

Enhancing Enterprise Adoption of Large Language Models (LLMs) through Symbolic AI: An Auditable and Explainable Approach

Rainbird: A Symbolic AI Platform to Scale LLMs Safely

Table of Contents

Table of Contents	2
Executive Summary	3
Introduction	4
Background on Large Language Models (LLMs)	4
The need for safe and responsible use of LLMs in enterprise applications	4
The role of symbolic AI in addressing LLM limitations	4
Overview of Symbolic AI	5
Definition and key concepts	5
Comparison of symbolic AI with connectionist AI (e.g. GPT-4)	5
Benefits of a combined approach	6
Rainbird AI: A Symbolic AI Platform for LLM Integration	7
Overview of Rainbird AI	7
Key features and benefits for enterprises	7
Case studies of successful Rainbird AI implementations	8
HFS Research: Rainbird makes complex privacy risk assessment a seamless experience (EY)	8
Gartner: Augment domain-expert decisions with knowledge graphs (BDO UK)	9
Combining Rainbird's Symbolic AI and LLMs for Enterprise Applications	10
The complementary nature of symbolic AI and LLMs	10
Potential use cases and industries	11
Challenges and Future Directions	13
Limitations of the current state	13
Ethical considerations and responsible AI practices	13
Future trends and research directions in symbolic AI and LLMs	14
Recommendations for Enterprises	15
Assessing the need for explainable and auditable AI solutions	15
Evaluating potential symbolic AI platforms, such as Rainbird.	15
Conclusion	16
Appendices	17
Glossary of terms	17
Acknowledgments	18

Executive Summary

The purpose of this white paper is to explore and advocate for the integration of symbolic AI, specifically Rainbird, as a means to enable enterprises to safely harness the power of large language models (LLMs) such as GPT-4.

By addressing the limitations of LLMs, particularly in terms of explainability and auditability, Rainbird provides a solution for organisations operating in regulated markets where decision-making transparency is crucial.

Symbolic AI is an approach that focuses on the manipulation of symbols and rules to represent knowledge and perform reasoning tasks. It relies on well-defined, explicit, and human-readable representations of knowledge. Rainbird is the latest generation of symbolic AI, designed to better address the challenges of modern AI applications. It combines the fundamental principles of symbolic AI, such as knowledge representation and rule-based reasoning, with a graph-based model and a powerful inference engine that allows for the integration of human priors and complex relationships.

As organisations increasingly rely on AI technologies for operational decision-making, incorporating symbolic AI into their AI strategy ensures that the reasoning behind decisions is transparent, traceable, and justifiable, especially in highly regulated industries where compliance is paramount.

By combining the strengths of LLMs with the explainable and auditable nature of Rainbird, enterprises can not only unlock the vast potential of LLMs but also maintain trust, accountability, and ethical standards, fostering responsible and sustainable AI adoption across diverse applications and sectors.

This white paper aims to outline the benefits of combining symbolic AI with LLMs, while also examining the unique capabilities of Rainbird and its application in various industries such as banking, insurance, financial services, tax, audit, advisory, and healthcare administration.

By presenting a comprehensive analysis of Rainbird's symbolic AI approach, we aspire to encourage enterprises to adopt a responsible and compliant AI strategy, ensuring the successful deployment of AI technologies for complex contextual decision-making processes. We highlight the critical role of symbolic AI in achieving these aims in regulated industries, its complementary nature with LLMs, and the potential use cases that can benefit from the combined approach.

It also underscores the importance of ethical considerations and responsible AI practices.

Introduction

Background on Large Language Models (LLMs)

Large Language Models (LLMs) represent a breakthrough in AI research, characterised by their ability to understand, generate, and manipulate human-like text. LLMs, such as GPT-4, rely on deep learning techniques and massive amounts of training data to achieve state-of-the-art performance in a wide range of natural language processing tasks. These models have demonstrated remarkable capabilities, including text summarization, translation, content generation, and question-answering.

However, despite their impressive performance, LLMs also present challenges for the enterprise. They are trained on closed data by third parties, suffer from biases, a lack of explainability, and difficulty in controlling their generated outputs. For this reason they are not well suited to making operational decisions where absolute confidence in the decision is important.

The need for safe and responsible use of LLMs in enterprise applications

As enterprises increasingly explore LLMs to enhance their business processes and decision-making, the need for safe and responsible use quickly becomes paramount. Ensuring ethical AI practices is crucial to maintaining trust, accountability, and compliance, especially in highly regulated industries. Responsible use of LLMs involves addressing potential biases, ensuring data privacy, and managing risks associated with a lack of transparency, unintended consequences or malicious use.

Only by focusing on these aspects can organisations harness the power of LLMs while adhering to industry regulations, ethical standards, and stakeholder expectations.

The role of symbolic AI in addressing LLM limitations

Symbolic AI plays a vital role in addressing some of the inherent limitations of LLMs, such as their lack of explainability and auditability. By representing knowledge in the form of symbols, rules, and relationships, symbolic AI systems provide transparent and traceable decision-making processes.

This is particularly important in regulated industries where compliance and trust are essential. Integrating symbolic AI with LLMs enables organisations to benefit from the advanced knowledge understanding capabilities of LLMs while maintaining the explainability and auditability required for responsible AI use.

As a result, enterprises can unlock the full potential of LLMs, ensuring that their AI-powered solutions remain ethical, transparent, and accountable.

Overview of Symbolic AI

Definition and key concepts

Traditional Symbolic AI often referred to as "Good Old-Fashioned Artificial Intelligence" (GOFAI), is an approach to AI that focuses on the manipulation of symbols and rules to represent knowledge and perform reasoning tasks. Symbolic AI relies on well-defined, explicit, and human-readable representations of problems and their solutions, typically employing techniques such as logic programming, semantic networks, and production systems. This approach dominated AI research during the mid-20th century, with a strong emphasis on top-down reasoning and the development of rule-based systems.

Rainbird is an evolution of traditional symbolic AI, designed to better address the challenges of modern AI applications. It combines the fundamental principles of symbolic AI, such as knowledge representation and rule-based reasoning, with a powerful inference engine and a graph-based model that allows for the integration of human priors and complex relationships. Unlike traditional symbolic AI solutions, Rainbird also handles probabilities, sparse or missing data and can operate at a massive scale, tasks not possible in the early 1990s.

Comparison of symbolic AI with connectionist AI (e.g. GPT-4)

Symbolic AI and connectionist AI, also known as neural network-based AI, represent two distinct paradigms in artificial intelligence. Both approaches have their strengths and weaknesses, making them suitable for different types of tasks and applications.

Symbolic AI (e.g. Rainbird):

- Knowledge representation: Symbolic AI represents knowledge in the form of symbols, rules, and relationships, making it explicit and human-readable.
- Visibility: Models are human-readable and can be easily interpreted and tested.
- Reasoning: It uses logic-based reasoning and well-defined algorithms to process information and solve problems.
- Explainability: Symbolic AI systems are inherently explainable and transparent, allowing users to understand and trace the decision-making process.
- Usability: Symbolic AI systems like Rainbird are no code and are easy to learn and use.
- Modularity: Symbolic AI systems are more modular, enabling easier modification and maintenance of individual components.
- Uncertainty and large data: While traditional approaches to symbolic AI struggled to handle uncertainty and large-scale data. Rainbird has evolved to handle these.
- Limitations: Symbolic AI cannot process unstructured data without additional natural language processing technology which LLMs like GPT-4 can now provide to symbolic AIs like Raibird.

Connectionist AI (e.g. GPT-4):

- Knowledge representation: Connectionist AI, often based on neural networks, is trained on closed data and represents knowledge implicitly through the learned weights and connections between artificial neurons.
- Visibility: Connectionist models are not human-readable.
- Ownership: Models are owned by their creators who invest exceptionally large sums of money building them. The models are not in the user's control.
- Reasoning: Connectionist AI processes information and solves problems through parallel computation and iterative learning from data.
- Explainability: Connectionist AI systems, particularly deep learning models, are often considered "black boxes" due to their lack of transparency and explainability.
- Adaptability: Neural networks can automatically learn and adapt to new data and situations, making them suitable for tasks that involve pattern recognition, natural language processing, and image classification.
- Limitations: This approach can require large amounts of data and computational power for training, and its lack of explainability may hinder its adoption in regulated industries or situations where accountability is crucial.

Benefits of a combined approach

Symbolic AI excels in tasks that demand explicit knowledge representation, logical reasoning, and explainability, while connectionist AI is better suited for tasks involving learning from data, pattern recognition, and adaptability.

The choice between these approaches depends on the specific requirements and constraints of the problem at hand with the need to explain outcomes being a key factor.

However, combining the strengths of both approaches, as demonstrated by the integration of Rainbird and GPT-4, can provide powerful, flexible, and explainable AI solutions for a wide range of enterprise applications.

Rainbird AI: A Symbolic AI Platform for LLM Integration

Overview of Rainbird AI

Rainbird is an advanced symbolic AI platform designed to address the unique challenges faced by modern enterprises, particularly in regulated industries. Combining the core principles of symbolic AI with innovative features, Rainbird enables organisations to build robust, explainable, and auditable AI models that facilitate complex decision-making processes.

These features include weighted conditions, optional or mandatory rules, and a range of expressions, providing flexibility and adaptability in modelling complex decision-making processes.

Rainbird facilitates the creation of AI models through a visual user interface. By seamlessly integrating with large language models such as GPT-4, Rainbird now also allows experts to develop AI models using natural language, providing a more accessible and intuitive approach to knowledge representation and rule-based reasoning.

It maintains the core advantages of traditional symbolic AI, such as explainability and auditability, while handling uncertainty and large datasets to better serve the needs of contemporary enterprise applications.

Key features and benefits for enterprises

- **Explainability:** Rainbird's symbolic AI approach ensures that the decision-making process is transparent and easily understandable. This is particularly important in regulated industries where accountability, compliance, and trust are paramount.
- **Auditability:** Rainbird's AI models are inherently auditable, allowing organisations to trace the chain of reasoning behind any decision. This helps identify potential issues or improvements in the decision-making process and then facilitates regulatory compliance.
- **Flexibility:** Rainbird's graph-based model and rule-based reasoning enable the platform to handle complex relationships and human priors, allowing for the development of AI models that can adapt to diverse and dynamic business environments.
- **Integration:** Rainbird's compatibility with large language models like GPT-4 enables the development of AI models using natural language, making it more accessible for domain experts to contribute their knowledge and expertise.

- **Visual User Interface:** Rainbird offers a user-friendly visual interface for building AI models, simplifying the process of knowledge representation and rule creation, and making it more approachable for non-technical users.
- **Scalability:** The platform is designed to handle large-scale applications, providing enterprises with the ability to scale their AI models as needed to accommodate growth and changing business requirements.

By offering a powerful combination of explainability, auditability, flexibility, and integration with LLMs, Rainbird AI empowers enterprises to harness the power of advanced AI technologies while maintaining accountability, trust, and compliance in their decision-making processes.

Case studies of successful Rainbird AI implementations

HFS Research: Rainbird makes complex privacy risk assessment a seamless experience (EY)

HFS Research recognised Rainbird as a OneOffice Hot Vendor for its alignment with the HFS Triple-A Trifecta, utilising automation, analytics, and AI to deliver value. Rainbird's capabilities and usability outperformed previous decision engines for many customers, making it an attractive choice for automating complex decision-making at scale.

A notable case study involves EY's development of the client-facing SaaS platform, EY Data Permissions Navigator (EY DPN), which manages privacy risks. Rainbird played a critical role in the platform, injecting decision intelligence directly into the privacy risk assessment process. EY DPN leverages Rainbird's advanced non-linear decision engine to automate privacy risk assessments, capturing expertise from data privacy specialists and replicating it at scale.

Rainbird's usability and the responsiveness of its engineering and support teams have been praised, contributing to a reduction of roughly 12 months in the platform's build time. The EY DPN now completes certain decision-making processes in around 15 minutes, which previously took several months. Rainbird's scalability and usability make it a powerful platform for automating complex decision-making processes, with EY's use case demonstrating its ability to handle complexity at scale.

“Complex processes need scalable decision engines that can easily cut through complexity, and Rainbird delivers on its promise.”

Ralph Aboujaoude Diaz, Practice Leader HFS Research



Gartner: Augment domain-expert decisions with knowledge graphs (BDO UK)

BDO UK aimed to augment and scale expert decision-making, which is complex, nonlinear, and requires deep domain knowledge.

The company focused on enhancing transparency, consistency, and auditability in decision-making, using Rainbird to translate the problem of determining optimal R&D tax credits into symbolic AI.

By collaborating with tax experts and data and analytics teams, BDO UK created a knowledge graph for R&D tax credits, incorporating nodes, relationships, data, rules, and weights, with Rainbird's inference engine determining the optimal path for a company's R&D tax claim.

This approach significantly reduced the time required for tax credit estimates, from nearly five hours to just a few seconds, allowing tax experts to focus on higher-value work.

The R&D tax credit knowledge graph was developed into a digital product, attracting 30 new clients and generating an 80% increase in the pipeline of advisory work.

“We turned to automation not only to productize expert decisions, but to also engage and retain our talent. Reducing the time spent on R&D tax benchmarking from five hours to a few seconds gives our experts more time for interesting, high-value work and better work-life balance.”

Dan Francis, Director of Innovation and Digital at BDO UK

Combining Rainbird's Symbolic AI and LLMs for Enterprise Applications

The complementary nature of symbolic AI and LLMs

Symbolic AI and LLMs like GPT-4 offer complementary strengths that can be harnessed together to create powerful and effective AI solutions. While LLMs excel at natural language understanding, generation, and complex pattern recognition, Rainbird provides a knowledge representation platform and an inference engine that can handle complex, nonlinear decision-making.

Rainbird enables the creation of knowledge graphs with explicit rules, conditions, and weights, making it easier to trace the reasoning behind a decision. This level of explainability is essential in regulated industries like finance, healthcare, and insurance, where understanding the rationale behind decisions is crucial for compliance, risk management, and trust-building.

LLMs like GPT-4 can quickly process vast amounts of textual data and generate Rainbird models by converting expert knowledge, captured through natural language interactions, into Rainbird's underlying knowledge representation language, RBLang. This accelerates the development of knowledge maps and makes it easier for domain experts to contribute to AI models without requiring specialised programming skills.

By combining the capabilities of LLMs and symbolic AI systems like Rainbird, enterprises can create AI solutions that are both powerful and transparent. LLMs can be used to gather, preprocess, and generate insights from unstructured data, while Rainbird can be employed to make complex, auditable decisions based on the insights derived from LLMs.

This complementary approach not only enhances the effectiveness of AI solutions but also ensures that they are safe, responsible, and suitable for deployment in mission-critical and regulated environments.

Potential use cases and industries

The combination of GPT-4 and Rainbird offers a powerful and explainable AI solution that can be employed in various regulated industries. Below are just a few examples, but Rainbird provide a labs process to identify and validate new use cases.

1. Banking and Finance

- Credit risk assessment and scoring
- Anti-money laundering (AML) and fraud detection
- Regulatory compliance monitoring and reporting
- Loan and mortgage underwriting
- Personalised financial advice and planning

2. Insurance

- Underwriting and policy pricing
- Claims management and fraud detection
- Customer segmentation and tailored product offerings
- Regulatory compliance and reporting
- Risk assessment and mitigation

3. Accounting

- Automating tax calculations
- Assessing and validating intercompany transactions
- International tax planning
- R&D tax credit analysis
- Risk-based audit planning
- Fraud detection and prevention
- Internal control evaluation
- Audit evidence review
- Financial due diligence
- Business valuation
- Corporate governance and risk management
- Financial planning and analysis

4. Healthcare

- Optimisation of healthcare processes and payments
- Clinical decision support and diagnostics
- Personalised treatment planning
- Patient risk stratification
- Regulatory compliance in medical device development
- Drug discovery and development compliance

5. Pharmaceuticals

- Clinical trial design and monitoring
- Regulatory submission and approval processes
- Pharmacovigilance and adverse event reporting
- Manufacturing quality control and compliance

6. Energy and Utilities

- Regulatory compliance monitoring and reporting
- Risk assessment and mitigation for infrastructure projects
- Environmental impact assessments
- Energy consumption forecasting and optimization
- Asset management and maintenance planning

7. Telecommunications

- Customer data privacy and security compliance
- Network infrastructure planning and regulatory compliance
- Spectrum allocation and management
- Fraud detection and prevention
- Regulatory reporting and audits

8. Transportation and Logistics

- Vehicle safety and emissions compliance
- Fleet maintenance planning and optimisation
- Route planning and optimisation for regulatory compliance
- Customs clearance and regulatory compliance
- Risk assessment for hazardous materials transportation

By integrating GPT-4 and Rainbird, organisations in regulated industries can benefit from advanced AI capabilities while ensuring explainability, transparency, and compliance with industry-specific regulations. This combined approach has the potential to enhance decision-making, optimise processes, and reduce risk across a wide range of applications.

Challenges and Future Directions

Limitations of the current state

While the integration of Rainbird and LLMs like GPT-4 offers significant power, there are some limitations to what is currently possible.

Incomplete knowledge is a concern, as GPT-4 may not always have up-to-date or comprehensive details for a specific domain or regulation, leading to gaps in the knowledge map that could impact accuracy and reliability. This can be mitigated by domain experts reviewing and adding to Rainbird models during the build process and ensuring that the knowledge required to automated decisions is properly represented.

The quality of natural language input is crucial, as LLMs rely on clear and unambiguous text to generate meaningful outputs. If the expert's explanation or instructions are unclear, GPT-4 may not accurately capture the intended meaning, leading to inaccuracies in the generated Rainbird model. Additionally, GPT-4 may struggle to interpret context accurately, especially in highly specialised or regulated domains, which could result in incorrect or irrelevant outputs when generating Rainbird models or when interacting with experts. This can be mitigated by using domain specific LLMs (for example Prometheus in healthcare), building a layer of "fine tuning" to an existing LLM and by domain experts carefully reviewing and testing the resulting Rainbird model.

Bias and ethical concerns are also significant, as LLMs may inadvertently introduce biases into the decision-making process based on the data they have been trained on or the expert input they receive. This may have unintended consequences, especially in regulated industries where ethical considerations and fairness are crucial. This is mitigated by Rainbird's transparency, testing capabilities and the ability to provide an audit trail for each decision.

Ethical considerations and responsible AI practices

When using a combined approach it is crucial to consider ethical considerations and responsible AI practices.

Ensuring that the decision-making process is transparent and explainable is crucial, particularly in regulated industries. Leveraging Rainbird's capabilities for explainability can help bridge the gap in LLMs' interpretability, making it possible for stakeholders to understand how a decision was reached.

LLMs can potentially introduce bias based on the data they are trained on, just as bias can come from the domain experts that build the models. To address this, it is essential to choose data carefully, involve domain experts in the development process and test models to ensure minimise bias in the generated models and decisions.

Other ethical considerations for the combined approach of symbolic AI and LLMs include data privacy, accountability, inclusivity, maintaining a human-in-the-loop approach, and regular monitoring and auditing.

Future trends and research directions in symbolic AI and LLMs

The future of Rainbird and LLMs will both be marked by continuous investment and advancements and innovative research. These developments will generate more and more effective AI systems that can better understand, reason, and interact with the world.

Rainbird is leveraging the reasoning capabilities of symbolic AI, along with the learning and generalisation power of LLMs. Such hybrid models will evolve to better capture the nuances of human cognition and facilitate more efficient, reliable, and explainable AI solutions.

This could lead to more effective decision-making and problem-solving in a wide range of applications, from natural language understanding and computer vision to robotics and complex optimization tasks.

Another research direction is the development of even more transparency and explainability. As the adoption of AI technologies grows, so does the demand for systems that provide clear explanations for their decisions, particularly in regulated industries. Researchers are exploring techniques to improve the interpretability of LLMs and better integrate them with the explainability features of symbolic AI systems like Rainbird.

Bias mitigation and fairness in AI models is another crucial area of research. As both symbolic AI and LLMs can potentially introduce bias into decision-making, future research will focus on developing strategies to detect and mitigate biases in AI systems. This may involve developing techniques to audit and refine models, and incorporate domain experts' input during the development process to minimise bias in generated models and decisions.

Finally, developing human-in-the-loop approaches that leverage the strengths of both humans and AI systems can lead to more accurate, reliable, and contextually appropriate decision-making. For example, Rainbird provides the ability for humans to automatically be called into a decision process to provide missing information.

Recommendations for Enterprises

Assessing the need for explainable and auditable AI solutions

Enterprises should begin by identifying the areas within their organisation where explainable AI can provide significant value. The approach is well suited to recurrent decision intensive processes where the efficiency, quality and auditability of the decision have value. Rainbird has an innovation lab and a number of thinking tools to help its client identify high value use cases.

Enterprises must evaluate the legal, ethical, and compliance implications of AI adoption and understand the importance of transparency and auditability in these decision-making processes.

Evaluating potential symbolic AI platforms, such as Rainbird.

When evaluating potential symbolic AI platforms like Rainbird AI, enterprises should consider the platform's capabilities, ease of integration, and scalability. Key factors to examine include the platform's ability to handle complex decision-making tasks, support for knowledge representation and inference, and the provision of explainable AI features.

Additionally, organisations should assess the platform's compatibility with their existing technology stack and the availability of resources, such as documentation and support, to facilitate successful implementation.

Conclusion

In conclusion, the integration of symbolic AI, such as Rainbird, with LLMs like GPT-4 offers significant opportunities for enterprises to harness the power of AI safely and responsibly.

By combining the strengths of both approaches, businesses can achieve explainability, auditability, and scalability, addressing limitations inherent in LLMs alone.

This white paper has highlighted the critical role of symbolic AI in regulated industries, its complementary nature with LLMs, and the potential use cases that can benefit from the combined approach.

It has also underscored the importance of ethical considerations, responsible AI practices, and the need for integrating these AI systems effectively within an organisation.

As enterprises continue to navigate the rapidly evolving AI landscape, adopting solutions like Rainbird can ensure they maximise the potential of LLMs while maintaining compliance and addressing ethical concerns.

By evaluating the suitability of symbolic AI platforms and fostering a culture of responsible AI use, organisations can make substantial strides in driving innovation, efficiency, and growth in an increasingly AI-driven world.

Appendices

Glossary of terms

Artificial Intelligence (AI): The development of computer systems that can perform tasks typically requiring human intelligence, such as visual perception, speech recognition, decision-making, and natural language understanding.

Symbolic AI: A subfield of AI that represents knowledge using symbols, such as logic, graphs, and rules.

Large Language Models (LLMs): AI models trained on vast amounts of text data to understand and generate human-like language. Examples include GPT-4 and its predecessors (GPT-3, GPT-2, etc.).

Rainbird AI: A symbolic AI platform that combines a knowledge representation system and an inference engine, allowing the creation of explainable and auditable AI models for complex decision-making.

Connectionist AI: A subfield of AI that focuses on artificial neural networks and deep learning, aiming to model information processing in a way that resembles the human brain.

Inference Engine: A component of a symbolic AI system that processes knowledge and rules to derive conclusions or make decisions based on the given data.

Knowledge Map: A graph-based knowledge representation that captures relationships between entities and concepts, allowing more efficient storage, querying, and reasoning.

RBLang: Rainbird's underlying knowledge representation language used to define the knowledge, rules, and relationships within a Rainbird model.

Explainable AI (XAI): AI systems that provide understandable and interpretable explanations for their decisions and outcomes, allowing for greater transparency and trust in their use.

Auditable AI: AI systems that can be inspected and evaluated for compliance with regulations, ethical guidelines, and organisational policies, ensuring responsible and safe use of AI technologies.

GDPR (General Data Protection Regulation): A comprehensive data privacy regulation enacted by the European Union, which sets rules for the collection, processing, storage, and sharing of personal data.

Human-in-the-Loop (HITL): An approach to AI system design that incorporates human expertise and judgement alongside AI components, ensuring critical decisions are not solely left to the AI and addressing potential limitations in the system.

Ontology: A formal representation of knowledge within a specific domain, consisting of concepts, entities, and the relationships between them, often used in symbolic AI systems.

Inclusivity: The practice of designing AI systems and applications that cater to the needs of diverse users and do not exclude any specific groups.

Acknowledgments

This paper was written by James Duez and Ben Taylor with drafting support from GPT-4.