Reasoning Era Begins | Sequoia Capital, https://www.sequoiacap.com/article/generative-ais-act-o1/
14. Reasoning models don't always say what they think - Anthropic, https://www.anthropic.com/research/reasoning-models-dont-say-think
15. GitHub - LAMDASZ-ML/Awesome-Neuro-Symbolic-Learning-with-LLM, https://github.com/LAMDASZ-ML/Awesome-Neuro-Symbolic-Learning-with-LLM
16. Improving LLM Reasoning through Scaling Inference Computation with Collaborative Verification | OpenReview, https://openreview.net/forum?id=Qyile3DctL
17. Improving LLM Reasoning through Scaling Inference Computation with Collaborative Verification - arXiv, https://arxiv.org/html/2410.05318v1
18. Generating Millions Of Lean Theorems With Proofs By Exploring State Transition Graphs, https://arxiv.org/html/2503.04772v1
19. Enhancing LLM Reasoning with Stepwise Natural Language Self-Critique - arXiv, https://arxiv.org/html/2503.17363v1

20. Neuro-Symbolic AI in 2024: A Systematic Review - arXiv, https://arxiv.org/pdf/2501.05435
21. Neural-Symbolic Reasoning over Knowledge Graphs: A Survey from a Query Perspective, https://arxiv.org/html/2412.10390v1
22. A Survey on Deep Learning for Theorem Proving - arXiv, https://arxiv.org/html/2404.09939v1
23. Large Language Models as Copilots for Theorem Proving in Lean - arXiv, https://arxiv.org/abs/2404.12534
24. Formal Theorem Proving by Rewarding LLMs to Decompose Proofs Hierarchically - arXiv, https://arxiv.org/html/2411.01829v1
25. https://en.wikipedia.org/wiki/Zermelo%E2%80%93Fraenkel_set_theory
26. [1808.09701] Comparison of Two Theorem Provers: Isabelle/HOL and Coq - arXiv, https://arxiv.org/abs/1808.09701
27. European approach to artificial intelligence | Shaping Europe's digital future, https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence
28. Artificial Intelligence (AI) in Science - European Commission - Research and innovation, https://research-and-innovation.ec.europa.eu/research-area/industrial-research-and-innovation/artificial-intelligence-ai-science_en
29. European research development and deployment of AI | Shaping Europe's digital future, https://digital-strategy.ec.europa.eu/en/policies/european-ai-research
30. New Horizon Europe Funding Boosts European Research in AI and Quantum Technologies, https://digital-strategy.ec.europa.eu/en/news/new-horizon-europe-funding-boosts-european-research-ai-and-quantum-technologies
31. Commission seeks feedback on the future Strategy for Artificial Intelligence in Science, https://research-and-innovation.ec.europa.eu/news/all-research-and-innovation-news/commission-seeks-feedback-future-strategy-artificial-intelligence-science-2025-04-10_en
32. Ethics guidelines for trustworthy AI | Shaping Europe's digital future - European Union, https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai
33. EU guidelines on ethics in artificial intelligence: Context and implementation - European Parliament, https://www.europarl.europa.eu/RegData/etudes/BRIE/2019/640163/EPRS_BRI(2019)640163_EN.pdf