Towards Neurocognitive-Inspired Intelligence

From AI's Structural Mimicry to Human-Like Functional Cognition

Golilarz, Al Khatib & Rahimi — arXiv, Oct 2025

Simplified for Everyone

This document explains in plain English how copying the brain's thinking process — not just its structure — could create AI that truly learns, reasons, and adapts. Same 25 pages as the original, no jargon left unexplained.

Original: arXiv:2510.13826

Simplified by: @muhamedfazalps7

1. What This Paper Is About

Artificial intelligence has made incredible progress. AI can recognize faces, understand speech, translate languages, and even write essays. But there's a problem: most AI systems are narrow specialists. They're great at one specific task but struggle when things change or when they encounter something new.

Humans, on the other hand, are remarkably flexible. We can learn a new skill in minutes, adapt to unfamiliar situations, understand context, and reason about cause and effect. We don't need millions of examples to learn — often one or two are enough.

This paper introduces a new framework called **Neurocognitive-Inspired Intelligence (NII)**. Instead of just copying the brain's structure (layers of artificial neurons), NII copies the brain's *thinking process* — how it perceives, remembers, pays attention, reasons, and adapts.

Think of it this way: traditional AI builds a brain-shaped computer. NII builds a computer that *thinks like* a brain.

2. Where Current AI Falls Short

Before explaining the solution, the paper identifies seven key problems with today's AI:

Problem 1: Poor Generalization

AI trained on one dataset often fails on slightly different data. A model that perfectly recognizes cats in bright daylight might fail completely when the lighting changes or the cat is partially hidden. Humans handle these variations effortlessly.

Problem 2: Data Inefficiency

Deep learning needs enormous amounts of labeled data. Training a state-of-the-art image classifier requires millions of examples. A child can learn what a dog is after seeing just one or two dogs. Current AI is wasteful with data.

Problem 3: No Long-Term Memory

AI models process information in isolation — each input is independent. They don't remember past interactions or build on previous knowledge. Humans accumulate wisdom over a lifetime.

Problem 4: Inflexibility

Once trained, most AI models are frozen. Adapting to a new task requires retraining from scratch. Humans can switch between tasks instantly, drawing on past experience.

Problem 5: Fragmented Processing

Current AI systems handle perception, language, and action as separate problems. The brain integrates all of these into a unified experience. A robot that sees a cup, understands the word 'cup,' and reaches to grab it is doing something that current AI struggles to coordinate.

Problem 6: Adversarial Fragility

Tiny, invisible changes to an image can fool AI completely. Add a few specially chosen pixels to a photo of a panda, and the AI might call it a gibbon with 99% confidence. Humans are immune to such tricks.

Problem 7: No Common Sense

AI lacks intuitive understanding of how the world works. It doesn't know that objects fall when dropped, that people get sad when bad things happen, or that a cup can't be in two places at once.

The Core Issue: Current AI mimics the brain's structure (neurons and connections) but ignores how the brain actually *thinks*. This paper argues that functional mimicry — copying the thinking process — is the key to better AI.

3. Why Biology Is the Blueprint

The human brain is the most impressive information-processing system ever created. It uses about 20 watts of power (same as a dim lightbulb) yet performs feats that supercomputers can't match. Understanding how it works gives us a roadmap for building better AI.

Embodied Cognition: Intelligence Needs a Body

One of the most important insights from cognitive science is that intelligence isn't just about abstract computation — it's deeply rooted in physical interaction with the world. When you pick up a cup, your brain combines visual information, touch feedback, weight estimation, and motor planning in a seamless loop. This is called **embodied cognition**.

Current AI is disembodied — it processes text and images but has no physical experience. NII argues that true intelligence requires connecting perception to action through feedback loops.

The Perception-Action Loop

In the brain, perception isn't passive — it's active and goal-directed. When you look for your keys, your eyes don't just passively scan the room. Your attention is guided by your goal (find keys), your memory (where did I last see them?), and your expectations (keys are usually on surfaces). This closed loop between sensing, thinking, and acting is fundamental to biological intelligence.

Predictive Processing: The Brain as a Prediction Machine

The brain constantly predicts what will happen next. When you catch a ball, your brain isn't reacting to each frame of the ball's trajectory — it's predicting where the ball will be and adjusting your hand accordingly. This is called **predictive coding**.

When predictions are wrong (prediction error), the brain updates its internal model. This error-driven learning is how the brain stays accurate without needing millions of training examples.

Memory: Not Just Storage, But a Thinking Tool

The brain has multiple memory systems that serve different purposes:

- **Working memory:** A mental scratchpad that holds information temporarily while you think about it
- **Episodic memory:** Records of specific events — what happened, when, and where
- **Semantic memory:** General knowledge and facts about the world
- **Procedural memory:** Skills and habits — how to ride a bike, type on a keyboard

Current AI has no equivalent of these systems. It processes information in one pass and forgets it. NII proposes explicit memory modules that store, organize, and retrieve information like the brain does.

Attention: The Brain's Filter

Your senses are flooded with information every second — millions of visual details, sounds, smells, and sensations. The brain uses **attention** to filter out the noise and focus on what matters. This isn't just a technical convenience — it's essential for intelligence.

Attention in the brain works two ways: bottom-up (something catches your eye) and top-down (you decide to look for something). Current AI attention (like in transformers) is mostly bottom-up. NII proposes attention mechanisms that combine both approaches.

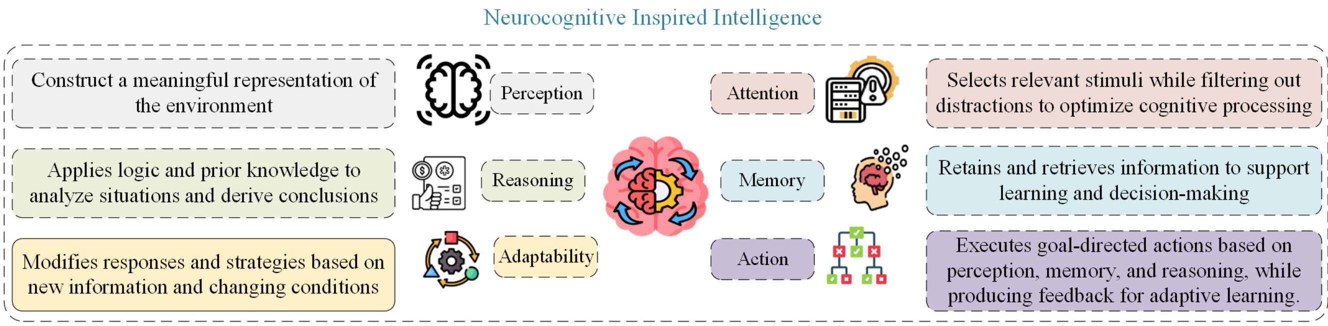
Neurocognitive Inspired Intelligence
Construct a meaningful representation of the environment
Perception
Applies logic and prior knowledge to analyze situations and derive conclusions
Reasoning
Modifies responses and strategies based on new information and changing conditions
Adaptability
Attention
Selects relevant stimuli while filtering out distractions to optimize cognitive processing
Memory
Retains and retrieves information to support learning and decision-making
Action
Executes goal-directed actions based on perception, memory, and reasoning, while producing feedback for adaptive learning.

4. The Neurocognitive-Inspired Intelligence Framework

The NII framework is a blueprint for building AI that thinks like a brain. It consists of seven interconnected modules, each inspired by a different brain function. These modules work together in a continuous loop, not a one-way pipeline.

Module 1: Perception Unit

This module converts raw sensory input (images, sounds, touch) into organized internal representations. It works like the visual and auditory cortices in the brain.

- **How it works:** Uses hierarchical feature extraction — first detecting edges and textures, then objects, then scenes and relationships
- **Key innovation:** Active perception — the system can direct its sensors to gather more information when the current view is unclear, just like moving your head to see better
- **Biology:** Inspired by the progression from simple cells in V1 to complex object recognition in the IT cortex

Module 2: Attention Mechanism

This module acts as a dynamic gatekeeper, deciding what information gets processed and what gets ignored. It's like the brain's spotlight of attention.

- **Bottom-up:** Sudden movement or bright colors automatically grab attention
- **Top-down:** Your current goal directs attention (looking for a specific person in a crowd)
- **Key innovation:** Attention that works across space, time, and meaning — not just where to look, but when to remember and what concept matters

Module 3: Memory Module

An explicit memory system with short-term (working) and long-term (episodic and semantic) components. Unlike current AI, which stores information implicitly in neural network weights, this module has dedicated storage that can be accessed and updated.

- **Working memory:** Temporary buffer for active thinking
- **Episodic memory:** Records of specific experiences with timestamps and context
- **Semantic memory:** Abstract knowledge and facts
- **Key innovation:** Memory consolidation — transferring important experiences to long-term storage, similar to how sleep helps consolidate human memories

Module 4: Learning Module

This module controls how stored representations change over time. It's the brain's plasticity — the ability to strengthen or weaken connections based on experience.

- **What it does:** Extracts patterns from experience, updates knowledge based on feedback, and adjusts reasoning strategies
- **Key innovation:** Prevents catastrophic forgetting — the tendency of AI to forget old knowledge when learning new things — using techniques inspired by synaptic consolidation

Module 5: Reasoning Engine

The cognitive core that combines pattern recognition (sub-symbolic) with logical reasoning (symbolic). This dual-process approach mirrors how humans think — fast intuitive responses combined with slow deliberate reasoning.

Sub-symbolic: Quick pattern matching using neural networks — like recognizing a face

Symbolic: Logical reasoning using rules and relationships — like solving a math problem

Key innovation: The two systems work together — patterns suggest hypotheses, logic tests them

Module 6: Adaptation Layer

The meta-controller that monitors performance and adjusts the entire system. It decides when to learn, when to explore, and when to play it safe.

What it does: Tracks prediction errors, manages learning rates, and reconfigures cognitive pathways

Key innovation: Continual learning — the system improves over its entire lifetime without forgetting what it already knows

Module 7: Action and Decision Unit

Converts cognitive output into behavior — whether that's moving a robot arm, generating speech, or making a decision. Critically, actions produce feedback that feeds back into the perception module, closing the cognitive loop.

Key innovation: Confidence-based gating — when uncertain, the system can delay action, request more information, or choose a safer option

> **The Key Difference:** Traditional AI is a pipeline — data goes in one end, decisions come out the other. NII is a loop — every action produces feedback that improves future perception and reasoning. This is how the brain works, and it's what makes biological intelligence so robust.

5. How NII Compares to Other Approaches

The paper includes a detailed comparison of three AI approaches. Here's a simplified version:

Feature	Traditional AI	Brain-Inspired AI	NII (This Paper)
Mimics	Statistical patterns	Brain structure	Brain thinking process
Transparency	Black box	Black box	Transparent by design
Adaptability	Fixed after training	Limited	Continuous learning
Memory	No explicit memory	Implicit in weights	Explicit working/long-term
Reasoning	Pattern matching only	Pattern matching only	Patterns + logic
Embodiment	Disembodied	Disembodied	Embodied with feedback loops
Common sense	None	None	World models + causal reasoning
Goal	Optimize for specific tasks	Replicate brain architecture	Think like a brain

Why This Matters

Traditional AI is like a calculator — amazing at math but useless for anything else. Brain-inspired AI builds a calculator that looks like a brain but still only does math. NII builds a system that *thinks* like a brain — flexible, adaptive, and capable of learning new things.

6. Real-World Applications

The paper demonstrates how NII could be applied in several practical domains:

Application 1: Resilient Robotics

Current robots are fragile — they break down when encountering unexpected situations. A neurocognitive robot would combine vision and touch to understand its environment, use memory to learn from past experiences, and adapt its behavior in real time.

> **Experiment:** The researchers tested vision-only vs. touch-only vs. combined (vision + touch) classification of surface materials. The combined system achieved 97% accuracy, compared to 88% for vision alone and 93% for touch alone
>
> **Why it matters:** This proves that multimodal perception — combining senses like the brain does — dramatically improves performance

Application 2: Healthcare Monitoring

NII could monitor elderly patients for early signs of cognitive decline. By tracking daily behavior patterns and speech variations, the system could detect anomalies that suggest dementia or mild cognitive impairment.

> **How it works:** The perception unit captures motion, audio, and video data. The attention mechanism highlights behavioral changes. The memory module tracks historical patterns. The reasoning engine infers whether changes indicate decline.

Application 3: Industrial Safety

In smart factories, NII could predict unsafe behavior or equipment failure before accidents happen. The system would analyze real-time sensor data, learn normal patterns, and flag deviations.

> **Key feature:** The adaptation layer continuously adjusts safety thresholds based on environmental feedback, ensuring the system stays responsive to site-specific risks.

Application 4: Personalized Education

NII could power intelligent tutoring systems that adapt to each student's learning style, attention span, and knowledge level. Unlike current systems that just adjust difficulty, NII would model the student's cognitive state and provide personalized feedback.

> **How it works:** The attention mechanism detects when a student is losing focus. The memory module tracks what the student has learned. The reasoning engine identifies misconceptions and creates individualized learning paths.

Application 5: Adversarial Robustness

The paper demonstrates how NII could resist adversarial attacks. By tracking how each added perturbation affects classification confidence, the system can dynamically adjust its decision-making.

> **Experiment:** The researchers incrementally added adversarial noise pixels to images and tracked classification probabilities. The system could monitor these changes in real time, suggesting potential for neurocognitive-inspired AI that adjusts dynamically to changing stimuli.

The Pattern: In every application, NII's advantage comes from the same sources: multimodal perception, explicit memory, adaptive attention, and closed-loop feedback. These are the mechanisms that make biological intelligence robust, and NII brings them to artificial systems.

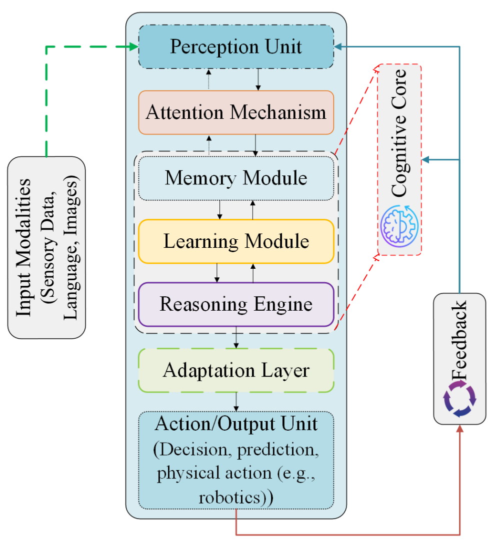

7. Building NII: Technical Details

The paper discusses several technical strategies for implementing NII:

Modular Design

Each module can be independently developed and optimized. This mirrors the brain's organization, where different regions handle different functions but work together seamlessly.

> **Examples:** SPAUN (a brain model with 2.5 million neurons) demonstrates how high-level cognition can emerge from modular interactions. LEABRA combines different learning rules to mimic cognitive behaviors.

Neuromorphic Hardware

Running NII on traditional computers would be expensive. Neuromorphic chips — hardware designed to mimic neural processing — could make NII practical.

- **Intel Loihi:** Simulates neural spikes for energy-efficient, event-driven computation
- **SpiNNaker:** Designed for real-time simulation of large-scale neural networks

Learning Strategies

NII uses several learning approaches inspired by biology:

- **Continual learning:** Learning new tasks without forgetting old ones, using elastic weight consolidation (protecting important connections)
- **Few-shot learning:** Learning from very few examples by leveraging memory and prior experience
- **Meta-learning:** Learning how to learn — developing strategies that improve over time

Evaluation

Conventional metrics (accuracy, F1-score) aren't enough for NII. The paper proposes new evaluation criteria that measure cognitive capabilities:

- **Generalization:** How well does the system handle new situations?
- **Memory retention:** Does it remember what it learned?
- **Attention dynamics:** Does it focus on the right things?
- **Explainability:** Can it explain its decisions?
- **Embodied interaction:** Does it learn from physical experience?

8. Challenges and Limitations

The paper is honest about the difficulties ahead:

Hardware Limitations

NII requires significant computational resources. The seven modules, each with their own parameters, demand more memory and processing power than current AI systems.

Data Scarcity

Multimodal datasets that combine vision, language, memory, and motor data are rare. Most AI research uses single-modality datasets (just images or just text).

Integration Complexity

Combining symbolic reasoning, neural networks, and sensorimotor control in a single system is extremely challenging. Each component uses different mathematical frameworks.

Simulation Challenges

Testing NII requires realistic simulated environments that support perception-action loops. Current simulators often lack the fidelity needed for embodied cognition research.

Continual Learning

Preventing catastrophic forgetting while continuously learning new tasks remains one of the hardest problems in AI. While NII proposes biological solutions, they're not yet proven at scale.

> **Bottom Line:** NII is a theoretical blueprint, not a finished system. It's a research agenda that invites future exploration, implementation, and validation. The authors are clear that achieving NII will require significant interdisciplinary collaboration.

9. How All Seven Modules Work Together

The real power of NII isn't in any single module — it's in how they all connect. Let's walk through a concrete example to see the framework in action.

Example: A Robot Making Coffee

Imagine a neurocognitive robot tasked with making coffee in an unfamiliar kitchen. Here's how each module contributes:

1. Perception Unit: The robot scans the kitchen, identifying the counter, cabinets, coffee maker, and various objects. It builds a 3D map of the environment using visual and tactile sensors.

2. Attention Mechanism: The robot focuses on the coffee-related objects — ignoring the toaster, radio, and other irrelevant items. It prioritizes the coffee maker and nearby cabinets.

3. Memory Module: The robot retrieves memories from previous kitchens it has worked in. It remembers that coffee is usually in upper cabinets and mugs are often near the sink.

4. Learning Module: As the robot explores, it notices this kitchen has a different layout than previous ones. It updates its internal model to account for this variation.

5. Reasoning Engine: The robot combines pattern recognition (identifying objects) with logical reasoning (coffee needs hot water, which requires the machine to be on). It creates a plan.

6. Adaptation Layer: The robot notices the coffee maker is different from what it expects. It slows down, examines the machine more carefully, and adjusts its approach.

7. Action Unit: The robot executes the plan — opening cabinets, finding coffee, measuring water, and operating the machine. Each action produces feedback that improves future performance.

This is fundamentally different from how current AI works. A traditional robot would follow a fixed script or rely on statistical patterns. A neurocognitive robot *thinks through* the task, adapting as it goes and learning for next time.

10. Why Feedback Loops Change Everything

The most important innovation in NII is the feedback loop. In traditional AI, information flows in one direction: input → processing → output. In NII, every output feeds back as new input.

Three Types of Feedback

- **Immediate feedback:** When the robot touches an object, the tactile information immediately updates its perception. This is like how touching a hot stove instantly changes your behavior.
- **Outcome feedback:** After completing an action, the system evaluates whether it succeeded. If the robot spilled the coffee, it adjusts its motor control for next time.
- **Long-term feedback:** Over many experiences, patterns emerge that update the system's fundamental knowledge. The robot learns that most kitchens have coffee near the sink.

How Feedback Enables Learning

In the brain, feedback happens through multiple mechanisms:

- **Dopamine signals:** When you receive a reward (or unexpectedly don't), dopamine neurons fire to update your value estimates. NII uses similar reward-prediction error signals.
- **Prediction errors:** When the brain predicts something will happen and it doesn't, the mismatch drives learning. NII incorporates this through its adaptation layer.
- **Social feedback:** Humans learn enormously from watching others. NII could incorporate demonstration learning where the system observes and imitates expert behavior.

Why Traditional AI Lacks This

Current AI systems are typically trained once and deployed. They don't improve through experience. When they encounter new situations, they either handle them based on training data or fail. NII changes this by creating systems that learn continuously from their own actions and outcomes.

11. The Power of Multiple Senses

One of the most compelling demonstrations in the paper is the vision-touch experiment. The researchers showed that combining multiple senses dramatically improves performance.

The Experiment

The team tested three approaches for classifying surface materials (concrete, plastic, wood, metal, etc.):

- **Vision only:** 88% accuracy — struggles with occlusion, poor lighting, and similar-looking surfaces
- **Touch only:** 93% accuracy — better than vision, but misses visual context
- **Vision + Touch:** 97% accuracy — combines the strengths of both senses

Why This Works

When you combine senses, each modality fills in gaps left by the others. Vision provides context (is this surface inside or outside?), while touch provides texture details (is this rough or smooth?). Together, they create a more complete picture than either alone.

This is exactly how the brain works. You don't just see a ball — you feel its weight, hear its bounce, and track its trajectory simultaneously. This multimodal integration makes perception robust and accurate.

Beyond Vision and Touch

The paper discusses how NII could integrate many more modalities:

- **Vision + Auditory:** Understanding speech in noisy environments by combining lip reading with audio
- **Vision + Proprioception:** Knowing where your body is in space while navigating
- **Vision + Language:** Grounding words in visual experience (understanding "red" by seeing red)
- **Vision + Thermal + Tactile:** Detecting hazards by combining heat, texture, and visual cues

Implications for AI

Most current AI works with single modalities — text models process text, image models process images. NII argues that true intelligence requires integrating multiple senses into a unified understanding. This is a fundamental shift in how we think about AI architecture.

12. Dual-Process Reasoning: Fast and Slow Thinking

The reasoning engine is perhaps the most innovative module in NII. It combines two types of thinking that mirror human cognition.

System 1: Fast Pattern Matching

This is the intuitive, automatic thinking that happens in milliseconds. When you see a face and instantly recognize a friend, that's System 1. In AI terms, this is pattern recognition using neural networks — fast, efficient, but sometimes wrong.

System 2: Slow Deliberate Reasoning

This is the careful, logical thinking that takes effort. When you solve a math problem or plan a route through a new city, that's System 2. In AI terms, this is symbolic reasoning — using rules, logic, and planning graphs.

How They Work Together

In NII, the two systems don't operate independently — they collaborate:

- **Pattern recognition suggests hypotheses:** "This looks like a coffee maker"
- **Logic tests the hypothesis:** "It has a water reservoir and filter basket, so yes, it's a coffee maker"
- **Logic generates new questions:** "What type of coffee does it make? How do I operate it?"
- **Pattern recognition provides answers:** "It looks like a drip machine based on its shape"

Why This Matters

Current AI is mostly System 1 — it recognizes patterns but can't reason logically about them. A chatbot might recognize that a question is about math but can't actually solve the problem. NII's dual-process approach could enable AI that both recognizes *and* reasons.

13. How NII Compares to Other Cognitive Systems

NII isn't the first attempt to build brain-like AI. The paper compares it to several existing approaches:

ACT-R and SOAR

These are classic cognitive architectures that attempt to model human cognition using symbolic rules. They're good at reasoning but struggle with perception and learning from data.

> **NII advantage:** Combines symbolic reasoning with neural pattern recognition, getting the best of both worlds

AlphaStar and Gato

These DeepMind systems demonstrate cross-modal functionality — handling multiple tasks with a single model. They're impressive but still operate as black boxes without transparent reasoning.

> **NII advantage:** Built-in transparency through explicit cognitive modules

Brain-Inspired Neural Networks

These systems copy the brain's structure (layers, connections, learning rules) but not its thinking process. They're still essentially pattern matchers.

> **NII advantage:** Copies the thinking process, not just the structure

Cognitive Neuroscience Models

These are detailed simulations of specific brain regions (like the hippocampus or prefrontal cortex). They're accurate but too specialized to build general AI.

> **NII advantage:** Integrates insights from many brain regions into a unified framework

> **NII's Unique Position:** It's not just another brain-inspired system. It's a comprehensive framework that integrates perception, memory, attention, reasoning, and action into a closed loop — something no existing system fully achieves.

14. What Needs to Happen Next

The paper identifies several concrete steps needed to move from theory to practice:

Step 1: Build Multimodal Datasets

NII requires datasets that combine vision, touch, audio, and motor data. Current datasets are single-modality. Creating rich, multimodal datasets is essential for training and evaluating NII systems.

Step 2: Develop Modular Toolkits

Researchers need software frameworks that make it easy to build and test individual NII modules. Toolkits like OpenCog and NengoDL are early attempts, but more comprehensive tools are needed.

Step 3: Create Realistic Testbeds

NII needs simulated environments that support perception-action loops. Current simulators often lack the physical realism and sensor fidelity needed for embodied cognition research.

Step 4: Invest in Neuromorphic Hardware

Running NII on traditional computers would be prohibitively expensive. Neuromorphic chips like Intel Loihi and SpiNNaker could make NII practical, but they need significant investment.

Step 5: Foster Interdisciplinary Collaboration

NII requires expertise from neuroscience, cognitive science, computer science, and engineering. No single research group has all these skills. New collaboration models are needed.

The Timeline: The paper suggests that basic NII components could be demonstrated within 3-5 years, with full systems taking 10-15 years. This is ambitious but achievable with sustained investment and collaboration.

15. Why This Matters for Everyone

NII isn't just an academic exercise — it has profound implications for society:

Healthcare

NII-powered systems could monitor patients continuously, detecting early signs of disease before symptoms appear. They could adapt to each patient's unique physiology and lifestyle.

Education

Intelligent tutors could understand how each student thinks, not just what they know. They could adapt teaching strategies in real time based on attention, memory, and comprehension.

Safety

Autonomous vehicles and robots that can reason about uncertainty, learn from mistakes, and adapt to new situations would be dramatically safer than current systems.

Accessibility

AI that truly understands context and intent could revolutionize assistive technology — helping people with disabilities communicate, navigate, and interact with the world.

Scientific Discovery

NII systems that can form hypotheses, design experiments, and learn from results could accelerate scientific progress across every field.

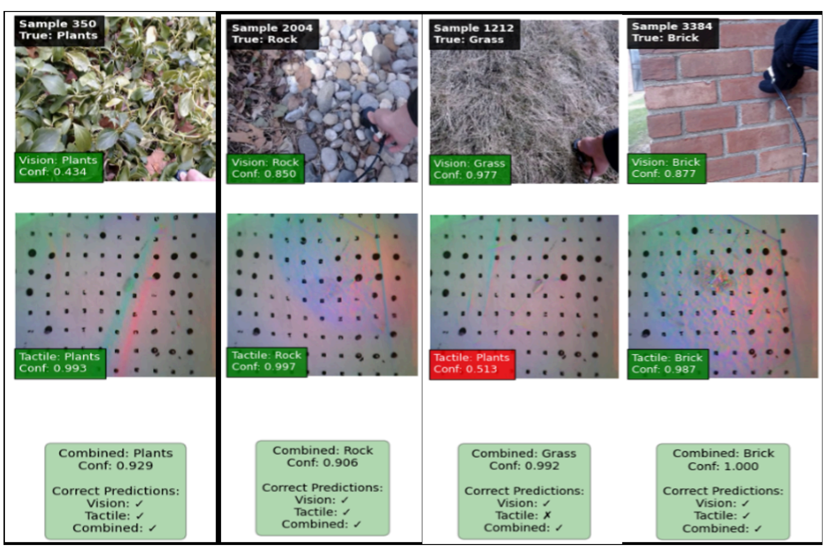

16. Conclusion: What This Means for the Future

This paper argues that the next breakthrough in AI won't come from bigger models or more data. It will come from fundamentally rethinking how AI thinks.

Current AI is like a sophisticated calculator — brilliant at specific tasks but clueless about the world. NII proposes building AI that perceives, remembers, reasons, and adapts like a human.

The three key insights are:

1. **Structure isn't enough:** Copying the brain's architecture (layers of neurons) isn't the same as copying how it thinks
2. **Cognition is a loop:** Real intelligence isn't a one-way pipeline — it's a continuous cycle of perception, reasoning, action, and feedback
3. **Integration is key:** The brain doesn't have separate systems for vision, language, and action — it integrates everything into a unified experience

The journey from theory to practice will be long and nonlinear. But by grounding AI in the principles of natural cognition, we move closer to building machines that don't just process information — they understand it.

> **The Ultimate Goal:** AI that doesn't just mimic intelligence but begins to sense and understand it. Machines that observe, plan, evaluate, and even experience in human-like ways. The neurocognitive intelligence framework could define the next era of artificial intelligence.

17. The Science Behind NII

NII isn't just a conceptual framework — it's grounded in decades of neuroscience and cognitive science research. Here are the key scientific principles that inform each module:

Perception: From V1 to IT Cortex

The brain's visual system processes information hierarchically. Simple cells in area V1 detect edges and basic features. As information flows through V2, V4, and into the inferotemporal (IT) cortex, representations become increasingly complex — from edges to textures to shapes to objects.

NII's perception unit mirrors this hierarchy, using convolutional filters at low levels and increasingly abstract representations at higher levels.

Attention: Prefrontal-Parietal Networks

In the brain, attention is controlled by networks connecting the prefrontal cortex (goals) with the posterior parietal cortex (spatial processing). The thalamus acts as a relay station, filtering sensory information based on relevance.

NII's attention mechanism mimics this dual structure, combining goal-directed modulation with stimulus-driven salience detection.

Memory: Hippocampal-Cortical Interactions

The hippocampus is crucial for forming new episodic memories, while the neocortex stores long-term semantic knowledge. During sleep, the hippocampus 'replays' experiences, transferring important information to cortical storage.

NII's memory module implements this dual-system architecture with explicit consolidation mechanisms.

Reasoning: Dual-Process Theory

Cognitive scientists distinguish between System 1 (fast, automatic) and System 2 (slow, deliberate) thinking. System 1 handles routine tasks efficiently. System 2 engages for novel or complex problems.

NII's reasoning engine implements both systems, allowing rapid pattern recognition when possible and deliberate reasoning when needed.

Adaptation: Neuromodulatory Systems

The brain uses chemical signals (dopamine, serotonin, norepinephrine) to regulate learning, attention, and motivation. These neuromodulators help the brain decide what to learn, when to explore, and when to play it safe.

NII's adaptation layer incorporates similar mechanisms, using reward signals and uncertainty estimates to guide learning and behavior.

18. Key Differences from Previous Approaches

To understand why NII represents a genuine advance, it helps to compare it with earlier attempts at brain-like AI:

NII vs. Deep Learning

Deep learning (the foundation of current AI) is excellent at pattern recognition but lacks the cognitive machinery for reasoning, memory, and adaptation. NII builds on deep learning's strengths while adding the missing pieces.

Deep learning: Recognizes patterns in data

NII: Recognizes patterns, reasons about them, remembers outcomes, and adapts behavior

NII vs. Symbolic AI

Symbolic AI (the original approach to AI) excels at logical reasoning but struggles with perception and learning from data. NII combines symbolic reasoning with neural pattern recognition.

Symbolic AI: Reasons with rules but can't see or learn from experience

NII: Reasons with rules AND perceives AND learns from experience

NII vs. Brain-Inspired AI

Brain-inspired AI copies the brain's structure (neurons, synapses, layers) but not its thinking process. NII focuses on functional mimicry — copying how the brain thinks, not just how it looks.

Brain-inspired AI: Looks like the brain but doesn't think like it

NII: Thinks like the brain, even if the implementation details differ

NII vs. Cognitive Architectures

Cognitive architectures (like ACT-R and SOAR) attempt to model human cognition but often lack the learning capabilities of modern AI. NII integrates cognitive principles with state-of-the-art machine learning.

Cognitive architectures: Model cognition but struggle with learning from data

NII: Models cognition AND learns from data AND adapts continuously

> **The Bottom Line:** NII isn't just another approach — it's a synthesis that combines the best aspects of multiple traditions: the pattern recognition of deep learning, the reasoning of symbolic AI, the principles of cognitive science, and the adaptability of biological intelligence.

18. Glossary of Key Terms

Neurocognitive-Inspired Intelligence (NII)

An AI framework that mimics the brain's thinking process — not just its structure — to create systems that perceive, reason, remember, and adapt.

Embodied Cognition

The theory that intelligence requires physical interaction with the world, not just abstract computation. The brain's thinking is shaped by having a body.

Perception-Action Loop

A closed cycle where sensing informs thinking, thinking guides action, and action produces new sensory information. Fundamental to biological intelligence.

Predictive Coding

A theory that the brain constantly predicts incoming sensory input and learns from prediction errors. Makes the brain efficient and adaptive.

Working Memory

A temporary mental workspace for holding and manipulating information during active thinking. Like a computer's RAM.

Episodic Memory

Memory of specific events — what happened, when, and where. Allows learning from personal experience.

Semantic Memory

General knowledge and facts about the world. Stores what you know, not how you learned it.

Attention Mechanism

A filter that selects which information gets processed and which gets ignored. Combines bottom-up (stimulus-driven) and top-down (goal-directed) selection.

Continual Learning

The ability to learn new tasks without forgetting old ones. A major challenge in current AI, solved in the brain through synaptic consolidation.

Catastrophic Forgetting

The tendency of neural networks to completely forget previously learned information when trained on new data.

Neuromorphic Computing

Computer hardware designed to mimic neural processing — using spikes, sparse computation, and event-driven architecture.

Symbolic Reasoning

Logical reasoning using rules, relationships, and abstract concepts. Contrasts with sub-symbolic (pattern-based) reasoning.

Sub-symbolic Reasoning

Pattern recognition using neural networks — fast, intuitive, but sometimes unreliable. Contrasts with symbolic (logical) reasoning.

Neuroplasticity

The brain's ability to reorganize its connections in response to experience, learning, or injury. The basis of biological adaptation.

Multi-modal Integration

Combining information from different senses (vision, touch, hearing) to create a richer understanding of the world.

Common Sense Reasoning

Intuitive understanding of physical and social realities that humans develop early in life but AI lacks.

World Model

An internal simulation of how the world works, used to predict outcomes and plan actions.

Adversarial Attack

Tiny, carefully crafted changes to input data that fool AI systems into making confident but wrong predictions.

Cognitive Architecture

A blueprint for organizing AI modules (perception, memory, reasoning) into a unified system that mimics human cognition.

Elastic Weight Consolidation (EWC)

A technique for preventing catastrophic forgetting by protecting important neural network weights during new learning.

Few-Shot Learning

Learning new concepts from just one or a few examples, like humans do.

Meta-Learning

Learning how to learn — developing strategies that improve learning efficiency over time.

Hebbian Learning

A learning rule where neurons that fire together strengthen their connections. The basis of biological learning.

Active Perception

Directed sensing where the system chooses what to look at or explore, rather than passively processing whatever input arrives.

About This Simplification

This document simplifies the paper "Towards Neurocognitive-Inspired Intelligence" (arXiv:2510.13826) by Golilarz, Al Khatib & Rahimi, originally 25 pages of technical content. The goal was to make the ideas accessible to a general audience without losing the essential content.

Simplified by: @muhamedfazalps7
Original paper: arXiv:2510.13826