

Lindenwood University

Digital Commons@Lindenwood University

Faculty Scholarship

Research and Scholarship

1-2024

Natural Language Processing and Neurosymbolic AI: The Role of Neural Networks with Knowledge-Guided Symbolic Approaches

Emily Barnes

James Hutson

Follow this and additional works at: <https://digitalcommons.lindenwood.edu/faculty-research-papers>



Part of the Artificial Intelligence and Robotics Commons

Original Article

Natural Language Processing and Neurosymbolic AI: The Role of Neural Networks with Knowledge-Guided Symbolic Approaches

Emily Barnes¹, James Hutson²

¹Lindenwood Global, Lindenwood University, MO, USA.

²Art History and Visual Culture, Lindenwood University, MO, USA.

¹ebarnes@lindenwood.edu

Received: 04 December 2024; Revised: 28 December 2024; Accepted: 10 January 2024; Published: 29 January 2024;

Abstract - Neurosymbolic AI (NeSy AI) represents a groundbreaking approach in the realm of Natural Language Processing (NLP), merging the pattern recognition of neural networks with the structured reasoning of symbolic AI to address the complexities of human language. This study investigates the effectiveness of neurosymbolic AI in providing nuanced understanding and contextually relevant responses, driven by the need to overcome the limitations of existing models in handling complex linguistic tasks and abstract reasoning. Employing a hybrid methodology that combines multimodal contextual modeling with rule-governed inferences and memory activations, the research delves into specific applications like Named Entity Recognition (NER), where architectures such as BiLSTM + CRF demonstrate improved accuracy by analyzing entire sentence contexts. The results affirm the potential of neurosymbolic AI in enhancing linguistic resolutions, semantic ambiguity resolution, and overall language understanding capabilities. Notably, the study showcases the significant strides in improving NER tasks, highlighting this approach's practical implications and effectiveness. The evolution of neurosymbolic AI, as indicated by this research, exemplifies the ongoing pursuit to create more sophisticated, accurate, and human-like interactions between machines and human language, promising a transformative impact on various sectors, including healthcare and education. The findings pave the way for future research and development in AI, pushing the boundaries of the role of technology in understanding and interacting with human language.

Keywords - Neurosymbolic AI, Natural Language Processing (NLP), Contextual modeling, Semantic ambiguity resolution, Named Entity Recognition (NER).

1. Introduction

Natural Language Processing (NLP) has long been a pivotal area in the realm of Artificial Intelligence (AI), offering profound insights into how machines understand and interact with human language. The evolution of NLP is a testament to the remarkable journey of AI from its inception to its current state, where the integration of neural networks with symbolic AI, known as Neurosymbolic AI (NeSy AI), or integrated neural-symbolic systems, is pushing the boundaries of technology [1]. Initially, the field of AI, emerging in the 1950s, was predominantly driven by symbolic approaches. These systems, rooted in logic and rule-based processes, laid the groundwork for early AI research, as seen in the development of expert systems and the initial forays into



machine understanding of language [2]. However, the limitations of these systems, particularly in handling the nuances and complexities of natural language, soon became apparent. In response, the latest advancements in NeSy AI are expanding the capabilities of NLP across various languages and applications, challenging the perception that English is the primary language for computers [3].

Moreover, integrating statistical methods like Machine Learning (ML) and data mining has brought AI-enabled devices closer to everyday life in sales and service sectors. Notably, NeSy AI models are excelling in domains like image and video reasoning, significantly enhancing Visual Question Answering (VQA) capabilities. Additionally, deep learning has contributed to successful implementations in machine translations, linguistic models, speech recognition, and automatic text generation [4]. The neural-symbolic systems also address crucial aspects of trust, safety, interpretability, and accountability in AI, making it an increasingly critical field in developing advanced NLP systems.

The 1980s and 1990s were a revolutionary period for NLP, marking significant advances and setting the stage for further development. The increased focus on machine learning and the expansion of computational capabilities accelerated the trend of using NLP in clinical text and other domains. Advances in ML, particularly non-sequential “parallel” methods, emerged during this time, offering new possibilities for language understanding. The period saw substantial growth in the ‘language industry’ and the ‘linguistic engineering’ field, with systems increasingly focusing on understanding users’ goals, intentions, and strategies [5]. Additionally, the late 1990s and early 2000s introduced full-fledged Bayesian machinery to NLP, enriching the field and accommodating shortcomings in the frequentist approach. Notably, statistical methods, introduced in the 1950s, were revived in the 1990s and proved effective in more ways than expected, encouraging new thinking about language and information processing.

These advances in neural networks and ML led to dramatic breakthroughs in NLP, exemplified by the introduction of models like BERT (2018) and GPT-3 (2020), which provided a powerful base for universal language understanding and generation [15]. BERT (Bidirectional Encoder Representations from Transformers) was first introduced in a paper by researchers from Google AI Language in 2018. The advance represented a significant shift in NLP by utilizing the transformer architecture to pre-train deep bidirectional representations from unlabeled text by joint conditioning on both left and right contexts in all layers. As such, BERT revolutionized various NLP tasks, including question-answering, language inference, and others, by enabling fine-tuning, prompting, and text-generation approaches. The introduction of BERT marked a new era in developing NLP models, emphasizing the importance of deep, bidirectional contextual word representations [6].

Another significant advance shaping language processing today was introduced in a paper by OpenAI in June 2020. GPT-3, also known as the third version of the Generative Pre-trained Transformer, has become known for its ability to produce text that is indistinguishable from that written by humans and has significantly contributed to various fields, demonstrating the versatility and power of the model in generating high-quality content. Since its launch, GPT-3 has been instrumental in various applications, showcasing the capabilities of large-scale language models in understanding and generating human-like text. The advancements in deep learning and AI significantly positively impacted NLP, enhancing vectorization, word embedding, classification, and automated speech recognition capabilities. These advances were made possible by the parallel developments in Deep Neural Networks (DNNs) post-2010, which became state-of-the-art algorithms in ML, speech recognition, and NLP, finding utility in various applications ranging from drug design to bioinformatics [7].

Moreover, the era witnessed the evolution of Graph Neural Networks (GNNs), revolutionizing data mining and ML tasks, including speech recognition and natural language understanding. Neural networks re-emerged as powerful machine-learning models, yielding state-of-the-art results in fields like image recognition speech

processing. They were applied to textual natural language signals, marking a significant shift in the computational treatment of language and its cognitive and psychological links. The profound impact of neural networks and ML during this transformative era laid the foundation for the sophisticated NLP systems of today. They steered the field towards more adaptive, parallel computing, influenced by pioneers like Warren Sturgis McCulloch (1898-1969) and others and paved the way for the current and future developments in AI and language processing [8]. The period was indeed a pivotal shift, marking the transition from conventional rule-based systems to more dynamic, contextually aware, and powerful AI models capable of understanding and interacting with human language in unprecedented ways.

Despite substantial progress in neural networks and machine learning during the 1980s and 1990s, neural network-based models continued to face challenges, particularly in handling tasks requiring complex reasoning or understanding abstract concepts. While adept at pattern recognition, these models struggled with intricate tasks, exposing a critical gap in capabilities. This spurred the development of NeSy AI, an approach blending the adaptability of neural networks with the explicit reasoning of symbolic AI, aiming for a holistic solution with nuanced understanding and contextually relevant responses. Crucial for effective human-AI interaction, this hybrid approach addresses comprehensive challenges in neurosymbolic programming, such as scalability, optimization stability, and deployment [9].

Furthermore, the development of NeSy AI, integrating neural network-based learning with symbolic knowledge representation and logical reasoning, aims to significantly enhance the trust, safety, interpretability, and accountability of AI systems, representing a crucial advancement in overcoming the limitations of purely neural network models and paving the way for more capable, reliable, and human-compatible AI solutions. Delving into the frontiers of integrated neural-symbolic systems, researchers encounter unique challenges and opportunities that shape the approach to advancing the field. The paper navigates these challenges by examining the historical evolution of NLP and AI, focusing on the transition from symbolic approaches to contemporary neurosymbolic integration. The goal is to explore the complex landscape of NeSy AI and its future implications, particularly its potential to revolutionize human-AI interactions in unprecedented ways, marking a pivotal era in the quest for AI systems capable of navigating the diverse and intricate demands of the real world [10].

Given the potential significance of NeSy AI, this treatment shall methodically review existing literature and advancements in the field, emphasizing the need for extensive, diverse datasets and considerable computational resources for training neural models alongside the ethical considerations vital in sensitive domains. It then delves into how NeSy AI addresses learning challenges in NLP through situational representations and the creation of flexible computational models and how integrating symbolic reasoning with deep learning produces highly predictive and comprehensible models. The paper thus aims to clarify the current state of NeSy AI, pinpoint research gaps and opportunities, and discuss the ethical and practical ramifications of its integration into AI. The anticipated contribution is a pathway towards more effective, ethical, and reliable AI systems, enhancing AI research and applications across various sectors. By tackling the challenges and harnessing the potential of these integrated neural-symbolic systems, this paper aspires to advance the development of sophisticated, human-compatible AI systems, leading to improved and more intuitive human-AI interactions.

2. Literature Review

The foundational theories of neural networks and symbolic AI, along with their integration into Neurosymbolic AI (NeSy AI), form the cornerstone of modern artificial intelligence. The development of neural networks, inspired by the biological neural networks in the human brain, has been deeply influenced by early work such as that of McCulloch and Pitts on artificial neurons published in 1943 [8]. This groundwork has been instrumental in advancing neural network architectures and learning algorithms. The introduction of the backpropagation algorithm by Rumelhart, Hinton, and Williams [11] also marked a significant advancement by

enabling the efficient training of multi-layer networks, leading to breakthroughs in complex pattern recognition tasks. Since then, neural networks have shown exceptional ability in learning from vast datasets, particularly in tasks involving complex pattern recognition like image and speech processing. This capability has been enhanced by associative processing principles, which aim to leverage meaning and combined sub-symbolic/symbolic operation, potentially improving artificial cognition. GNNs have further evolved as models for neural-symbolic computing across various domains, contributing to improved explainability, interpretability, and trust in AI systems [12].

In parallel, symbolic AI, emphasizing logic-based, high-level symbolic representations for problem-solving, has been integral to tasks requiring explicit, interpretable reasoning. Originating from the foundational works of researchers like John McCarthy and Marvin Minsky in the 1950s and 1960s, it utilizes human-understandable symbols and logic for knowledge representation and reasoning, offering high interpretability but limited flexibility with unstructured data [13]. The advent of Knowledge-infused Learning (K-iL) marks a significant development, integrating deep learning with structured knowledge to enhance learning processes and advance neuro-symbolic learning approaches. This integration forms the basis of NeSy AI, a blend of neural network-based learning with symbolic reasoning, emerging as a critical solution to ensure trust, safety, interpretability, and accountability in AI systems. The approach seeks to overcome the rigidity of traditional symbolic AI while preserving its explicit reasoning capabilities, aiming for a more adaptable, reliable, and human-aligned AI future.

The history of symbolic AI is marked by two primary paradigms: symbolism, focusing on logical processes, and connectionism, centered around artificial neural networks. NeSy AI models, which integrate these paradigms, promise high accuracy and reduced training data needs, albeit with certain limitations in parallelism due to their intricate control flow and operational intensity [14]. However, the integration of neural-symbolic computing has been effective in creating AI systems that are both explainable and accountable, leveraging the best of machine learning and reasoning. A seminal work in this area is *Neural-Symbolic Cognitive Reasoning* by Garcez, Lamb, and Gabbay [13], which explores the possibilities of combining connectionist learning with symbolic reasoning. Their work, along with others like the overview of capsule networks by Dombetzki [15], illustrates the ongoing efforts to create systems that not only learn from large datasets but also reason and make decisions based on structured knowledge and rules. In the realm of data science, symbolic AI offers valuable contributions by providing clear and interpretable data representations, thereby enhancing natural language analysis.

Despite its contributions, AI, including symbolic AI, encounters epistemological and methodological challenges, often inherited from psychological paradigms dominated by empiricism. This leads to issues in meaning, abstraction, generalization, and the development of higher cognitive functions. Nonetheless, the abilities of symbolic AI have significantly shaped the AI landscape despite limitations in handling non-symbolic processes and certain rigidity. Research in this domain has yielded innovative applications and theoretical progress, with developments like the Neuro-Vector-Symbolic Architecture (NVSA) improving accuracy and speed in complex tasks by integrating deep neural networks with symbolic AI [16]. Further, neural-symbolic computing is recognized for merging ML and reasoning in a principled manner, offering both explainable and accountable solutions and thus reinforcing the role of symbolic AI as a vital contributor to the advancement of artificial intelligence and its ethical grounding.

The historical evolution of AI, spanning from the inception of neural networks to the development of symbolic AI and their subsequent convergence in NeSy AI, illustrates an enduring quest to replicate and enhance human cognitive capabilities (Table 1). Research reveals that neural networks and symbolic AI, initially perceived as disparate or even competing paradigms, have evolved to be recognized as complementary forces. Collectively, they underpin a more comprehensive and nuanced approach to artificial intelligence.

Table 1. Comparative analysis of key terms in neurosymbolic AI and related technologies

Key Term	Definition	Strengths	Weaknesses
Neurosymbolic AI (NeSy AI)	AI that integrates neural network-based learning with symbolic reasoning to create more holistic AI solutions.	Enhanced reasoning and interpretability	Complexity in integration
		Capability to handle complex tasks	It may require large computational resources
Natural Language Processing (NLP)	A field of AI that focuses on the interaction between computers and human language, particularly how to program computers to process and analyze large amounts of natural language data.	Wide application in text analysis	Challenges in understanding context and sarcasm
		Improved human-computer interaction	Language ambiguity
Named Entity Recognition (NER)	A subtask of information extraction that seeks to locate and classify named entities mentioned in unstructured text into predefined categories.	Essential for data extraction and organization	Can struggle with ambiguous entities
		Useful in various NLP applications	Dependent on the quality and scope of training data
Bidirectional Encoder Representations from Transformers (BERT)	A transformer-based machine learning technique for NLP that pre-trains deep bidirectional representations from the text by joint conditioning on both left and right contexts.	High performance on various NLP tasks	Requires significant computational power
		Deep understanding of language context	Can be prone to biases in training data
Generative Pre-trained Transformer (GPT)	A type of language model that uses deep learning to produce human-like text.	Advanced text generation capabilities	High computational cost
		Flexibility and adaptability in various tasks	Ethical concerns with generated content
Graph Neural Networks (GNNs)	A type of neural network that directly works on the graph structure, applying convolution over the graph nodes and edges.	Effective in handling relational data	Complexity in understanding and implementation
		Useful in recommendation systems and social network analysis	Sensitive to graph quality and structure
Neural Network	A series of algorithms that attempt to recognize underlying relationships in a data set through a process that mimics how the human brain operates.	Excellent at pattern recognition and classification	Requires large datasets for training
		Adaptive learning abilities	Black-box nature makes them hard to interpret
Symbolic AI	A type of AI that uses symbolic representations of problems and logic to solve them, akin to human deductive reasoning.	High interpretability and explainability	Struggles with handling fuzzy, unstructured data
		Well-defined logic and rules	Rigid and lacks learning adaptability

The evolution from distinct methodologies to an integrated framework is a testament to the adaptability of the field and its relentless pursuit of more sophisticated, ethically grounded, and human-centric AI systems. Continued exploration into the intricacies and applications of NeSy AI positions it to significantly reshape the domain of AI, heralding advancements in capability, interpretability, and reliability within complex real-world contexts. The journey from foundational theories to current innovations underscores an ongoing progression towards crafting AI systems that are not only powerful but also attuned to augmenting human-AI synergy. This evolution propels the discourse into the next section, which delves into the methodological challenges inherent in NeSy AI, critically examining the barriers and solutions in this emerging field.

2.1. Methodological Challenges in Neurosymbolic AI

Integrating neural networks with symbolic AI in neurosymbolic AI presents several methodological challenges that are pivotal to advancing this field. One of the primary challenges lies in effectively combining the learning capabilities of neural networks with the reasoning and interpretability of symbolic AI. This integration is not merely about connecting two systems but about creating a seamless interaction where each complements the other's strengths and weaknesses. As highlighted in *Neural-Symbolic Learning Systems: Foundations and Applications* by Garcez, Lamb, and Gabbay [13], this integration requires an intricate balance between the empirical learning methods of neural networks and the deductive reasoning methods of symbolic AI. Another significant challenge is ensuring that the integrated system is transparent and interpretable. Neural networks, particularly deep learning models, are often criticized for their 'black box' nature, making it difficult to understand how they arrive at specific conclusions. This opaqueness stands in stark contrast to the transparent nature of symbolic systems. Garcez et al. [14] emphasize the importance of developing systems where the reasoning process is understandable and explainable. The challenge is maintaining this level of transparency even as the system becomes more complex through integration.

Data representation forms another methodological hurdle. Neural networks require data in a format vastly different from that used by symbolic AI systems. While neural networks work with numerical data, symbolic AI operates on high-level, human-readable symbols. Bridging this data representation gap without losing each approach's nuances is a complex task. Hitzler and van Harmelen [1] discuss the previous difficulties in creating a common representational framework compatible with both paradigms.

Furthermore, the scalability of neurosymbolic systems is a critical challenge. As the complexity of tasks increases, the computational resources required to train and run these systems escalate. This issue has been highlighted over the past two decades, as seen in Bengio et al.'s work [18] on deep learning, which requires extensive computational power and large datasets for effective training. As such, Graziani et al. [17] recently noted that integrating symbolic AI could potentially exacerbate this issue, as symbolic systems often involve complex logic and large knowledge bases.

Addressing these challenges is essential for the practical application of NeSy AI systems. For instance, in developing AI for healthcare diagnostics, it is crucial not only that the AI can learn from medical data (neural approach) but also that its diagnostic decisions can be understood and rationalized (symbolic approach). The balance between these two aspects is critical for creating AI systems that are both effective and trusted by users. Addressing these challenges will require a concerted effort from researchers across disciplines, combining empirical and theoretical insights to create AI systems that are both powerful and user-friendly. The goal is to develop AI that can learn and adapt like a neural network while reasoning and explaining like a symbolic system.

2.1.1. Neural Networks in NLP

Neural networks have significantly revolutionized NLP, enabling groundbreaking advancements in how machines understand and generate human language. Before the advent of neural networks, NLP relied heavily

on rule-based and statistical methods, which were limited in their ability to process natural language flexibly and context-awarely [20]. Neural networks, particularly deep learning models, introduced a paradigm shift using layered architectures that mimic human brain functioning to process, learn from, and generate language. These models excel in identifying patterns in large datasets, allowing them to understand syntax, semantics, and even some aspects of pragmatics in language. The shift has enabled machines to perform complex language tasks such as translation, summarization, and sentiment analysis with unprecedented accuracy and efficiency, transforming NLP from a niche research area into a cornerstone of modern AI applications.

Key breakthroughs in neural network-based NLP have largely been driven by advancements in model architecture and learning algorithms. Transformer models, introduced in the seminal paper “Attention is All You Need” by Vaswani et al. [21], marked a significant milestone. They eschewed the sequential processing of traditional recurrent neural networks for a parallel approach, using mechanisms called attention and self-attention to process different parts of the input data simultaneously. This innovation significantly improved the efficiency and effectiveness of NLP models, enabling the handling of longer sequences of data and a better understanding of context. The Generative Pretrained Transformer (GPT) series exemplifies the power of transformer models. Each iteration, from GPT to GPT-3 and beyond, has showcased increasingly sophisticated language understanding and generation capabilities, handling tasks like language translation, question-answering, and creative content generation with remarkable proficiency. These models have enhanced existing applications and opened new avenues for human-AI interaction in fields ranging from customer service to creative industries.

Despite these advances, purely neural network-based NLP systems come with inherent limitations. As with deep learning models, one of the primary challenges is the ‘black box’ nature of these models. Due to their complexity and the opaque manner in which they process information, understanding how they arrive at specific conclusions or outputs can be challenging, raising concerns about interpretability and trustworthiness. Additionally, while these models are excellent at pattern recognition, their reliance on large datasets for training can lead to biases in their outputs, reflecting biases present in the training data. This issue can lead to ethical concerns, especially in sensitive applications like sentiment analysis or content moderation.

Furthermore, neural network models often require extensive computational resources for training and operation, making them less accessible for researchers and organizations with limited resources. Lastly, their performance is heavily dependent on the quality and quantity of training data, and they can struggle with tasks that involve abstract reasoning, common-sense knowledge, or understanding of novel concepts not well-represented in the training data [22]. These limitations highlight the need for integrating other AI approaches, such as symbolic AI, to address the gaps and enhance the overall effectiveness of NLP systems.

Integrating Symbolic Approaches

Symbolic AI, often regarded as the initial form of AI, relies on manipulating high-level, human-readable symbols to perform tasks involving knowledge representation and reasoning. The approach, deeply rooted in the traditions of logic and philosophy, uses symbols to represent objects, facts, and relationships within a given domain. It employs logical rules to infer new facts or make decisions based on these representations. This type of AI excels in tasks requiring explicit reasoning, handling abstract concepts, and applying well-defined rules and knowledge. Unlike neural networks, which learn from data, symbolic AI systems are programmed with a set of predefined rules and knowledge bases.

This method allows for a clear and interpretable decision-making process, making it particularly useful in domains where explainability is crucial, such as medical diagnosis or legal analysis. The combination of neural networks with symbolic methods in AI presents a compelling approach to overcome the limitations inherent in

each. Neural networks are adept at learning patterns from data, handling ambiguous or incomplete information, and processing large-scale datasets. However, they lack the ability to reason explicitly or handle abstract concepts well. On the other hand, Symbolic AI provides these capabilities but struggles with learning from data and handling noisy or unstructured information. By integrating these two approaches, NeSy AI aims to leverage both strengths. The fusion allows for systems that cannot only learn and adapt based on data but also reason and make decisions based on structured knowledge and rules. The advantages of such integration include improved interpretability, more robust handling of novel situations, and the ability to incorporate common-sense reasoning. However, challenges remain, such as integrating learning and reasoning in a seamless manner, managing the complexity of combining two fundamentally different approaches, and ensuring that the system outputs are accurate and explainable.

Several case studies and examples demonstrate the practical application and benefits of NeSy AI. One notable example is Project Debater by IBM, a system in development since 2012 that can debate humans on complex topics [23]. It combines neural networks for processing and understanding natural language with symbolic AI for structuring arguments and applying rules of debate. The integration allows the system to not only comprehend and generate language but also to form coherent, logical arguments based on structured knowledge. Another example is in the field of medical diagnosis, where neurosymbolic systems combine the pattern recognition capabilities of neural networks (e.g., identifying patterns in medical images) with the rule-based reasoning of symbolic AI (e.g., applying medical knowledge to diagnose conditions) [24]. This approach enhances the system's accuracy and provides healthcare professionals with explainable AI-based recommendations. These are but a few examples that highlight the potential of NeSy AI to transform various sectors by providing more advanced, interpretable, and reliable AI systems.

3. Discussion

3.1. Enhancing Human-AI Interaction

At the same time, the advent of neurosymbolic AI represents a significant leap in enhancing human-AI interaction, fundamentally changing how humans and machines communicate and collaborate. Traditional AI systems, primarily based on neural networks, have made strides in understanding and responding to human language. However, their interactions often lack depth in terms of reasoning and context-awareness, which are crucial for more meaningful communication. NeSy AI, on the other hand, with its integration of neural networks and symbolic reasoning, brings a more nuanced understanding of language and logic to AI systems. This hybrid approach allows AI not only to process and analyze large volumes of data but also to apply logical reasoning and contextual understanding akin to human thought processes. As a result, AI systems can engage in more sophisticated dialogues, understand abstract concepts, and provide responses that are not just data-driven but also contextually relevant and logically sound. The advancement leads to more natural, intuitive, and effective human-AI interactions, paving the way for AI systems to become more integral and trusted partners in various aspects of life and work.

The future of prompting and prompt engineering in AI is also being reshaped by the advancements in NeSy AI. Sophisticated AI models that integrate neural networks with symbolic reasoning transform how AI understands and responds to human queries. Prompting, the process of providing AI with an initial input or query to elicit a response, becomes more intricate and potent with these advanced models. They can interpret prompts not just at a surface level but with an understanding of deeper meanings, intentions, and contexts. This means that AI can provide more accurate, relevant, and nuanced responses, moving beyond literal interpretations to grasp the subtleties and complexities of human language. In prompt engineering, where the goal is to design prompts that effectively guide AI responses, NeSy AI enables the creation of more sophisticated prompts that can leverage the reasoning capabilities of the system. These new abilities lead to interactions where users can communicate with AI more naturally and in a less structured way, akin to human conversation, making AI

systems more user-friendly and accessible to a broader audience. Beyond advancing conversational engagement, facial engagement and emotional intelligence are emerging frontiers in AI where NeSy AI plays a crucial role. Interpreting and responding to non-verbal cues, such as facial expressions and emotional signals, is a complex task that requires not only the recognition of patterns but also the understanding of subtle context and cultural nuances. NeSy AI, with its combination of pattern recognition (neural networks) and symbolic reasoning, is uniquely suited for this challenge. It allows AI systems to detect facial expressions and emotional cues through neural network-based image and signal processing and interpret these cues within the appropriate context using symbolic reasoning. This capability enables AI systems to respond in a manner that is more empathetic, appropriate, and emotionally intelligent. Such advancements in AI could revolutionize fields like customer service, therapy, and education, where understanding and responding to emotional cues are crucial. By bridging the gap between emotional understanding and AI, NeSy AI, in particular, is set to create more empathetic, responsive, and human-like interactions between machines and humans.

3.1.1. Applications and Implications

NeSy AI is finding practical applications across a diverse range of fields, each benefiting from its enhanced capability for both data-driven learning and logical reasoning. In healthcare, for instance, these integrated neural-symbolic systems are being used to improve diagnostic accuracy and personalize treatment plans. As noted, by combining the pattern recognition abilities of neural networks with the structured, rule-based reasoning of symbolic AI, these systems provide healthcare professionals with more accurate, comprehensive, and explainable diagnostic and treatment recommendations. In customer service, NeSy AI aids in creating more sophisticated and context-aware chatbots and virtual assistants. These systems can understand and respond to customer queries with pre-programmed answers and by reasoning and adapting to the context of each interaction, significantly improving customer experience. In education, these connectionist-symbolic systems are being employed to develop adaptive learning systems that tailor educational content to individual learning styles and students' progress, understanding not just the input data but also the logic and principles behind educational methods.

While there are demonstrable benefits for integrating such systems, ethical considerations and potential risks are paramount in deploying them. One of the primary concerns is privacy, especially in applications like healthcare and education, where sensitive personal data are involved. Ensuring these systems adhere to strict data privacy and security standards is crucial. Another significant concern is AI bias. While integrating symbolic AI with neural networks can help mitigate some forms of bias by incorporating rule-based logic and ethical guidelines, the risk of biases being introduced through training data or symbolic rules still exists. Having diverse, representative datasets and transparent, ethical rule-setting in developing these systems is essential. Additionally, there are concerns about the interpretability and accountability of AI decisions, especially in critical applications. Ensuring that NeSy AI systems are effective, transparent, and accountable is vital for maintaining public trust and ensuring ethical AI usage.

Looking towards the future, NeSy AI is poised to continue evolving and significantly influencing human-AI interaction with potential benefits in a range of areas (Table 2). As research in this field progresses, we can expect these systems to become more sophisticated, with enhanced abilities to understand and reason in increasingly human-like ways. This evolution will likely lead to AI systems that are more intuitive, adaptable, and capable of handling complex, nuanced tasks across various domains. In the future, we might see widespread adoption of such hybrid neural-symbolic computing systems in areas like autonomous vehicles, where AI must make rapid, logical decisions, or in environmental monitoring and management, where AI can help make predictions and decisions based on data patterns and environmental regulations. Integrating emotional intelligence and cultural context understanding in AI systems will also become more refined, leading to more personalized and empathetic interactions.

Table 2. Neurosymbolic AI: Core arguments and future use case potentials

Aspect	Argument Description	Potential Future Use Cases
Integration of Paradigms	Neurosymbolic AI integrates neural networks with symbolic AI for a holistic approach to problem-solving.	Enhanced robotic reasoning
		Sophisticated virtual assistants
Cognitive Mimicry	Aims to mimic and augment human cognitive abilities for improved AI interactions.	Adaptive educational platforms
		Advanced cognitive health diagnostics
Complementarity	Neural Networks and Symbolic AI are complementary, leveraging the strengths of both for better AI systems.	Complex decision-making systems
		Integrated system diagnostics and repair
Human-AI Synergy	Enhancing human-AI synergy to create more natural and effective interactions.	Collaborative creative design tools
		Interactive, personalized customer service
Ethical AI	Addressing ethical considerations by promoting transparency and accountability in AI.	Fair and unbiased hiring tools
		Transparent decision-making in finance and healthcare
Advanced Applications	The potential for Neurosymbolic AI in various advanced applications.	Precision medicine
		Autonomous vehicles
		Smart city management

However, as these technologies advance, the importance of addressing ethical challenges and societal impacts will become ever more critical. Balancing technological innovation with responsible AI development will be key to harnessing the full potential of NeSy AI in enhancing human-AI interaction.

Challenges and Future Research

One of the primary challenges in the field of NeSy AI and NLP is data availability and quality. While neural networks require large datasets for effective training, the availability of such data, particularly datasets that are diverse, unbiased, and representative of real-world scenarios, is often limited. The scarcity of quality data can lead to models that do not perform well across different demographics or perpetuate existing training data biases. Another challenge is the computational resources required for training and running sophisticated neurosymbolic models. These systems often need significant processing power and memory, making them inaccessible for researchers or organizations with limited resources [13]. Additionally, interpretability remains a critical challenge. While symbolic AI components offer some level of transparency, the neural network components of these hybrid systems are often seen as 'black boxes', making it difficult to understand how they arrive at certain conclusions or decisions [10]. The lack of transparency can be a significant hurdle, especially in applications where explainability is crucial for trust and accountability, such as in healthcare or legal contexts.

Future research and development in these areas could take several directions to address these challenges. One area of focus could be the development of more efficient training algorithms and models that require less computational power, making the technology more accessible and sustainable. Another key area could be enhancing the interpretability of these systems. This could involve creating hybrid models that not only combine the strengths of neural networks and symbolic AI but also incorporate mechanisms for explaining their reasoning processes in a human-understandable manner. Research could also explore ways to reduce bias in AI systems, including developing methods for creating more diverse and representative datasets and algorithms that can

identify and mitigate biases in training data and model outputs. Additionally, there is a need for further exploration into integrating emotional intelligence and cultural context in NeSy AI to make these systems more adept at handling the subtleties of human communication and interaction. In the realm of NLP, future research could focus on improving the ability of AI systems to understand and generate not just text but more complex forms of communication, such as sarcasm, humor, and cultural references. Advancements in this area could lead to more nuanced and sophisticated human-AI interactions.

There's also potential for exploring the application of Neurosymbolic AI in low-resource languages and dialects, which have traditionally been underserved in the field of NLP. This would not only broaden the reach and utility of NLP technologies but also help preserve and understand less common languages. Overall, the future research directions in Neurosymbolic AI and NLP promise to address current limitations and open up new possibilities for more advanced, equitable, and user-friendly AI systems.

4. Conclusion

The exploration of neurosymbolic AI in the realm of natural language processing has highlighted a significant advancement in the field of artificial intelligence and its interaction with humans. The integration of neural networks and symbolic AI approaches has led to the development of systems that not only process vast amounts of data but also apply logical reasoning and contextual understanding akin to human cognitive processes. The fusion of capabilities has resulted in AI systems capable of more intuitive, effective, and sophisticated human-AI interactions. These advancements are evident in various applications, ranging from customer service chatbots to complex medical diagnostic tools, where AI now provides more accurate, context-aware, and user-friendly interactions. The challenges identified, such as data availability, computational resource needs, and interpretability, underscore the ongoing journey in AI development. Nonetheless, the progress made so far offers a glimpse into a future where AI can work alongside humans in a more integrated and beneficial manner.

The significance of integrating neural networks with knowledge-guided symbolic approaches cannot be overstated. This integration marks a pivotal shift in the development of AI systems, moving beyond the limitations of purely data-driven, neural network-based models. By adding the dimension of symbolic reasoning, AI systems are not just reactive but also proactive in their interactions, capable of understanding and reasoning in previously challenging ways. This advancement enhances the ability of AI to handle complex, nuanced, and context-sensitive tasks, making it a more effective tool in a wide range of applications. The enhanced interpretability and ethical considerations that come with symbolic AI also contribute to building trust and reliability in AI systems, an essential factor as these technologies become more prevalent in everyday life.

Looking forward, the future of AI, particularly in the context of human life, is poised for transformative changes. As neurosymbolic AI continues to evolve, we can anticipate more seamless and integrated human-AI interactions, where AI can understand and respond to human needs and behaviors in a more empathetic and contextually relevant manner. The potential for AI to augment human capabilities, enhance decision-making processes, and contribute to various sectors of society is immense. However, it is crucial to proceed with a balanced approach, addressing ethical, privacy, and inclusivity concerns to ensure that the benefits of AI are accessible and equitable. The journey of AI is one of continuous learning and adaptation, reflecting not just technological advancement but also the evolving understanding of what it means to augment human intelligence and capabilities. The future of AI, intertwined with human endeavours, holds promise for a more efficient, informed, and interconnected world.

Authors' Contributions

Conceptualization, E. Barnes; Methodology, E. Barnes; Validation, E. Barnes; Investigation, J. Hutson – Original Draft Preparation, J. Hutson; Writing – Review & Editing, J. Hutson.; Visualization, J. Hutson.

References

- [1] Pascal Hitzler, and Frank van Harmelen, "A Reasonable Semantic Web," *Semantic Web*, vol. 1, no. 1-2, pp. 39-44, 2010. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Robin Cowan, "Expert Systems: Aspects of and Limitations to the Codifiability of Knowledge," *Research Policy*, vol. 30, no. 9, pp. 1355-1372, 2001. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [3] Kyle Hamilton et al., "Is Neuro-Symbolic AI Meeting its Promise in Natural Language Processing? A Structured Review," *arXiv*, pp. 1-27, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] Rabinandan Kishor, "Neuro-Symbolic AI: Bringing a New Era of Machine Learning," *International Journal of Research Publication and Reviews*, vol. 3, no. 12, pp. 2326-2336, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] Vamsi Krishna Vedantam, "The Survey: Advances in Natural Language Processing Using Deep Learning," *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, vol. 12, no. 4, pp. 1035-1040, 2021. [[Google Scholar](#)] [[Publisher Link](#)]
- [6] D. Deepa, "Bidirectional Encoder Representations from Transformers (BERT) Language Model for Sentiment Analysis Task," *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, vol. 12, no. 7, pp. 1708-1721, 2021. [[Google Scholar](#)] [[Publisher Link](#)]
- [7] Wahab Khan et al., "Exploring the Frontiers of Deep Learning and Natural Language Processing: A Comprehensive Overview of Key Challenges and Emerging Trends," *Natural Language Processing Journal*, vol. 4, pp. 1-31, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Michael A. Arbib, "Warren McCulloch's Search for the Logic of the Nervous System," *Perspectives in Biology and Medicine*, vol. 43, no. 2, pp. 193-216, 2000. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [9] Artur d'Avila Garcez, and Luis Lamb, "Neurosymbolic AI: The 3rd Wave," *Artificial Intelligence Review*, pp. 1-37, 2020. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [10] Amit Sheth, Kaushik Roy, and Manas Gaur, "Neurosymbolic Artificial Intelligence (Why, What, and How)," *IEEE Intelligent Systems*, vol. 38, no. 3, pp. 56-62, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning Representations by Back-Propagating Errors," *Nature*, pp. 533-536, 1986. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Alexiei Dingli, and David Farrugia, *Neuro-Symbolic AI: Design Transparent and Trustworthy Systems that Understand the World as You do*, Packt Publishing, pp. 1-196, 2023. [[Google Scholar](#)] [[Publisher Link](#)]
- [13] Artur S. D'Avila Garcez, Luis C. Lamb, and Dov M. Gabbay, *Neural-Symbolic Cognitive Reasoning*, Springer, pp. 1-197, 2009. [[Google Scholar](#)] [[Publisher Link](#)]
- [14] Artur d'Avila Garcez et al., "Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning," *arXiv*, pp. 1-27, 2019. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [15] Luca Alessandro Dombetzki, "An Overview over Capsule Networks," *Network Architectures and Services*, pp. 89-95, 2018. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Michael Hersche et al., "A Neuro-Vector-Symbolic Architecture for Solving Raven's Progressive Matrices," *Nature Machine Intelligence*, vol. 5, pp. 363-375, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [17] Mara Graziani et al., "A Global Taxonomy of Interpretable AI: Unifying the Terminology for the Technical and Social Sciences," *Artificial Intelligence Review*, vol. 56, pp. 3473-3504, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [18] Emmanuel Bengio et al., "Conditional Computation in Neural Networks for Faster Models," *arXiv*, pp. 1-12, 2016. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [19] Mara Graziani et al., "A Global Taxonomy of Interpretable AI: Unifying the Terminology for the Technical and Social Sciences," *Artificial Intelligence Review*, vol. 56, pp. 3473-3504, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] Hao Zheng et al., "Improving Deep Neural Networks Using Softplus Units," *2015 International Joint Conference on Neural Networks (IJCNN)*, Killarney, pp. 1-4, 2015. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [21] Ashish Vaswani et al., "Attention is all you need," *31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA, pp. 1-11, 2017. [[Google Scholar](#)] [[Publisher Link](#)]

- [22] Dongran Yu et al., "A Survey on Neural-Symbolic Learning Systems," *Neural Networks*, vol. 166, pp. 105-126, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [23] Roy Bar-Haim et al., "Project Debater APIs: Decomposing the AI Grand Challenge," *arXiv*, pp. 1-8, 2021. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [24] Samriddhi Srivastav et al., "ChatGPT in Radiology: The Advantages and Limitations of Artificial Intelligence for Medical Imaging Diagnosis," *Cureus*, vol. 15, no. 7, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [25] Judith E. Dayhoff, and James M. DeLeo, "Artificial Neural Networks: Opening the Black Box," *Cancer: Interdisciplinary International Journal of the American Cancer Society*, vol. 91, no. S8, pp. 1615-1635, 2001. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]