The NeoAGI Hyperloop Thesis: Knowledge Base as AGI Instrument

This module defines The NeoAGI Hyperloop Thesis: Knowledge Base as AGI Instrument as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page sets the identity of the whole edition and converts the previous knowledge-base engine into a neoAGI-oriented operating doctrine. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Available prior corpus scan: 10463 lines, 1087 headings, 2017537 characters. Dominant terms include what:2757, concept:2554, asks:2338, user:2151, receives:1725, failure:1696, becomes:1672, mechanism:1665, unit:1421, metric:1218, should:1212, between:1205, first:1204, evidence:1142, implementation:1087, when:1081, retrieval:1037, language:1036, second:1034, difference:1034.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. The page also imports the lesson of Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, The NeoAGI Hyperloop Thesis: Knowledge Base as AGI Instrument must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce The NeoAGI Hyperloop Thesis: Knowledge Base as AGI Instrument to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Certainty Ladder for AGI-Oriented Knowledge Engineering

This module defines Certainty Ladder for AGI-Oriented Knowledge Engineering as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page prevents epistemic inflation by separating established engineering from Moussaif invention, operational hypothesis, and speculative extension. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Available prior corpus scan: 10463 lines, 1087 headings, 2017537 characters. Dominant terms include what:2757, concept:2554, asks:2338, user:2151, receives:1725, failure:1696, becomes:1672, mechanism:1665, unit:1421, metric:1218, should:1212, between:1205, first:1204, evidence:1142, implementation:1087, when:1081, retrieval:1037, language:1036, second:1034, difference:1034.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. The page also imports the lesson of Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Certainty Ladder for AGI-Oriented Knowledge Engineering must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Certainty Ladder for AGI-Oriented Knowledge Engineering to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Line-by-Line Ingestion of the V2 Corpus and Its Limits

This module defines Line-by-Line Ingestion of the V2 Corpus and Its Limits as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page treats older books, prompts, and project memories as raw ore that must be decomposed into claims, mechanisms, constraints, and reusable objects. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Available prior corpus scan: 10463 lines, 1087 headings, 2017537 characters. Dominant terms include what:2757, concept:2554, asks:2338, user:2151, receives:1725, failure:1696, becomes:1672, mechanism:1665, unit:1421, metric:1218, should:1212, between:1205, first:1204, evidence:1142, implementation:1087, when:1081, retrieval:1037, language:1036, second:1034, difference:1034.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. The page also imports the lesson of Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Line-by-Line Ingestion of the V2 Corpus and Its Limits must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Line-by-Line Ingestion of the V2 Corpus and Its Limits to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

From Books to Operational Knowledge Objects

This module defines From Books to Operational Knowledge Objects as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page treats older books, prompts, and project memories as raw ore that must be decomposed into claims, mechanisms, constraints, and reusable objects. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Available prior corpus scan: 10463 lines, 1087 headings, 2017537 characters. Dominant terms include what:2757, concept:2554, asks:2338, user:2151, receives:1725, failure:1696, becomes:1672, mechanism:1665, unit:1421, metric:1218, should:1212, between:1205, first:1204, evidence:1142, implementation:1087, when:1081, retrieval:1037, language:1036, second:1034, difference:1034.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. The page also imports the lesson of Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, From Books to Operational Knowledge Objects must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce From Books to Operational Knowledge Objects to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

How MICAI, MCAI, MIGAGI, and neoAGI Are Integrated

This module defines How MICAI, MCAI, MIGAGI, and neoAGI Are Integrated as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page treats older books, prompts, and project memories as raw ore that must be decomposed into claims, mechanisms, constraints, and reusable objects. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Available prior corpus scan: 10463 lines, 1087 headings, 2017537 characters. Dominant terms include what:2757, concept:2554, asks:2338, user:2151, receives:1725, failure:1696, becomes:1672, mechanism:1665, unit:1421, metric:1218, should:1212, between:1205, first:1204, evidence:1142, implementation:1087, when:1081, retrieval:1037, language:1036, second:1034, difference:1034.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. The page also imports the lesson of Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, How MICAI, MCAI, MIGAGI, and neoAGI Are Integrated must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce How MICAI, MCAI, MIGAGI, and neoAGI Are Integrated to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Non-Claim Rule: Why This Book Does Not Declare AGI Solved

This module defines Non-Claim Rule: Why This Book Does Not Declare AGI Solved as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page treats older books, prompts, and project memories as raw ore that must be decomposed into claims, mechanisms, constraints, and reusable objects. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. The page also imports the lesson of Schick et al. (2023), Toolformer: language models learning when and how to use external tools. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Non-Claim Rule: Why This Book Does Not Declare AGI Solved must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Boundary rule: this edition refuses the easy error of mistaking ambition for evidence. It uses strong language for engineering orientation, but its strongest scientific claims are framed as models to be measured and destroyed. The more inventive a concept is, the more aggressive its court must become.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Non-Claim Rule: Why This Book Does Not Declare AGI Solved to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

RDF, OWL, SHACL, PROV-O as the Machine Spine

This module defines RDF, OWL, SHACL, PROV-O as the Machine Spine as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page hardens the knowledge base with formal structure, provenance, validation, and machine-readable commitments. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. The page also imports the lesson of Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, RDF, OWL, SHACL, PROV-O as the Machine Spine must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Schema rule: every domain object must have a stable identifier, type, source, claim level, conflict state, update time, owner agent, and test state. RDF supplies triples, OWL supplies ontology structure, SHACL supplies validation, and PROV-O supplies provenance. The Moussaif extension adds vessel type, MCD pressure, court result, novelty status, and action closure.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce RDF, OWL, SHACL, PROV-O as the Machine Spine to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

FAIR Knowledge Assets for LLM Thinking Models

This module defines FAIR Knowledge Assets for LLM Thinking Models as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page hardens the knowledge base with formal structure, provenance, validation, and machine-readable commitments. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. The page also imports the lesson of Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, FAIR Knowledge Assets for LLM Thinking Models must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Schema rule: every domain object must have a stable identifier, type, source, claim level, conflict state, update time, owner agent, and test state. RDF supplies triples, OWL supplies ontology structure, SHACL supplies validation, and PROV-O supplies provenance. The Moussaif extension adds vessel type, MCD pressure, court result, novelty status, and action closure.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce FAIR Knowledge Assets for LLM Thinking Models to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Claim Ontology: Fact, Consensus, Dispute, Hypothesis, Invention

This module defines Claim Ontology: Fact, Consensus, Dispute, Hypothesis, Invention as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page hardens the knowledge base with formal structure, provenance, validation, and machine-readable commitments. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. The page also imports the lesson of NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Claim Ontology: Fact, Consensus, Dispute, Hypothesis, Invention must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Schema rule: every domain object must have a stable identifier, type, source, claim level, conflict state, update time, owner agent, and test state. RDF supplies triples, OWL supplies ontology structure, SHACL supplies validation, and PROV-O supplies provenance. The Moussaif extension adds vessel type, MCD pressure, court result, novelty status, and action closure.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Claim Ontology: Fact, Consensus, Dispute, Hypothesis, Invention to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Vascular Knowledge Graphs and Multi-Vessel Provenance

This module defines Vascular Knowledge Graphs and Multi-Vessel Provenance as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page hardens the knowledge base with formal structure, provenance, validation, and machine-readable commitments. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. The page also imports the lesson of OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Vascular Knowledge Graphs and Multi-Vessel Provenance must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Schema rule: every domain object must have a stable identifier, type, source, claim level, conflict state, update time, owner agent, and test state. RDF supplies triples, OWL supplies ontology structure, SHACL supplies validation, and PROV-O supplies provenance. The Moussaif extension adds vessel type, MCD pressure, court result, novelty status, and action closure.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Vascular Knowledge Graphs and Multi-Vessel Provenance to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Definition of AGI Hyperloops for Thinking Models

This module defines Definition of AGI Hyperloops for Thinking Models as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. The page also imports the lesson of MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Definition of AGI Hyperloops for Thinking Models must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Definition of AGI Hyperloops for Thinking Models to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 1: Difference Extraction and Problem Compression

This module defines Loop 1: Difference Extraction and Problem Compression as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Schick et al. (2023), Toolformer: language models learning when and how to use external tools. The page also imports the lesson of Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 1: Difference Extraction and Problem Compression must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 1: Difference Extraction and Problem Compression to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 2: Retrieval Membrane and Context Governance

This module defines Loop 2: Retrieval Membrane and Context Governance as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. The page also imports the lesson of SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 2: Retrieval Membrane and Context Governance must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 2: Retrieval Membrane and Context Governance to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 3: Thought Search, Branching, and Backtracking

This module defines Loop 3: Thought Search, Branching, and Backtracking as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. The page also imports the lesson of GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 3: Thought Search, Branching, and Backtracking must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 3: Thought Search, Branching, and Backtracking to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 4: Critique, Falsification, and Self-Revision

This module defines Loop 4: Critique, Falsification, and Self-Revision as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. The page also imports the lesson of ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 4: Critique, Falsification, and Self-Revision must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 4: Critique, Falsification, and Self-Revision to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 5: Tool Action, World Effect, and Feedback

This module defines Loop 5: Tool Action, World Effect, and Feedback as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. The page also imports the lesson of Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 5: Tool Action, World Effect, and Feedback must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 5: Tool Action, World Effect, and Feedback to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 6: Memory Distillation and Skill Library Growth

This module defines Loop 6: Memory Distillation and Skill Library Growth as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. The page also imports the lesson of W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 6: Memory Distillation and Skill Library Growth must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 6: Memory Distillation and Skill Library Growth to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 7: Novelty Entropy and Inventive Synthesis

This module defines Loop 7: Novelty Entropy and Inventive Synthesis as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. The page also imports the lesson of W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 7: Novelty Entropy and Inventive Synthesis must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 7: Novelty Entropy and Inventive Synthesis to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 8: Benchmark Court and External Calibration

This module defines Loop 8: Benchmark Court and External Calibration as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. The page also imports the lesson of W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 8: Benchmark Court and External Calibration must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 8: Benchmark Court and External Calibration to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Loop 9: Governance, Safety, and Refusal Boundaries

This module defines Loop 9: Governance, Safety, and Refusal Boundaries as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page defines one loop in the higher control circuit that lets an LLM thinking model iterate beyond one-pass answer generation. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. The page also imports the lesson of W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Loop 9: Governance, Safety, and Refusal Boundaries must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Loop rule: a hyperloop is not a longer answer. It is a control circuit that forces the model to pass through difference extraction, retrieval, thought search, critique, tool action, memory distillation, novelty mutation, benchmark calibration, and governance. The loop is useful only when the next iteration has a better state than the previous one.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Loop 9: Governance, Safety, and Refusal Boundaries to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

MCD as Meta-Complexity Differentiation for Reasoning Control

This module defines MCD as Meta-Complexity Differentiation for Reasoning Control as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. The page also imports the lesson of Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, MCD as Meta-Complexity Differentiation for Reasoning Control must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce MCD as Meta-Complexity Differentiation for Reasoning Control to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

HAKIM-MCD as a Governance Architecture

This module defines HAKIM-MCD as a Governance Architecture as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. The page also imports the lesson of Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, HAKIM-MCD as a Governance Architecture must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce HAKIM-MCD as a Governance Architecture to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Vascular Superposition in Knowledge and Thought

This module defines Vascular Superposition in Knowledge and Thought as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. The page also imports the lesson of Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Vascular Superposition in Knowledge and Thought must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Vascular Superposition in Knowledge and Thought to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Inner Time Engine and Multi-Step Cognitive Patience

This module defines Inner Time Engine and Multi-Step Cognitive Patience as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. The page also imports the lesson of Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Inner Time Engine and Multi-Step Cognitive Patience must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Inner Time Engine and Multi-Step Cognitive Patience to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Perturbation Court and Recursive Destruction

This module defines Perturbation Court and Recursive Destruction as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. The page also imports the lesson of Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Perturbation Court and Recursive Destruction must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Perturbation Court and Recursive Destruction to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

World-Effect Loop and Action Closure

This module defines World-Effect Loop and Action Closure as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. The page also imports the lesson of Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, World-Effect Loop and Action Closure must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce World-Effect Loop and Action Closure to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Self-Reconstruction Memory and Experience Distillation

This module defines Self-Reconstruction Memory and Experience Distillation as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. The page also imports the lesson of Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Self-Reconstruction Memory and Experience Distillation must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Self-Reconstruction Memory and Experience Distillation to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

MIGAGI: Meta-Integrative Governance for AGI Loops

This module defines MIGAGI: Meta-Integrative Governance for AGI Loops as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page integrates the user theory family as generative hypotheses and architectures, not as already proven AGI science. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. The page also imports the lesson of Schick et al. (2023), Toolformer: language models learning when and how to use external tools. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, MIGAGI: Meta-Integrative Governance for AGI Loops must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Moussaif theory rule: use Vascular Superposition as a lens for multiple vessels of cognition - source, concept, mechanism, value, action, memory, critique, and goal. Use MCD to measure whether the system can distinguish these vessels rather than collapse them into one vague narrative. Use HAKIM-MCD to govern selection, refusal, and reconstruction.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce MIGAGI: Meta-Integrative Governance for AGI Loops to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

RAG Is Memory, Not Intelligence: The Missing Loops

This module defines RAG Is Memory, Not Intelligence: The Missing Loops as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. The page also imports the lesson of Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, RAG Is Memory, Not Intelligence: The Missing Loops must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce RAG Is Memory, Not Intelligence: The Missing Loops to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Self-RAG as Retrieval Self-Governance

This module defines Self-RAG as Retrieval Self-Governance as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. The page also imports the lesson of Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Self-RAG as Retrieval Self-Governance must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Self-RAG as Retrieval Self-Governance to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

GraphRAG as Global Coherence Mechanism

This module defines GraphRAG as Global Coherence Mechanism as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. The page also imports the lesson of NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, GraphRAG as Global Coherence Mechanism must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce GraphRAG as Global Coherence Mechanism to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

ReAct as the First Operational Reasoning-Action Alternation

This module defines ReAct as the First Operational Reasoning-Action Alternation as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. The page also imports the lesson of OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, ReAct as the First Operational Reasoning-Action Alternation must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

• Instrument card - Input: messy sources, prior Moussaif concepts,
• current benchmark needs, and an LLM thinking model.
• Transformation: split claims, route vessels, retrieve evidence,
• generate alternatives, run court, compress survivors.
• Output: a validated knowledge object, a loop policy, a metric
• reading, and a next experiment.
• Failure signal: impressive wording without provenance,
• mechanism, falsifier, or operational consequence.
• Upgrade path: convert the weak object into schema, test,
• dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce ReAct as the First Operational Reasoning-Action Alternation to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Tree of Thoughts and Graph of Thoughts under MCD Control

This module defines Tree of Thoughts and Graph of Thoughts under MCD Control as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. The page also imports the lesson of MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Tree of Thoughts and Graph of Thoughts under MCD Control must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Tree of Thoughts and Graph of Thoughts under MCD Control to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Toolformer, Tool Use, and API Consciousness Limits

This module defines Toolformer, Tool Use, and API Consciousness Limits as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Schick et al. (2023), Toolformer: language models learning when and how to use external tools. The page also imports the lesson of Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Toolformer, Tool Use, and API Consciousness Limits must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Toolformer, Tool Use, and API Consciousness Limits to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Reflexion and Verbal Reinforcement as Memory Without Weights

This module defines Reflexion and Verbal Reinforcement as Memory Without Weights as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. The page also imports the lesson of SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Reflexion and Verbal Reinforcement as Memory Without Weights must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Reflexion and Verbal Reinforcement as Memory Without Weights to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Voyager and the Skill Library Principle

This module defines Voyager and the Skill Library Principle as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page maps external research into the Moussaif loop vocabulary and identifies the missing mechanism that our architecture adds. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. The page also imports the lesson of GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Voyager and the Skill Library Principle must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research mapping rule: existing methods usually improve one axis - retrieval, reasoning, tool use, reflection, or skill accumulation. The hyperloop thesis is that the missing object is the disciplined composition of these axes under provenance, benchmark, and memory governance. The system must know when to retrieve, when to think, when to act, when to stop, and when to downgrade a claim.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Voyager and the Skill Library Principle to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Inventive Entropy: Raising Search Without Losing Truth

This module defines Inventive Entropy: Raising Search Without Losing Truth as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. The page also imports the lesson of ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Inventive Entropy: Raising Search Without Losing Truth must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Inventive Entropy: Raising Search Without Losing Truth to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Negative-Space Mining: What the Field Does Not Ask

This module defines Negative-Space Mining: What the Field Does Not Ask as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. The page also imports the lesson of Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Negative-Space Mining: What the Field Does Not Ask must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Negative-Space Mining: What the Field Does Not Ask to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Cross-Domain Transplantation Without Category Confusion

This module defines Cross-Domain Transplantation Without Category Confusion as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. The page also imports the lesson of W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Cross-Domain Transplantation Without Category Confusion must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Cross-Domain Transplantation Without Category Confusion to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Edge-Case Manipulation and Boundary Intelligence

This module defines Edge-Case Manipulation and Boundary Intelligence as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. The page also imports the lesson of W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Edge-Case Manipulation and Boundary Intelligence must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Edge-Case Manipulation and Boundary Intelligence to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

The Moussaif Novelty-Survivability Index

This module defines The Moussaif Novelty-Survivability Index as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. The page also imports the lesson of W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, The Moussaif Novelty-Survivability Index must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce The Moussaif Novelty-Survivability Index to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Anti-Slop Laws for Generative Knowledge Bases

This module defines Anti-Slop Laws for Generative Knowledge Bases as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. The page also imports the lesson of W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Anti-Slop Laws for Generative Knowledge Bases must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Anti-Slop Laws for Generative Knowledge Bases to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Inventive Knowledge Engineering Patterns

This module defines Inventive Knowledge Engineering Patterns as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. The page also imports the lesson of Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Inventive Knowledge Engineering Patterns must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Inventive Knowledge Engineering Patterns to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

The Hyperskill Layer: Skills About Generating Skills

This module defines The Hyperskill Layer: Skills About Generating Skills as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns creativity from aesthetic improvisation into controlled search, mutation, survival testing, and implementation pressure. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. The page also imports the lesson of Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, The Hyperskill Layer: Skills About Generating Skills must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Inventive entropy rule: novelty is useful only when it survives mechanism, evidence, and adversarial compression. The system deliberately mutates frames through inversion, negative-space mining, cross-domain transplant, edge-case manipulation, and first-principles reduction. But every mutation is forced to state its use condition, failure mode, and test path.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce The Hyperskill Layer: Skills About Generating Skills to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

AGI-Usefulness Is Not Leaderboard Score

This module defines AGI-Usefulness Is Not Leaderboard Score as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. The page also imports the lesson of Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, AGI-Usefulness Is Not Leaderboard Score must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce AGI-Usefulness Is Not Leaderboard Score to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

HELM, SWE-bench, GAIA, and ARC as Four Test Realms

This module defines HELM, SWE-bench, GAIA, and ARC as Four Test Realms as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. The page also imports the lesson of Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, HELM, SWE-bench, GAIA, and ARC as Four Test Realms must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce HELM, SWE-bench, GAIA, and ARC as Four Test Realms to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

MCD-Bench: Measuring Differentiated Reasoning

This module defines MCD-Bench: Measuring Differentiated Reasoning as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. The page also imports the lesson of Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, MCD-Bench: Measuring Differentiated Reasoning must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce MCD-Bench: Measuring Differentiated Reasoning to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Vascular Governance Benchmark for Agentic Systems

This module defines Vascular Governance Benchmark for Agentic Systems as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. The page also imports the lesson of Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Vascular Governance Benchmark for Agentic Systems must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Vascular Governance Benchmark for Agentic Systems to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Novelty Engine Benchmark for Non-Trivial Discovery

This module defines Novelty Engine Benchmark for Non-Trivial Discovery as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. The page also imports the lesson of Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Novelty Engine Benchmark for Non-Trivial Discovery must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Novelty Engine Benchmark for Non-Trivial Discovery to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Knowledge-Base Regression Tests and Decay Audits

This module defines Knowledge-Base Regression Tests and Decay Audits as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. The page also imports the lesson of Schick et al. (2023), Toolformer: language models learning when and how to use external tools. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Knowledge-Base Regression Tests and Decay Audits must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Knowledge-Base Regression Tests and Decay Audits to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Cost, Latency, Context, and Reasoning-Budget Governance

This module defines Cost, Latency, Context, and Reasoning-Budget Governance as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page converts claims into evaluation designs and resists benchmark worship by measuring transfer, cost, calibration, and failure modes. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. The page also imports the lesson of Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Cost, Latency, Context, and Reasoning-Budget Governance must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Evaluation rule: a leaderboard result is not a theory of intelligence. The system must compare task success, transfer, cost per correct solution, context economy, calibration, robustness, contamination risk, and reproducibility. A model that wins by overfitting the benchmark fails the AGI-usability standard even if the score looks impressive.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Cost, Latency, Context, and Reasoning-Budget Governance to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

NeoAGI Knowledge Base OS: Minimal Architecture

This module defines NeoAGI Knowledge Base OS: Minimal Architecture as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. The page also imports the lesson of Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, NeoAGI Knowledge Base OS: Minimal Architecture must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce NeoAGI Knowledge Base OS: Minimal Architecture to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Schema Pack: Claim, Source, Loop, Tool, Benchmark, Risk

This module defines Schema Pack: Claim, Source, Loop, Tool, Benchmark, Risk as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. The page also imports the lesson of NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Schema Pack: Claim, Source, Loop, Tool, Benchmark, Risk must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Schema Pack: Claim, Source, Loop, Tool, Benchmark, Risk to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Prompt Compiler for Hyperloop Execution

This module defines Prompt Compiler for Hyperloop Execution as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. The page also imports the lesson of OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Prompt Compiler for Hyperloop Execution must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Prompt Compiler for Hyperloop Execution to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Agent Roles: Librarian, Critic, Synthesizer, Verifier, Operator

This module defines Agent Roles: Librarian, Critic, Synthesizer, Verifier, Operator as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. The page also imports the lesson of MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Agent Roles: Librarian, Critic, Synthesizer, Verifier, Operator must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Agent Roles: Librarian, Critic, Synthesizer, Verifier, Operator to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Context Routing and Retrieval Membrane Policy

This module defines Context Routing and Retrieval Membrane Policy as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Schick et al. (2023), Toolformer: language models learning when and how to use external tools. The page also imports the lesson of Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Context Routing and Retrieval Membrane Policy must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Context Routing and Retrieval Membrane Policy to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Knowledge Graph Build Pipeline

This module defines Knowledge Graph Build Pipeline as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Shinn et al. (2023), Reflexion: verbal reinforcement and episodic memory for language agents. The page also imports the lesson of SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Knowledge Graph Build Pipeline must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Knowledge Graph Build Pipeline to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Skill Library and Tool Registry

This module defines Skill Library and Tool Registry as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wang et al. (2023), Voyager: lifelong LLM agent with curriculum, skill library, and self-verification. The page also imports the lesson of GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Skill Library and Tool Registry must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Skill Library and Tool Registry to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Dashboard Metrics for Scientific and Product Use

This module defines Dashboard Metrics for Scientific and Product Use as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page specifies an executable layer: schemas, agents, routing rules, dashboards, validators, and update loops. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): govern, map, measure, and manage AI risks. The page also imports the lesson of ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Dashboard Metrics for Scientific and Product Use must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Implementation rule: the book must compile into an operating system. The OS has a source ingestor, ontology builder, claim splitter, retrieval membrane, loop compiler, court engine, benchmark runner, memory distiller, tool registry, and dashboard. Every chapter should become executable as a schema, prompt, workflow, metric, or dataset.

• Instrument card - Input: messy sources, prior Moussaif concepts,
• current benchmark needs, and an LLM thinking model.
• Transformation: split claims, route vessels, retrieve evidence,
• generate alternatives, run court, compress survivors.
• Output: a validated knowledge object, a loop policy, a metric
• reading, and a next experiment.
• Failure signal: impressive wording without provenance,
• mechanism, falsifier, or operational consequence.
• Upgrade path: convert the weak object into schema, test,
• dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Dashboard Metrics for Scientific and Product Use to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Prompt Injection and Knowledge-Base Poisoning

This module defines Prompt Injection and Knowledge-Base Poisoning as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page attacks the system from the adversarial side: injection, poisoning, laundering, drift, overclaim, and governance collapse. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: OWASP Top 10 for LLM Applications: prompt injection, data leakage, insecure tool use, and related application risks. The page also imports the lesson of Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Prompt Injection and Knowledge-Base Poisoning must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Security rule: a knowledge base that can guide agents can also be poisoned, over-trusted, or exploited. The court must test prompt injection, malicious retrieval, source laundering, fake consensus, tool escalation, benchmark gaming, and ontology drift. Safety is not external to intelligence; it is one vessel of intelligence.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Prompt Injection and Knowledge-Base Poisoning to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Provenance Defense and Source Laundering Detection

This module defines Provenance Defense and Source Laundering Detection as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page attacks the system from the adversarial side: injection, poisoning, laundering, drift, overclaim, and governance collapse. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: MITRE ATLAS: adversarial tactics and techniques against AI-enabled systems. The page also imports the lesson of W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Provenance Defense and Source Laundering Detection must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Security rule: a knowledge base that can guide agents can also be poisoned, over-trusted, or exploited. The court must test prompt injection, malicious retrieval, source laundering, fake consensus, tool escalation, benchmark gaming, and ontology drift. Safety is not external to intelligence; it is one vessel of intelligence.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Provenance Defense and Source Laundering Detection to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Ontology Drift, Benchmark Overfitting, and False AGI

This module defines Ontology Drift, Benchmark Overfitting, and False AGI as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page attacks the system from the adversarial side: injection, poisoning, laundering, drift, overclaim, and governance collapse. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Stanford HELM (2022 onward): holistic, transparent, scenario-and-metric evaluation of language models. The page also imports the lesson of W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Ontology Drift, Benchmark Overfitting, and False AGI must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Security rule: a knowledge base that can guide agents can also be poisoned, over-trusted, or exploited. The court must test prompt injection, malicious retrieval, source laundering, fake consensus, tool escalation, benchmark gaming, and ontology drift. Safety is not external to intelligence; it is one vessel of intelligence.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Ontology Drift, Benchmark Overfitting, and False AGI to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Recursive Court of Destruction for the Entire System

This module defines Recursive Court of Destruction for the Entire System as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page attacks the system from the adversarial side: injection, poisoning, laundering, drift, overclaim, and governance collapse. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: SWE-bench (2023 onward): repository-level software issue resolution benchmark for language agents. The page also imports the lesson of W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Recursive Court of Destruction for the Entire System must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Security rule: a knowledge base that can guide agents can also be poisoned, over-trusted, or exploited. The court must test prompt injection, malicious retrieval, source laundering, fake consensus, tool escalation, benchmark gaming, and ontology drift. Safety is not external to intelligence; it is one vessel of intelligence.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Recursive Court of Destruction for the Entire System to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Twelve-Year Research Map for neoAGI Knowledge Engineering

This module defines Twelve-Year Research Map for neoAGI Knowledge Engineering as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns the book into a long research agenda with testable hypotheses, experiments, datasets, and falsification conditions. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: GAIA (2023): real-world assistant benchmark requiring reasoning, web use, multimodality, and tool use. The page also imports the lesson of W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Twelve-Year Research Map for neoAGI Knowledge Engineering must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research rule: the program advances only when each bold claim becomes a hypothesis, each hypothesis becomes a protocol, each protocol becomes data, and each dataset becomes a revised theory. The book therefore ends not with applause but with work: experiments, ablations, failure taxonomies, reproducibility packages, and domain pilots.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Twelve-Year Research Map for neoAGI Knowledge Engineering to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Twenty Research Hypotheses for LLM-Only Thinking Systems

This module defines Twenty Research Hypotheses for LLM-Only Thinking Systems as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns the book into a long research agenda with testable hypotheses, experiments, datasets, and falsification conditions. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: ARC Prize / ARC-AGI: few-shot abstraction and agentic generalization benchmarks; 2025 reports emphasize remaining gaps between systems and human-like efficient learning. The page also imports the lesson of Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Twenty Research Hypotheses for LLM-Only Thinking Systems must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research rule: the program advances only when each bold claim becomes a hypothesis, each hypothesis becomes a protocol, each protocol becomes data, and each dataset becomes a revised theory. The book therefore ends not with applause but with work: experiments, ablations, failure taxonomies, reproducibility packages, and domain pilots.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Twenty Research Hypotheses for LLM-Only Thinking Systems to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Ten Experiments That Can Disprove the Program

This module defines Ten Experiments That Can Disprove the Program as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns the book into a long research agenda with testable hypotheses, experiments, datasets, and falsification conditions. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Recent surveys on test-time scaling and LLM agents (2025-2026): inference compute, self-correction, search, adaptive budgets, and agent evaluation remain active research frontiers. The page also imports the lesson of Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Ten Experiments That Can Disprove the Program must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research rule: the program advances only when each bold claim becomes a hypothesis, each hypothesis becomes a protocol, each protocol becomes data, and each dataset becomes a revised theory. The book therefore ends not with applause but with work: experiments, ablations, failure taxonomies, reproducibility packages, and domain pilots.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Ten Experiments That Can Disprove the Program to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

This module defines From Book to Dataset, Library, Dashboard, and Agent OS as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns the book into a long research agenda with testable hypotheses, experiments, datasets, and falsification conditions. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is MCD-K, measuring differentiation, hierarchy, mechanism clarity, action closure, temporality, counter-model awareness, falsifiability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C RDF 1.1 Concepts (2014): RDF graphs as subject-predicate-object triples for machine-readable descriptions. The page also imports the lesson of Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and critique through reflection tokens. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, From Book to Dataset, Library, Dashboard, and Agent OS must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research rule: the program advances only when each bold claim becomes a hypothesis, each hypothesis becomes a protocol, each protocol becomes data, and each dataset becomes a revised theory. The book therefore ends not with applause but with work: experiments, ablations, failure taxonomies, reproducibility packages, and domain pilots.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce From Book to Dataset, Library, Dashboard, and Agent OS to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

How to Conquer Fields Without Slop: The Field-Invasion Protocol

This module defines How to Conquer Fields Without Slop: The Field-Invasion Protocol as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page turns the book into a long research agenda with testable hypotheses, experiments, datasets, and falsification conditions. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is AGI-Usability Coefficient, measuring machine readability, task routing, uncertainty tags, provenance, tool affordance, benchmark traceability. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C OWL 2 Overview (2012): formal ontologies with classes, properties, individuals, and data values. The page also imports the lesson of Edge et al. (2024), GraphRAG: local-to-global graph retrieval and community summaries for query-focused summarization. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, How to Conquer Fields Without Slop: The Field-Invasion Protocol must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Research rule: the program advances only when each bold claim becomes a hypothesis, each hypothesis becomes a protocol, each protocol becomes data, and each dataset becomes a revised theory. The book therefore ends not with applause but with work: experiments, ablations, failure taxonomies, reproducibility packages, and domain pilots.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce How to Conquer Fields Without Slop: The Field-Invasion Protocol to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

The Final NeoAGI Hyperloop Model in One Page

This module defines The Final NeoAGI Hyperloop Model in One Page as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page compresses the edition into final models, quality judgment, source register, and next-build obligations. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Novelty-Survivability Index, measuring novelty, mechanistic plausibility, transfer value, adversarial survival, implementation readiness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation through shapes and constraints. The page also imports the lesson of Yao et al. (2022/2023), ReAct: interleaving reasoning traces and task-specific actions. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, The Final NeoAGI Hyperloop Model in One Page must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Closure rule: the goal is not to proclaim final AGI. The goal is to close the engineering vocabulary around what LLM thinking models can do when surrounded by structured knowledge, recursive critique, retrieval governance, novelty pressure, tool action, and memory. The remaining ignorance is part of the model, not an embarrassment to hide.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce The Final NeoAGI Hyperloop Model in One Page to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Quality Evaluation and Remaining Ignorance

This module defines Quality Evaluation and Remaining Ignorance as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page compresses the edition into final models, quality judgment, source register, and next-build obligations. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Retrieval Membrane Score, measuring source fit, query reformulation gain, citation density, context economy, conflict awareness. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: W3C PROV-O (2013): provenance classes and properties for entities, activities, and agents. The page also imports the lesson of Yao et al. (2023), Tree of Thoughts: deliberate search over coherent thought units. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Quality Evaluation and Remaining Ignorance must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Closure rule: the goal is not to proclaim final AGI. The goal is to close the engineering vocabulary around what LLM thinking models can do when surrounded by structured knowledge, recursive critique, retrieval governance, novelty pressure, tool action, and memory. The remaining ignorance is part of the model, not an embarrassment to hide.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Quality Evaluation and Remaining Ignorance to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

NEOAGI-KB

Reference Spine and Citation Register

This module defines Reference Spine and Citation Register as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page compresses the edition into final models, quality judgment, source register, and next-build obligations. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Court Survival Ratio, measuring percentage of claims that survive attack, revision, retest, and implementation pressure. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Wilkinson et al. (2016), FAIR Guiding Principles: findability, accessibility, interoperability, and reusability for data, algorithms, tools, and workflows. The page also imports the lesson of Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph structures over LLM thoughts and transformations. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Reference Spine and Citation Register must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Closure rule: the goal is not to proclaim final AGI. The goal is to close the engineering vocabulary around what LLM thinking models can do when surrounded by structured knowledge, recursive critique, retrieval governance, novelty pressure, tool action, and memory. The remaining ignorance is part of the model, not an embarrassment to hide.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Reference Spine and Citation Register to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

Reference spine used in this extension:
- W3C RDF 1.1 Concepts (2014): RDF graphs as
- subject-predicate-object triples for machine-readable descriptions.
- W3C OWL 2 Overview (2012): formal ontologies with classes,
- properties, individuals, and data values.
- W3C SHACL (2017; SHACL 1.2 Core draft 2026): graph validation
- through shapes and constraints.
- W3C PROV-O (2013): provenance classes and properties for
- entities, activities, and agents.
- Wilkinson et al. (2016), FAIR Guiding Principles: findability,
- accessibility, interoperability, and reusability for data, algorithms, tools, and workflows.
- Lewis et al. (2020), Retrieval-Augmented Generation: combining
- parametric memory and non-parametric memory for knowledge-intensive generation.
- Asai et al. (2023), Self-RAG: adaptive retrieval, generation, and
- critique through reflection tokens.
- Edge et al. (2024), GraphRAG: local-to-global graph retrieval and
- community summaries for query-focused summarization.
- Yao et al. (2022/2023), ReAct: interleaving reasoning traces and
- task-specific actions.
- Yao et al. (2023), Tree of Thoughts: deliberate search over
- coherent thought units.
- Besta et al. (2023/2024), Graph of Thoughts: arbitrary graph
- structures over LLM thoughts and transformations.
- Schick et al. (2023), Toolformer: language models learning w
- and how to use external tools.
- Shinn et al. (2023), Reflexion: verbal reinforcement and episodic
- memory for language agents.
- Wang et al. (2023), Voyager: lifelong LLM agent with curricu
- skill library, and self-verification.
- NIST AI RMF 1.0 (2023) and Generative AI Profile (2024): gov
- map, measure, and manage AI risks.
- OWASP Top 10 for LLM Applications: prompt injection, data
- leakage, insecure tool use, and related application risks.
- MITRE ATLAS: adversarial tactics and techniques against
- AI-enabled systems.
- Stanford HELM (2022 onward): holistic, transparent,
- scenario-and-metric evaluation of language models.
- SWE-bench (2023 onward): repository-level software issue
- resolution benchmark for language agents.
- GAIA (2023): real-world assistant benchmark requiring reasoning,
- web use, multimodality, and tool use.

NEOAGI-KB

Construction Log for Version III and Next Edition

This module defines Construction Log for Version III and Next Edition as part of the NeoAGI Hyperloop Knowledge-Base Engineering layer. Its function is not to add decorative terminology, but to convert a prior book-scale knowledge base into an operational instrument for LLM thinking models. The page compresses the edition into final models, quality judgment, source register, and next-build obligations. The central rule is: a knowledge base becomes AGI-relevant only when it changes what the model can retrieve, distinguish, critique, test, remember, and do across repeated loops.

The inherited MICAI/MCAI/MIGAGI frame is used as a disciplined hypothesis family. MCAI names the meta-complexity agency view; MICAI names the integrative cognitive-artificial intelligence laboratory; MIGAGI names the governance-oriented extension toward general intelligence loops. None of these labels is treated as established science. They are formalized as architectures to be stress-tested against public research, standards, benchmarks, and adversarial failure cases. Prior Moussaif corpus basis: V2 supplies the generative knowledge-base engine, the Closure Codex, MCD-K, the Recursive Court of Destruction, the Epistemic Immune System, and the AGI-usability vocabulary. This page injects a neoAGI hyperloop specialization into that base.

Operationally, this module turns the page topic into ten linked objects: claim, mechanism, evidence tier, retrieval boundary, ontology object, critique gate, action protocol, benchmark, memory update, security constraint. A paragraph is not accepted as knowledge until it can be converted into at least one of those objects. A claim that cannot receive provenance remains weak. A concept that cannot be differentiated remains rhetorical. A mechanism that cannot face a counter-model remains immature. A protocol that cannot be benchmarked remains aspiration.

The relevant metric here is Knowledge Metabolism Rate, measuring speed and quality with which raw sources become validated, reusable, linked knowledge objects. The score is never used as absolute truth. It is used as a pressure device: low score demands redesign, high score demands attack, stable score demands transfer to a new domain. The metric therefore behaves like a governance instrument, not a vanity number.

Research spine: Lewis et al. (2020), Retrieval-Augmented Generation: combining parametric memory and non-parametric memory for knowledge-intensive generation. The page also imports the lesson of Schick et al. (2023), Toolformer: language models learning when and how to use external tools. Together these references constrain the Moussaif invention: formal knowledge must be represented, retrieved, reflected upon, acted through, benchmarked, and protected. The novelty engine cannot be allowed to bypass evidence; it must route novelty through structure.

Protocol: 1) extract the smallest difference that makes the task non-trivial; 2) route the difference to a vessel: source, concept, mechanism, action, risk, or benchmark; 2) retrieve only what changes the reasoning state rather than what merely increases volume; 2) represent the result as a claim object with provenance, uncertainty, and use conditions; 2) generate competing interpretations before selecting a path; 2) attack the selected path through counterexamples, adversarial prompts, and benchmark pressure; 2) compress the surviving result into memory, schema, tool instruction, and next experiment. This protocol is the minimal hyperloop. If any step is skipped, the system falls back into ordinary prompt-response behavior. If all steps run recursively, the model begins to approximate a research process: not omniscience, not guaranteed AGI, but a stronger form of controlled cognitive work.

Implementation consequence: in a real neoAGI knowledge base, Construction Log for Version III and Next Edition must compile into a reusable artifact. The artifact may be a SHACL shape, a prompt segment, a benchmark rubric, an agent role card, a tool policy, a memory update rule, or a dashboard metric. A page that remains only prose is considered half-built. The book therefore treats writing as the first representation layer, not the final product.

Domain-conquest use: to enter a new field without slop, the system first builds a domain claim map, then separates vocabulary from mechanisms, mechanisms from evidence, evidence from benchmarkable behavior, and benchmarkable behavior from commercial or scientific deployment. The model must ask what experts measure, what outsiders misunderstand, what sources are canonical, what edges cases break the field, and what hidden variable reorganizes the map.

Ablation design: remove retrieval and observe hallucination; remove graph structure and observe coherence loss; remove the court and observe overclaim; remove novelty entropy and observe stagnation; remove provenance and observe source laundering; remove benchmark pressure and observe self-praise; remove memory distillation and observe repeated rediscovery. These ablations define why the hyperloop is a system rather than a list of tricks.

Knowledge-object schema: id, title, vessel, source set, claim level, mechanism statement, opposing model, uncertainty, test, risk, owner agent, last update, dependent nodes, failure examples, and compression seal. The compression seal is the one sentence that survives after evidence, critique, and implementation pressure. Without such a seal, long context becomes storage rather than intelligence.

Closure rule: the goal is not to proclaim final AGI. The goal is to close the engineering vocabulary around what LLM thinking models can do when surrounded by structured knowledge, recursive critique, retrieval governance, novelty pressure, tool action, and memory. The remaining ignorance is part of the model, not an embarrassment to hide.

- Instrument card - Input: messy sources, prior Moussaif concepts,
- current benchmark needs, and an LLM thinking model.
- Transformation: split claims, route vessels, retrieve evidence,
- generate alternatives, run court, compress survivors.
- Output: a validated knowledge object, a loop policy, a metric
- reading, and a next experiment.
- Failure signal: impressive wording without provenance,
- mechanism, falsifier, or operational consequence.
- Upgrade path: convert the weak object into schema, test,
- dashboard widget, or rejection note.

Destruction test for this page: Can a hostile reviewer reduce Construction Log for Version III and Next Edition to a renamed existing method? If yes, the page must state the exact prior method, the added Moussaif constraint, and the measurable difference. Can the page be implemented by an agent without extra explanation? If no, it must be rewritten as a schema or protocol. Can the page be disproved? If no, it is downgraded from model to metaphor.

Construction scan: the available V2 Markdown contained 10463 lines, 1087 headings, and 2017537 characters. The first headings were: THE MOUSSAIF GENERATIVE KNOWLEDGE BASE ENGINE; Methodological Warning: Evidence, Synthesis, and Invention; Detailed Table of Contents; Part I - The Knowledge Base Becomes a Cognitive Reactor; Part II - Evidence Discipline, Epistemic Tiers, and Anti-Summary Reading; Part III - Ontologies, Knowledge Graphs, Provenance, and Machine-Usable Structure; Part IV - RAG, GraphRAG, Self-RAG, Long Context, and Agentic Retrieval; Part V - Creative Entropy, Constructive Chaos, and Invention Under Law; Part VI - Moussaif Synthetic Theories for Generative Knowledge Bases; Part VII - Operational Protocols for Research-Warrior Knowledge Production; Part VIII - Applications: Products, Websites, Education, Research, and Agents; Part IX - The Court of Destruction and Regenerative Revision. This does not mean every unmounted page or book was accessible; it means the active runtime corpus was read and converted into a Version III extension.

NEOAGI-KB