VIEW POINT

THE HYBRID AI OPPORTUNITY FOR BUSINESSES

Abstract

Hybrid or Neuro-Symbolic AI is a sub-field of Artificial Intelligence that combines classical Symbolic AI techniques with Neural Networks. It has the potential to significantly boost democratization of AI, help Enterprises scale their digitalization initiatives and achieve business differentiation. Hybrid AI provides an excellent platform and construct for Responsible Human-in-the-Loop (HITL) AI implementations that inherently support interpretation and scrutiny by Business-IT organizations as well as compliance and regulatory authorities.

Infosys®
Navigate your next

“We can only see a short distance ahead, but we can see plenty there that needs to be done.” - A. Turing

AI is at its second inflection point and is ready for mainstream adoption at the scale needed for an enterprise. It is expected to create significant new capabilities, bring in efficiencies and extract the most out of technology investments in Cloud, Big Data, IoT, Robotics, etc. This acceleration in adoption is attributed to several converging factors such as technology affordability, democratization, and extreme digitalization necessitated by several disruptions of global ramifications. There is a significant increase in the footprint of AI across the systems as enterprise use cases continue to grow. AI will prove to be a significant value lever for an enterprise on its digital transformation journey by transforming back office systems and preparing them for the new and personalized consumer experiences. Infosys recently unveiled its Applied AI offering that comprehensively covers the discovery of the right use cases, rapid experimental verification via the AI Living Labs, and de-risked enterprise scale implementations to realize the true benefits of AI.

The need for Hybrid or Neuro-Symbolic AI

Humans understand and respond to the real world by processing the world as symbols while technology predominantly works on numbers & data. In broad terms, let’s consider how intelligence develops in a child – it relies on sensory capabilities and neurons to understand the real world and communicate to the brain to correctly interpret and respond to real-world stimuli. It naturally learns by experience, reinforces, abstracts to a higher order of thinking and transfers this for use in similar or often different situations. Abilities that are natural to a child are however difficult to realize in practice using technology. But the quest remains strong for generalizing AI and reaching levels of or exceeding human capability with attempts combining various sub-fields of AI. Capabilities such as Natural Language Processing, Speech, Image Recognition & Optical Character Recognition help in understanding and translating the real world to computer readable formats. Experts define the ontology and rules that work with higher order concepts encapsulated and manipulated as symbols that reflect various levels of abstraction and context with respect to the real world. Symbolic AI covers the generalization of working with symbols and the corresponding rules while Neural Networks are the data driven counterparts. These complementary techniques form the field of AI known as Neuro-Symbolic or Hybrid AI. There are encouraging attempts at human like ability to understand and interact with the real world as demonstrated by the Neuro Symbolic Concept Learner at MIT-IBM Watson AI Lab.

Let’s now consider other real-world implementation of AI at varied levels of maturity. Autonomous Vehicles sense and respond in real-time to the world around it.

On-board sensors continuously gather real-time information about the surroundings and then models use the data to continuously interpret and command the vehicle to respond appropriately. To achieve this level of autonomy, AI techniques such as Deep Neural Networks are the best fit for interpreting the stimuli e.g. real-world conditions such as an identifying lighting conditions, the presence of an obstacle, how far the obstacle is, etc. Symbolic AI then enables data-driven decision-making on slowing down, the rate of slow down, to what speed, when to stop and for how long, whether the vehicle needs to move around the obstacle from the left or the right, etc. The Symbolic AI is typically encoded in Expert Systems with ontologies, decision trees and rules that define the Autonomous Vehicle's operating rules.

In the Healthcare context, Virtual Health Assistants need to understand utterances that are in the medical lexicon, infer the intent and respond based on medically relevant information on symptoms, health records, diagnoses, etc. To achieve this, a set of ontologies relevant for the Healthcare context need to be established while AI/ML/Deep Learning can provide the right interpretation of the natural interactions via voice, text, and video. In this scenario as well, we see a need for multiple AI techniques to work together to achieve the overarching goals.

In the banking scenario today, customer onboarding occurs within minutes if not seconds while it used to take days to complete such processes not very long ago. Exciting AI technology is already at work albeit in carefully designed pockets of implementation for e.g. Chatbots for customer assistance and self-service, workflows such as Know your Customer (KYC) and Fraud Prevention that are partially accelerated using Computer Vision algorithms for Facial Recognition and Optical Character Recognition (OCR) on Identity proofs, etc. There is a potential for more e.g. real-time Credit Risk Assessments that need extensive insight from in-house knowledge repositories as well as from third-party databases. The chatbots would need to be primed with Financial Services related terminology, the rules of response need to be configured in the system to ensure a structured interaction via the chatbot channel. In this highly regulated industry, there is little leeway to allow AI to devise its own new responses in an unstructured manner. Contrast this with scenarios where Alexa would respond with a joke or could even give a wrong answer on a consumer's home echo device without serious consequences.

As seen from the varied use cases, we opine that a combination of Symbolic and Sub-symbolic approaches needs to come together to realize real-world use cases. AI needs to be maintainable, interpretable, and explainable. Business should know how confident the AI is on the decision and it should provide evidence of why, for example, a Credit Risk Algorithm accepted or rejected an application. AI needs to be ethical as well i.e. free of bias and discrimination. Any possibility of bias being introduced by the programmer or the data needs to be identified and addressed immediately using Human-in-the-Loop (HITL) techniques across the life cycle of AIOps and execution of the AI in live production scenarios. The Hybrid AI approach provides multiple benefits including enhancing outcomes of existing Neural Networks as seen from illustrations in the below table.

Business Objectives	Intelligence Techniques Employed to Achieve the Goals
Recommendations for "Complete the Look" in Online Shopping In an online apparel shopping scenario, an intelligent system mimics a sales representative who is always looking for opportunities to cross-sell. The engine visually analyzes the consumer's choices and interests to suggest matching garments and accessories.	A hybrid approach here combines an expert system to encode SKU catalog and relations, color matching rules and ranking. Run time Computer Vision & DL algorithms (Image processing / Deep Learning) extract features from the selection catalog or an image provided by the consumer. The algorithms then adaptively find matching colors based on nearest neighbor principles and provide recommendations as close to what seasoned sales representatives provide.
Multi-Factor Adaptive Authentication for Customer Identification In the Digital world, customer identification is crucial to ensure products and services are reaching the intended recipients. Business and end user confidence in correct identification is an important foundational capability for compliance reasons as well as to tackle increasing incidents of cyber malpractice.	Continuous physiological and behavioral biometrics identification based on typing patterns, mouse movement, gaze, gait, face recognition, etc. were used to monitor activity in real time. The encoded rules dynamically decide on offering the right combination and sequence of authentication challenges based on the real-world conditions and the thresholds defined by compliance teams. Deep Learning models were executed real time to derive physical & behavioral patterns. Overall accuracy levels increased from 75-85% to consistently around 85-95%.
Synthetic Transactional Data Generation for training Deep Learning Models (Financial use cases) Large volumes of data are needed to effectively train and improve Deep Learning models. The objective here is to create data that mimics real world transactions at volumes needed for Deep Learning models.	Generative Adversarial Networks (GAN) were initially used to create synthetic data after training on a small sample set. The results were of reasonable quality and required manual cleanup. Hybrid AI techniques combined the GAN models with expert defined data models and rules to improve the quality of data and accuracy levels in generation resulting in very high quality of synthetic data that was a close match to real world transactional records improving the overall quality of training and testing.

Similar projects provided several encouraging insights including the ability to build a team and rapidly scale up on AI technologies, design the Hybrid AI solutions and achieve accuracies that are closer or better than human success rates. The symbolic vs. sub-symbolic footprint in an AI solution will of course vary depending on the nature of the use case. Simple custom programming approaches may suffice in leveraging Neural Network models to achieve specific & narrow objectives. More complex implementations will leverage Knowledge graphs, Rule engines as well as custom programming

with high level languages. The selection of programming languages could also vary significantly ranging from popular high-level languages and functional approaches to those that inherently support parallel execution. General Purpose Graphics Processing Units (GP-GPU) will be commonplace infrastructure for speeding up processing to near real-time. Languages that had niche roles earlier in scientific computing domains may also get to see mainstream adoption.

Evolution of Hybrid AI

The Hybrid AI thinking dates to 1969 when Marvin Minsky and Seymour Papert in their book *Perceptrons* argued that a connectionist neural network-based approach would never be sufficient to imbue machines with genuine intelligence.

Early implementations of AI, also referred to as **Symbolic (or Classical) AI**, used expert-built ontologies and knowledge representations, decision logic, rules and later extended to statistical models and machine learning (Classical ML) based methods. These provided an almost magical differentiation to the pioneers who embraced these technologies. When encoded into Expert systems or Decision support systems, these AI implementations automated tasks or assisted novice users in mimicking experts. Symbolic approaches were easier to understand, engineer and lent themselves well to regulatory and legal scrutiny. However, these approaches faced a challenge in dealing with the variety and scale of unstructured content.

Neural Networks (NN) & Deep Learning (DL) are best suited for working with unstructured data. Also known as sub-symbolic AI, this is proving itself in the real world where information is unstructured and where the goal could also be discovery of hitherto unknown correlations in vast amounts of Text, Documents, Audio and Video. However, DL provides a narrow field of intelligence in very specific areas to identify patterns and classify. It also needs substantial investments in infrastructure and annotated data to train and execute. There's also a degree of uncertainty in the outcomes with accuracy levels varying based on conditions in which the DL models execute. There's a concern on Explainability and the introduction of bias in data or the models, etc. which has sparked extensive research on Explainability and Ethical AI to reduce the gaps perceived by businesses and regulatory bodies.

Hybrid or Neuro-Symbolic AI aims at combining the two complementary techniques of Symbolic and Sub-Symbolic AI bringing forth a greater potential while minimizing the gaps of each individual technique. This would allow businesses to achieve their AI goals while retaining the ability to withstand compliance related scrutiny. This is largely due to the interpretability offered by the symbolic approach which provides a governance layer over and above the relatively less interpretable and oftentimes unpredictable or fuzzy outcomes of the Neural Network models.

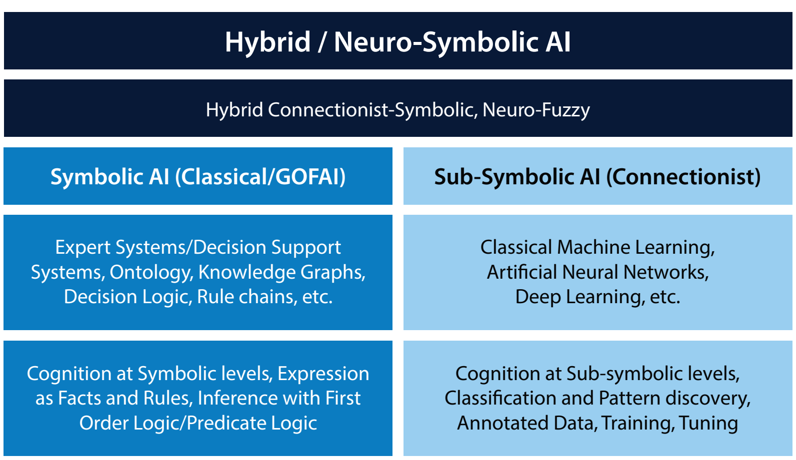

Figure: Neuro-Symbolic AI

A Hybrid AI approach will ensure responsible implementations and continued differentiation for businesses while the general trend of commoditization of AI will help adoption and scale. The added advantage of the 'Hybrid AI' thought process is its potential to bring in a much needed 'method to the madness' in AI initiatives.

Authors

Prasad Joshi

Senior Vice President and Head,
Infosys Center for Emerging Technology Solutions (iCETS)

Prasad Joshi is a Senior Vice President at Infosys and heads the Infosys Center for Emerging Technology Solutions (iCETS). With more than 20 years of experience in software technology, IT, business process outsourcing, and multicultural team management, he defines the strategic program, sets up co-creation labs with key clients, and leads the innovation ecosystem at Infosys. Prasad is on the customer advisory board for Microsoft Azure and is also an invited member for the industry advisory board at the Georgia Tech Information Security Center.

Sanjay Mohan

Senior Principal Technology Architect,
Infosys Center for Emerging Technology Solutions (iCETS)

Sanjay Mohan is a senior principal technology architect at Infosys' Center for Emerging Technology Solutions. He has more than 20 years of experience across delivering projects for global clients as well as technology and product innovation. He has a special interest in Research and Development for advanced solutions leveraging AI, Computer Vision algorithms, Autonomous technologies and hardware across Humanoids, Drones, Brain Computer Interfaces etc. to solve some of the new and unique business problems.

About Us

The incubation center of Infosys called 'Infosys Center for Emerging Technology Solutions' (iCETS) focuses on incubation of NextGen services and offerings by identifying and building technology capabilities to accelerate innovation. The current areas of incubation include AI & ML, Blockchain, Computer Vision, Conversational interfaces, AR-VR, Deep Learning, Advanced Analytics using video, speech, text and much more.

To know more, please reach out to iCETS@infosys.com.

For more information, contact askus@infosys.com