


# LUMINAI™

## Cognitive Reasoning Engine

Harness the Power of Cognitive AI

## ABSTRACT

Machine learning has come to represent AI in the popular imagination. But machine learning, as with other conventional numerical AI techniques, exhibits rigidity, brittleness, and lack of explainability. These limitations make conventional numerical AI approaches insufficient to solve complex problems that require flexible reasoning and decision-making. For more demanding applications dealing with high-value industrial, medical, and financial assets, a new generation of AI technology is required. Here we present the application of advanced Cognitive AI, utilizing symbolic reasoning in combination with numerical AI to power knowledge-based reasoning systems that can think more like humans and explain their answers clearly.



“ If you think of AI as a stack of capability, at the base layer of the pyramid is data analytics, and at the pinnacle is cognitive computing. Cognitive computing mimics human reasoning and will use knowledge provided by domain experts, not just data, to understand situations, solve problems and recommend actions. Cognitive ‘agents’ differ from standard artificial intelligence applications as they can adapt and become smarter over time as they interact with more experts, problems and data. Cognitive ‘agents’ are able to operate in complex situations where uncertainty exists and data may not be as prevalent as we would like, and in such situations, they approach problems much as a human would. ”

Paul Stone- BP Digital Innovation

## INTRODUCTION

### AI – FROM THE R&D LAB TO THE FACTORY FLOOR

As a technology category, Artificial Intelligence is unfortunately named. For many years, AI systems have been referred to as AI, no matter what was under the hood. Today, the AI landscape has gotten more complex with new phrases such as Machine Learning, Deep Learning, Numeric AI, Symbolic AI, Hybrid AI, Cognitive AI, Expert Systems, etc. Simply put, they all describe a computer program that exhibits some degree of intelligence, regardless of the underlying mechanisms.

In the last ten years, AI has become virtually synonymous with numeric machine learning or deep learning approaches. As neural networks and deep learning move from the lab into production in mission-critical fields from medicine to driverless cars, we must recognize the very real limitations of such software code and statistical systems, which are not the learning and thinking intelligences we might imagine them to be.

Many AI researchers have applied AI techniques to solve theoretical exercises such as winning board games from chess to Go to Atari. But industrial companies around the world are actively seeking to put AI technologies to work in service of increased efficiency, reduced waste, predictive maintenance, process optimization, and tangible ROI. What AI technology (or combination of technologies) is best suited to handle complex mission-critical problems?

## BEYOND LIMITS TECHNOLOGY IP

- + Beyond Limits has filed 4 patent applications in the 2017-2020 timeframe. Named inventors: Beyond Limits / Drs. Shahram Farhadi Nia, Azarang Golmohammadi, Zack Nolan, and Mark L. James.
- + CTO Mark James owns two patents for an advanced AI system called BEAM (Beacon-based Exception Analysis for Multimissions) developed for NASA deep space missions while he worked at the Caltech Jet Propulsion Laboratory.
- + Beyond Limits has exclusive licenses to 70 IP blocks (the term NASA uses to describe technologies, rather than “patents”) of technologies developed for NASA space missions by the Caltech Jet Propulsion Laboratory.
- + Since 2014, Beyond Limits has internally developed an additional 134 technology IP blocks to solve complex problems for our customers. These innovations are bound by confidentiality and have not been externally published.
- + Technical publications describing applications of our technology, including 54 papers authored by Beyond Limits CTO Mark James and 420 technical publications describing our licensed technologies are located here: <https://bit.ly/2TsOWPe>

## WHAT COMES AFTER NEURAL NETWORKS?

From a high-level, artificial intelligence can be broken into a combination of two broad camps: Numeric (training from data) and Symbolic (education using book-like knowledge). In a nutshell, Numeric systems find features in data, whereas Symbolic systems find meaning from the features.

In a numeric approach using neural networks and deep learning, a computer model learns to perform classification tasks directly from images, text, or sound. They are trained using a large set of labeled data, so if you want to train it to recognize a cat, you would feed it large numbers of cat-centric images. The symbolic approach says the best way to teach an AI is to feed it human-readable information related to what you think it needs to know. For example, if you wanted to create an AI to assist a doctor you would feed it a collection of medical textbooks so it could reason from logical principles.

Neural nets do not think. Beyond Limits technology does think; and (in addition) our systems generate an audit trail so they can explain themselves and be applied to high-value assets. Our systems, like the LUMINAI™ Cognitive AI Engine, allows you to represent conflicting information, and reason in the presence of missing and misleading data. Even when a particular knowledge base is incomplete, it's able to repurpose it in pieces by converting it and understanding the knowledge at a conceptual level then repurposing it to solve the problem.

- + **Numeric systems do not think. They find patterns in data. They answer the “What” questions.**
- + **Symbolic systems find meaning in patterns. They answer the “Why” questions.**
- + **Cognitive systems think. They use numeric and symbolic techniques to answer the “What do I do next” questions and solve complex problems.**

At the moment, deep learning systems are in vogue and much work has been done to make them as intelligent and fast as possible. Even though these systems significantly speed up numeric-based approaches, unfortunately, they only help with half of the problem – they cannot speed up the symbolic AI algorithms. One of their problems is that deep learning cannot be used for single-event phenomena or where an explanation of how they derived their answers is required (e.g., medical, high-value assets, fintech).

In this light, the term Symbolic AI deserves new importance as it highlights the fact that the underlying mechanism in AI is symbolic in nature rather than purely numeric. Taking it a step further, Hybrid AI is an AI system that possesses both symbolic and numeric elements to exhibit intelligence. The term Cognitive AI is used to refer to a Hybrid AI system that includes bio-inspired algorithms to provide the resulting system with an extra degree of intelligence.



“ Cognitive Artificial Intelligence has the human-like ability to perceive, understand, correlate, learn, teach, reason, and solve problems faster than conventional AI solutions. ”

Kimberly Gilbert, PhD - Director of Technical Commercial Engineering

## LIMITATIONS OF CONVENTIONAL AI LIKE MACHINE LEARNING AND DEEP LEARNING

Historically, symbolic and numeric computation have pursued different lines of evolution, have been written in different languages, and are generally seen to be competitive rather than complementary techniques. Even when both were used to solve a problem, ad hoc methods were used to transfer the data between them. Since the methods appear to address different needs, an influential position has emerged which argues that the methods need to be combined.




Figure 1. At Beyond Limits, we have combined numeric and symbolic into a new generation of technology: Cognitive AI.

## THE COGNITIVE DIFFERENCE

A neural network of today no more “learns” or “reasons” about the world than a linear regression of the past; it merely induces patterns through statistics. Those patterns may be more opaque, more mediated, and more automatic than historical approaches and capable of representing more complex statistical phenomena, but they are still merely mathematical incarnations, not intelligent entities, no matter how spectacular their results.

Whether neural network, Naïve Bayes, or simply linear regression, data scientists train their machine learning models on carefully constructed piles of training examples then claim their algorithms have “learned” about the world. In reality, machine learning is merely another form of machine instruction, different from purely expert manual coding of rules, but still guided with the algorithms and workflows manually tuned for each application.

Why does this matter? It matters because as we increasingly deploy AI systems into mission-critical applications directly affecting human life, from driverless cars to medicine, we must understand their very real limitations and brittleness in order to properly understand their risks.

### TWO SIDES OF THE COGNITIVE BRAIN

Our favorite definition for cognitive intelligence:

*“The ability to learn new things, recall information, think rationally, apply knowledge, and solve problems.” (Kaplan & Sadock, 1991)*

From our perspective at Beyond Limits, AI can be divided into two distinct areas: Numeric-based AI and Symbolic-based AI. Unlike conventional AI providers, we use both for a true best-in-class approach.

#### Numeric-based AI

Numeric-based systems (machine learning, neural nets, and deep learning, etc.), are trained from thousands of examples to reduce a problem to sophisticated pattern matching rather than a reasoning-based process. They are wonderful when you have lots of data to train a system. These systems, unfortunately, are “black boxes” that cannot explain how they arrive at their answers, so they literally don’t know what they know. Nor do they know what they don’t know.

#### Symbolic-based AI

Symbolic-based AI systems are educated from declarative knowledge (case-based reasoning, rule-based inference, etc.) much as a person would be, and deeply think about each fact that is presented to them. Symbolic systems can reason through ambiguity and missing information with human-like intuition and can provide a detailed explanation for how they arrived at their answers (known as Explainable AI or XAI). Symbolic systems can also make changes to their knowledge base in real-time, making the system smarter.

Both approaches are valuable AI techniques and typically both are required for complex problems that involve massive, disparate data sets and high-value assets like those found in energy, manufacturing, finance, and healthcare.

# BEYOND LIMITS LUMINAI™ COGNITIVE AI REASONING ENGINE

Against this backdrop, the LUMINAI™ system was developed at Beyond Limits. LUMINAI™ utilizes a symbolic reasoner, which means it uses the outputs from sensors and neural nets and applies its *education* to reason about what it sees. LUMINAI™ is a new approach to reasoning based on opportunistic self-discovery monitoring, bi-modal cognitive-based reasoning, and autonomous self-discovery to resolve ambiguities.

Since we do not live in a perfect world, LUMINAI™ combines numeric AI-based approaches with symbolic ones to efficiently manage uncertainty in the presence of missing and misleading data. A complete approach to reasoning under uncertainty requires support for incremental and interactive formulation and revision of, as well as reasoning with, models of the problem domain capable of representing our uncertainty. We present a hybrid reasoning scheme that combines symbolic and numeric methods for uncertainty management to provide efficient and effective support for each of these tasks.




Figure 2. LUMINAI™ Cognitive AI Engine - a hybrid reasoning scheme which combines symbolic and numeric methods for uncertainty management and explainability.

LUMINAI™ has the capability to autonomously shift through corridors of data to discover plausible facts and scenarios using interpretations from diverse data sources. It avoids the usual computational complexities of traditional systems by using a technique called autonomic monitoring that enables serendipitous discovery during cognitive forensic analysis.

It resolves ambiguity, peruses promising hypotheses, and activates user-supplied instruments (e.g., sensors, tools, probes, etc.) to gather supplementary data to follow a hunch or narrow its scope. One of its advanced capabilities involves the ability to simultaneously and efficiently reason in both the forward (e.g., *Is this valuable?*) and backward (e.g., *How do I find something valuable?*) directions (bi-modal reasoning) to dissect and tackle a problem from heterogeneous perspectives while simultaneously reasoning and learning.

## WHAT IS COGNITIVE AI?

Cognitive Artificial Intelligence technology goes beyond conventional AI by utilizing a unique hybrid combination of machine learning numeric approaches alongside higher-order symbolic techniques; a method that delivers cognitive reasoning and intelligence resembling that of human intuition. Our Cognitive AI engines are trained on data (with the ability to still function regardless of missing or misleading data) and educated by expert human knowledge to provide clear guidance for everyday people, leading to a faster, more profound decision-making process in complex scenarios.




Figure 3. Hybrid Approach That Combines Numeric with Symbolic Technologies

## BIOLOGICALLY INSPIRED PREDICTIVE SYSTEM

At Beyond Limits, we define *Cognitive Intelligence* as the idea of drawing insights from a cognitive model according to consistent analytic principles (deductive, inductive, and abductive rules) that allow us to construct new beliefs from known inviolable truths (axioms). We combine a system that is educated with a trained neural net-based system, providing a significant improvement in reducing the amount of time and data required to have a system that exhibits a high degree of cognition. This approach works first by having the LUMINA<sup>AI</sup>™ system, which is educated, being taught the high-level aspects of the domain, and then training a neural net on identifying the features of the domain.

For example, rather than training a system for every possible configuration of a human, you first educate it with the information that a human has a head, torso, arms and legs, and the subsequent relevance for each of those features. Then you train the neural net to recognize those features individually and the contains from the symbolic knowledge only considers realistic instances of a human. This unique combination provides a hybrid approach of fusing grounded a priori knowledge coupled with heuristic learning-based pattern matching. The combination of symbolic education and data-centric training reduces the number of ground instances that are necessary to have a high degree of recognition. We believe this approach approximates the observed reasoning processes of how a person would solve a problem in unfamiliar or unknown territory.

- + **Since we do not live in a perfect world, LUMINA<sup>AI</sup> combines numeric AI-based approaches with symbolic ones to efficiently manage uncertainty in the presence of missing and misleading data.**




Figure 4. LUMINAI™ application of data and knowledge

## HYBRID APPROACH TO TRAINING AND LEARNING

Our approach to building a hybrid AI system like LUMINAI™ is to use a cognitive reasoner – the element that is responsible for figuring out how to solve a problem. The reasoner is educated with domain expertise and then trains neural nets to identify features from data. The resulting output is actionable intelligence annotated with human and machine understandable audit trails.





Figure 5. Beyond Limits LUMINAI™ Cognitive Reasoner high-level architecture

For a reasoner to do anything, it needs to be fitted with an Education, i.e., knowledge about a domain and how to solve a problem. A Cognitive Education usually takes the form of a combination of symbolic and numeric knowledge. Symbolic knowledge is encoded as facts, rules, algorithms, experience, heuristics, analogies, examples, and workflows. Numeric knowledge is encoded as training sets from ML algorithms, e.g., Deep Learning, Support Vector Machines, Bayesian Networks.

- + It uses the outputs from Deep Learning elements as its inputs, and reasons using deductive, abductive, and inductive logic coupled with case-based approximate reasoning to perform its function. It solves problems using general and domain-specific knowledge, which means it can think outside of the box in the presence of missing and misleading information to intelligently explain and refine its solutions.

- + It uses the aggregated results from its knowledge to derive actionable insight that is situationally relevant and, if data is insufficient to solve a problem then it becomes an expert toolsmith to generate workflows to synthesize derived sources of data.
- + Its strength over other approaches (e.g., simple rules, Deep Learning alone, etc.) is that it provides a domain-independent method to generate valuable results with respect to completeness and truth.
- + Building a system that is both educated with declarative knowledge and trained with data provides a significant improvement in considerably reducing the amount of time and data required to have a system that exhibits a high degree of cognition.
- + This hybrid approach does not reduce the value of neural nets but adds symbolic layers to reduce the limitations of neural net-only approaches and augmenting them with computational intelligence (reasoning) and complete explainability of their results. This is important because neural nets alone do not think but efficiently extract labels from data.



“ Fulfilling the promise and massive potential of artificial intelligence will require human-like reasoning, far beyond what most AI systems can currently provide. Beyond Limits has taken the proven AI used in deep space and commercialized it for use in solving our world’s most pressing problems. ”

AJ Abdallat - Beyond Limits CEO

## COGNITIVE AI IS NOT 1970s EXPERT SYSTEMS

In the 1970s, AI systems represented knowledge using rules embedded inside of expert systems, and these systems are the simplest example of artificial intelligence. They used ordinary IF-THEN rules as their only means to encode wisdom and were designed to solve a specific problem, not general classes of problems. This approach implemented a solution as essentially a dynamic decision tree, which required a domain expert to encode all the necessary steps in advance to solve the problem. These systems were designed to mimic the behavior of an expert; but because they were unable to think, they were very brittle and broke when the data or questions changed slightly. They weren’t even able to adapt to small changes in the way questions were posed or shifts in the data.

Rule-based Reasoning (RBR) and Cognitive-based Reasoning (CBR) are two complementary alternatives for building knowledge-based “intelligent” decision-support systems. The first approach is closely related to expert systems. Expert Systems (ES) are typically defined as computer programs that emulate the decision-making ability of a human expert. The power of an ES is derived from the presence of a knowledge base filled with expert knowledge, mostly in symbolic form. In addition, there is a generic problem-solving mechanism used as the inference engine. Some other typical features of expert systems include uncertainty processing, dialogue mode of the consultation, and explanation abilities. Beside ES dedicated to specific applications, “empty” expert systems (also called “shells”) have been developed, which can be coupled with an arbitrary knowledge base encoded in an appropriate format. Research in the area of expert systems started in the mid-1970s. Classical examples of early systems that influenced other researchers are Mycin and Prospector.

The knowledge of an expert is usually represented in the form of IF-THEN rules, which are applied in a deductive way: if the condition of a rule is satisfied, then this rule can be applied to either

derive some conclusion or to perform the respective actions. The central point of all these systems was the compositional approach to inference, allowing us to compose the contributions of multiple rules (leading to the same conclusion) using a uniform combination function, regardless of their mutual dependencies. This approach was later subjected to criticism by most of the uncertainty-processing community, which resulted in ES research grounds dominated by probabilistic approaches. However, although probabilistic reasoning is more sounder, it is also much harder for casual users to understand than rule-based reasoning.

The opportunity for practical applications of probabilistic systems that require direct capture of human expertise (rather than to learn from a large set of cases) is thus significantly limited. Many developers of expert system applications decide to ignore uncertainty modeling altogether or only adopt scoring schemes, which unfortunately have zero capacity for explanation and thus little potential for sharing.


The so-called knowledge acquisition bottleneck (i.e., the problem of eliciting the domain-specific knowledge from experts in the form of sufficiently general rules) has several workarounds. One of them is to use machine learning techniques to acquire knowledge from data representing situations successfully solved in the past. In addition, to augment this approach with case-based reasoning, where the knowledge is represented in the form of prototype problems and their solutions (so-called "cases"). This seems to be a more psychologically plausible model of human reasoning than using rules as we do in classical rule-based (expert) systems.

## BUILDING A COGNITIVE SYSTEM

### Step 1: Train ML Models to Reduce Raw Data to Features

- + Train ML models to recognize the phenomena you wish to detect; whose output is labeled Features
- + This is a data-intensive task requiring potentially lots of example data from different perspectives
- + It is difficult because if you overfit, you lose the ability to be general, and if you underfit, you miss the phenomena
- + The addition of a symbolic reasoner reduces the number and complexity of ML models and simplifies training significantly

### DEEP LEARNING APPROACH



Step 1 - Train on Pictures to Produce Features

### Step 2: Educate the Reasoner with Domain Knowledge

- + You educate a symbolic reasoner with domain expertise; whose output is actionable intelligence with audit trails


- + Knowledge is encoded as facts, rules, algorithms, experience, heuristics, analogies, examples, and workflows
- + Uses Features as its inputs, and reasons using deductive, abductive, and inductive logic coupled with case-based approximate reasoning

## Educate Using Textbook-like Knowledge *(one of several possible ways)*

1. A Family has one or two Adults and sometimes have one or more Children and/or Pets. Having Children increases the chances the Parents are Married. The Children in a Family may or may not be related to both Adults.
2. An Adult is a Male or Female Person whose age is 18 years or older. Some family members do not marry one another.
3. If an Adult has a Child, then they are called a Parent. A Child usually lives with one or two Adults. The concept of having a Child may be as an Offspring or having Legal Responsibility for the Child.
4. A Child is an offspring from an Adult that is always less in age than their Parent. Over time, Children become Adults.
5. Being Married is a legal contract to be a Family. Children typically do not marry one another.
6. A Pet is a Dog that is owned by an Adult or Child.
7. A Family Setting might be when two or more members of a Family are within close physical proximity of one another.

Step 2 - Educate the Reasoner

## A HYBRID APPROACH IS BETTER



**Results from Deep Learning**

1. Adult female
2. Adult male
3. Young male
4. Young female
5. Dog
6. No background image
7. Overlapping objects

Deep Learning labels objects to produce relevant information (features) along with confidence values and geospatial placements

+

**Results from Knowledge-based Reasoning**

1. Maybe the adult female is a mother
2. Maybe the adult male is a father
3. Maybe the young female is their child
4. Maybe the young male is their child
5. The dog is maybe their family pet
6. This could be a family setting
7. This could suggest married with children

Knowledge-based reasoning uses features, combines them with context and reasons with domain expertise to generate answers

=

**Advantages**

- + A system that is both trained from data and educated using knowledge provides a quantum leap in problem solving capability
- + A hybrid approach improves the value of ML by adding a reasoning layer that reduces the amount of data required to solve problems
- + The hybrid approach coupled with a cognitive audit trail provides transparency of results to promote acceptance and enables deployment to high-value assets

# STEPS TO DELIVER A SOLUTION TO A CLIENT




Figure 7. Steps to build a Cognitive AI client solution

## COGNITIVE SYSTEMS THINK FOR THEMSELVES

Our LUMINAI™ cognitive reasoning system represents wisdom in a variety of forms, such as training from example data and education from domain knowledge and experience. We educate them using a rich variety of different forms of knowledge, and they deduce their own steps to solve the problem using heuristic approaches, where their knowledge is represented as defeasible tidbits of wisdom, such as facts, rules, algorithms, strategies, heuristics, prior examples, analogies, and workflows. These systems are fundamentally different from expert systems because they do think, which means they can solve general classes of problems not explicitly programmed, just like people do.




Figure 8. AI technologies used by LUMINAI™ Cognitive Reasoning Engine

## HUMAN-LIKE REASONING

The combination of symbolic education and data-centric training reduces the number of ground instances that are necessary to have a high degree of recognition. We believe this approach approximates the observed reasoning processes of how a person would solve a problem in unfamiliar or unknown territory. But providing a descriptive audit trail of how the answer was solved opens up a whole new class of opportunities for AI being applied to high-value assets (e.g., energy, healthcare, fintech).

The education of LUMINAI™ is completely different from the old school approach of using rule-based inferences, which have clearly been shown to be useful in solving the simplest kinds of problems. For example, consider the problem of isolating partially occluded people in a complex scene. Rather than training a neural net for every possible configuration of a human, you first educate a symbolic reasoner on what constitutes a human (e.g., head, torso, arms, legs) and augment it with the geospatial constraints of the elements and then have the reasoner compose the neural nets to efficiently recognize the sufficient number of parts. This synergistic combination provides a hybrid approach of fusing grounded a priori knowledge coupled with pattern matching from neural nets, producing not only an answer but the additional benefit of an audit trail for how it got the answer.

The symbolic approach says that the best way to teach an AI is to feed it human-readable information related to what you think it needs to know. If you want to create an AI to replace a doctor, you feed it a ton of medical textbooks and it answers questions by looking up the answers from those textbooks.

The non-symbolic approach admits that human-based information formats are not always the best fit for AI and encourages feeding raw information into the AI that it can analyze and construct its own implicit knowledge about.




Figure 9. How knowledge is encoded in LUMINAI™

### Knowledge is represented in one of two possible forms:

- + We either work with the client's domain experts to directly encode their knowledge in Homer or to implement a Domain Specific Language (DSL) so they can encode their own knowledge directly
  - + The internal language of LUMINAI™ is called Homer, which is compiled to native machine code for maximal performance
  - + A client-defined Domain Specific Language (DSL) is automatically translated by LUMINAI™ to Homer
  - + Knowledge is represented using facts, rules, algorithms, experience, heuristics, analogies, examples, and workflows; LUMINAI™ automatically extracts those representations from the DSL and encodes them into Homer
  - + Audit trails are handled in the same way. Either they are displayed as Homer or in terms of the client's DSL

# HOW IT WORKS – THE HYBRID APPROACH

LUMINA<sup>TM</sup> is based on symbolic techniques adapted from case-based approximate reasoning systems, Assumption-based Truth Maintenance Systems (ATMS), and numeric methods adapted from the Dempster/Shافر theory of evidence, as extended in Baldwin's Support Logic Programming system. The hybridization is achieved by a proprietary symbolic algebra system for uncertainty calculations.

This technique has several major advantages over conventional methods for performing inference with numeric certainty estimates in addition to the ability to dynamically determine hypothesis spaces. This includes improved management of dependent and partially independent evidence, faster run-time evaluation of propositional certainties, the ability to query the certainty value of a proposition from multiple perspectives, and the ability to incrementally extend or revise domain models.

LUMINA<sup>TM</sup> provides a closed-loop solution that can simultaneously reason and learn, so it does not need to be taken off-line for it to be updated with new symbolic knowledge. Because of this, it can take advice from both people and other AI systems so its results can be critiqued, extended, changed, or tuned in real-time to provide the autonomous capabilities of learning and optimization. At the highest level, its approach is a classic and familiar one, but it is also the one most frequently employed by people: (1) Acquire data, (2) Adjust the data to be useful, (3) Propose a solution, (4) Use it if it is good enough, (5) If wrong, either fix it internally or seek advice and try again (see Figure 5).




Figure 10. How LUMINA<sup>TM</sup> works

Here, we give the system multiple models of “rational behavior” while simultaneously endowing the system with the means of using and adapting over time the “best” of these presumably rationale models. In this way, the LUMINAI™ version of reasoning and learning is a particular brand of “bio-inspired rationality”. In this world, LUMINAI™ logic takes “economic” concerns into the fold. LUMINAI™ is instructed to use an initial set of models, propose hypothetical extensions, use tools to measure the results of experiments, and integrate the new information as evidence for or against the validity of the models available. LUMINAI™ is always trying to learn, but only taking advantage of the “useful” parts of its learning: those parts that allow for relatively better performance in a system-wide sense. This allows LUMINAI™ to be smart but tends to ensure there is no blind optimization of potentially dangerous goals.

Put more specifically, LUMINAI™ runs on a meta-model of reasoning and learning that enables it to encode several potentially contradictory or incomplete models for how to do different kinds of work. By applying these and using them “in-situ,” LUMINAI™ can interpret the results of sensors and machine learning-based perceptual mechanisms to encode new knowledge about the models as grounded in the environment. Over time, the experience accumulated through this reasoning, testing, and learning process allows LUMINAI™ to make decisions about which models are worthwhile to continue learning about and how to balance this continued pursuit of the learning activity with other survival or operational goals.



*Is Your AI Smart Enough To Reason Like an Expert?*


## MAIN BENEFITS OF BEYOND LIMITS COGNITIVE AI SYSTEMS

- + Augment human decision-making with better information
- + Accelerate analysis and planning
- + Fast-track time-to-market
- + Reduce waste
- + Elevate Formulator/Chemist support and performance
- + Improve productivity
- + Reduce tedious, dangerous, or repetitive work

## HYPOTHETICAL PATHS TO PREDICT THE FUTURE

The key to cognitive intelligence in AI is to begin with an initial set of models, then propose hypothetical extensions to the models. For example, imagine if your model has five parameters defined to frame a problem and lead to a solution, but you only have data for three out of the five. You need to make a decision, but you don't want to guess. A Cognitive AI system looks for guidance from encoded human expert knowledge and experience, from observations of available sensor data and from external data sources. It can then model hypothetical paths, predict future events, and propose courses of action to make smart decisions and resolve the issue, even with less-than-ideal data.

In remote, disconnected environments, like a Rover on the surface of Mars, an oil rig in the frozen Arctic or a sweltering desert, conditions are never ideal. Data is frequently unavailable. Human experts are not always on-site or able to intervene. But an AI system that is cognitive, always-on, and always paying close attention, can step in. Advanced Cognitive AI systems can capture information from IoT sensors, outside data sources, historical experience, institutional knowledge, plus expert humans, and then reason solutions in real-time, anticipate problems and resolve them in record time; even before they happen.



## WHAT IS EXPLAINABLE AI?

Commonly used conventional AI techniques like machine learning, deep learning, and neural networks are *black box* approaches that have no explainable understanding of how they arrive at answers. Advanced solutions, like those provided by Beyond Limits' Cognitive AI engines, deliver clear audit trails that explain the reasoning behind their recommendations showing all evidence, risks, certainties, and ambiguities. These transparent audit trails are designed to be understood by people and interpreted by machines. Explainable AI (XAI) is the key to building trust in artificial intelligence solutions with humans in the loop.

## DECISION SUPPORT – AI WORKING FOR HUMANS

The best Cognitive AI systems augment human decision-making and amplify human talent. They provide lightning-fast always-on monitoring, assessment, and autonomous problem-solving in environments often hostile to people. They never lose focus, never get tired, and never get distracted. When built to solve problems for industrial companies, Cognitive AI systems free people from dangerous or repetitive tasks so they can focus on higher-value work. In energy production, they improve efficiency and reduce waste. In power generation, they predict demand to calibrate supply.

With super-human processing power informed by big data and human expertise, Beyond Limits Cognitive AI systems represent new ways of thinking and working with new insights into how to go beyond conventional AI.

## AUTOMATION ON THE PATH TO AUTONOMY

The key to implementing this technology in industrial and IoT settings revolves around the degree to which an AI system is entrusted with autonomous decision-making. Using thermostats as an example, here is a simplified evolution of control systems from manual control to automation to autonomous decision-making and operation.




Figure 11. Example of evolution from manual control to cognitive control

## INDUSTRIES TAKING ADVANTAGE OF BEYOND LIMITS COGNITIVE AI

Primary users of Beyond Limits Cognitive AI systems include heavy industries: Energy, Power & Natural Resources, and Manufacturing & IoT. Other industries include Healthcare, Automotive, Finance, and Logistics. The world's most demanding industries are appreciating the need to go beyond conventional AI and use Cognitive AI solutions to solve some of their most complex problems. They are using Beyond Limits solutions to accelerate analysis and planning, boost efficiency, fast-track time-to-market, reduce waste, elevate performance and operator support, improve operating conditions, and realize unprecedented operational visibility and system planning agility.

For most professionals, the problem doesn't center around too little or too much information. The problem is more about getting the right information to the right people at the right time to make better decisions. An AI system powered by Beyond Limits cognitive reasoning thinks through the situation on behalf of the person and provides recommended courses of action in a clear, evidence-based audit trail that can be queried for answers. In other words, AI that can explain how it thinks so people can make vital decisions based on a clear understanding of risks and probabilities. The transformative speed of AI analysis accelerates time-to-decision so people can dramatically increase their capacity.

## APPLIED COGNITIVE AI – ROI IN THE INDUSTRIAL WORLD

Nothing is more fundamental to modern life than the energy it takes to power homes, businesses, transportation, food industries, and more. The goal is to find better ways to produce energy while emitting less carbon. To this end, Beyond Limits deploys Cognitive AI products and systems across the energy value chain to reduce waste, increase efficiency, and drive profitability at sustainable production levels.

- + **Subsurface Systems:** interprets vast data sets to understand complex geological structures, fluid dynamics, and seismic activities thousands of feet below the surface. The systems are designed to precisely identify where to explore next, advising reservoir engineers on optimized well locations and predicting in-fill production rates to accelerate months of manual analytic work and reduce unproductive wells, potentially saving millions.
- + **Production Systems:** designed to keep wells flowing and healthy by analyzing data in the light of decades of encoded human knowledge for more efficient 24/7 facility operations, no matter how remote.
- + **Refinery Systems:** monitors an entire refinery, analyzing performance to identify problems, predict remedies, and operate with less waste. The system comprehends the entire big picture operation down to device-level details and understands the relationships between subsystem units.
- + **Formulation Systems:** draws on decades of human expertise to advise lubricant engineers to meet stringent standards and testing protocols by evaluating thousands of options in seconds to recommend the optimal solution, saving significant time and money.
- + **Inspection Systems:** Cognitive AI embedded into robots to power autonomous operations, even in dangerous conditions, so large industrial facilities can operate safely and predictably around-the-clock.
- + **Power Systems:** helps manage operations of new LNG power plants. The world's first power plant guided by Cognitive AI technology, designed to balance supply/demand and improve reliability for economic and industrial development.

## SEE THE BIG PICTURE AND THE SMALLEST DETAIL

Beyond Limits has deployed LUMINAI™-powered systems to monitor entire refinery facilities, analyzing performance to identify problems, predict remedies, and operate with less waste. The system comprehends the entire big picture operation down to device-level details and understands the relationships between subsystem units. This level of automated data collection and analysis benefits operators trying to attain daily planning goals.


For example, the operator is given the directive that the refinery plan calls for 90% propane production from a specific distillation column, but the AI system notices that only 75% is being produced. That's a problem because they're losing money or reducing profitability. So the AI system would say: "this specific sensor is detecting that you're not producing as much propane as you should, so to increase production efficiency, the AI system recommends that you increase the reflex ratio in that particular distillation column." That's one case.

But what if the column load is already too high, that system is constrained, and the mitigation step does not apply? Then the system would say: "Your propane level is too low. To increase the

propane level, you need to increase the bottom of a distillation column upstream so that remedy will work its way downstream and eventually increase the amount of propane.” An operator who’s been working at the plant for 30 years might look at the problem and say: “Well, the first thing I always do is increase the reflux ratio.” The AI system will give an explanation: “I looked into that - but it’s already fully constrained, so you can’t use that mitigation. Try mitigation step #2.” So, in this example, the AI solution provides clear recommendations and a clear explanation of its reasoning to inform the operator.



Figure 12. Beyond Limits Cognitive AI understands the smallest details and the big picture, from individual sensors to the entire production facility.



## IoT AND BEYOND

For us, a smart IoT device isn’t just a tiny sensor controlling the lights in an office. It could be a refinery the size of 10 football fields, a supertanker in the middle of the Indian Ocean, an oil rig in the North Sea, a robotic crane the size of a skyscraper moving containers in a port, or an autonomous Rover operating on the surface of Mars, 150 million miles from Earth. At that distance, real-time radio control from the earth is not possible. There is no cloud. Sensor data can’t always be trusted. Systems fail in extreme conditions, yet the mission must not fail.

Our embedded solutions can be deployed at the edge for critical real-time applications and handle limited, missing, or messy data with ease. Above all, Beyond Limits Cognitive AI solutions feature explainability, providing transparency to insights, clear evidence, and a query-able audit trail for informed decision-making, even at the extreme edge.

## About Beyond Limits

Beyond Limits is a pioneering Artificial Intelligence engineering company creating advanced software systems that go beyond conventional AI. Founded in 2014, Beyond Limits is helping companies solve tough, complex, mission-critical problems and transform their business. The company applies a unique hybrid approach to AI, combining numeric AI techniques like machine learning with higher order symbolic AI and expert human knowledge to produce actionable information. The result is faster, better decisions that reduce risk, decrease waste, and increase efficiency.

## Beyond Limits' Relationship with the Caltech Jet Propulsion Laboratory

Beyond Limits' technology evolved out of ambitious NASA projects created by Caltech's Jet Propulsion Laboratory and the legendary AI Reasoning Laboratory. A vital reason the 2012 Curiosity Rover's mission to Mars was successful was due to the intelligent Cognitive AI technology that helped keep it safe when a potentially mission-critical scenario arose along its journey. In 2014 Beyond Limits formed (with Caltech's Jet Propulsion Laboratory as a founding investor) to commercialize powerful technologies used for deep space missions, brought down to Earth to solve some of this world's toughest challenges.

For more information about Beyond Limits Cognitive AI, visit [www.beyond.ai](http://www.beyond.ai)

©2020 Beyond Limits | All Rights Reserved