

FUZZY COGNITIVE MAPS BRIDGING TRANSPARENCY AND PERFORMANCE IN HYBRID AI SYSTEMS

Maikel Leon

Department of Business Technology, Miami Herbert Business School, University of
Miami, Florida, USA.

ABSTRACT

The combination of symbolic and sub-symbolic Artificial Intelligence (AI) provides an excellent opportunity for innovations that merge the interpretability of the former with the learning capabilities of the latter. This paper presents Fuzzy Cognitive Maps (FCMs) as a hybrid and flexible model that combines the strengths of both paradigms and proposes them as a feasible solution to the challenges of explainability and interpretability in AI systems without losing working feasibility. FCMs have emerged as a robust framework for representing causal knowledge and facilitating intuitive and justifiable decision-making processes, but there is much more to explore. FCMs can handle the inherent uncertainty and vagueness present in real-world scenarios, allowing for a more natural approach to problem-solving in combination with the learning and adaptation capabilities of sub-symbolic AI. FCMs are an ideal choice for applications requiring high levels of explainability and interpretability.

KEYWORDS

Fuzzy Cognitive Maps, Symbolic AI, Sub-symbolic AI, Explainable AI, and Interpretable AI.

1. INTRODUCTION

In the rapidly changing world of AI, two distinct paradigms have emerged, each with its unique approach to modeling intelligence and solving problems. On one side of the spectrum lies symbolic AI, an approach grounded in representing knowledge through explicit symbols and rules, mirroring the logical structures of human thought. Conversely, subsymbolic AI eschews these clear-cut representations for a more opaque yet powerful method of learning directly from data, embodying the patterns and statistical correlations that underpin intelligence in a way that's often incomprehensible to human observers. These paradigms, seemingly at odds, represent the dual paths through which AI has sought to replicate or surpass human cognitive capabilities. Yet, as we delve deeper into the strengths and limitations inherent in each approach, a compelling narrative emerges—one that suggests the future of AI may not rest on the supremacy of one paradigm over the other but on the synergy of both. Symbolic AI relies on the manipulation of symbols and execution of logical operations to perform tasks, solve problems, and make decisions. This approach, foundational to early AI research, excels in domains where rules are well-defined and outcomes are predictable. Its transparency and interpretability, where every decision can be traced through a logical chain of reasoning, offer clear advantages in applications demanding explainability and compliance with regulatory standards.

However, symbolic AI's rigidity, reliance on exhaustive rule sets, and difficulty encoding commonsense knowledge have limited its applicability to complex, real-world problems where

ambiguity and uncertainty are the norms. Symbolic AI is a reasoning-oriented field that relies on classical logic (usually monotonic) and assumes that logic makes machines intelligent. For instance, if you ask yourself, with this paradigm in mind, "What is an apple?" the answer will be that an apple is "a fruit," "has red, yellow, or green color," or "has a roundish shape." These descriptions are symbolic because we utilize symbols (color, shape, and kind) to describe an apple. Between the 50s and the 80s, it was the dominant AI paradigm. Regarding the implementation of symbolic AI, one of the oldest yet still the most popular logic programming languages is Prolog (its roots are in first-order logic) [1].

On the other hand, sub-symbolic AI, which includes neural networks and deep learning, takes a different approach. By learning directly from large amounts of data, sub-symbolic AI models develop an internal understanding of the world that is useful for tasks such as pattern recognition, language processing, and predictive modeling, often outperforming humans. However, this capability comes at the expense of transparency, leading to the "black box" problem where the rationale behind a model's decision is not easily understandable or explainable. The central idea of the sub-symbolic approach is that the ability to produce a good model with limited experience leads to a model's success. Instead of clearly defined, human-readable relationships, we create less explainable mathematical equations to solve problems. Neural networks, ensemble models, regression models, decision trees, and support vector machines are some of the most popular sub-symbolic AI models, especially in developing machine learning models. During the 80s, the sub-symbolic AI paradigm took over symbolic AI's position as the leading subfield [2].

The dichotomy between symbolic and sub-symbolic AI has led to a vibrant discourse on the future direction of AI research and application. Within this discourse, Fuzzy Cognitive Maps (FCMs) emerge as a fascinating hybrid technique, combining the explicit knowledge representation of symbolic AI with the adaptability and learning capabilities of sub-symbolic AI [3]. FCMs use fuzzy logic to handle uncertainty and model complex systems through networks of concepts and causal relationships. This connects the deterministic world of symbols with the probabilistic nature of sub-symbolic learning. Combining symbolic and sub-symbolic approaches can unleash unprecedented capabilities as we enter a new era in AI. Symbolic AI's interpretability, structured knowledge representation, and the learning efficiency and adaptability of sub-symbolic AI can pave the way for more advanced, flexible, and reliable AI systems. This article explores the contrasting strengths and weaknesses of symbolic versus sub-symbolic AI, highlights FCMs as a prime example of hybrid AI techniques, and speculates on a future where AI's full potential is realized through the harmonious integration of both paradigms [4]. In doing so, we may find that the future of AI is not a question of either/or but a confluence of both, harnessing the best of what each approach has to offer.

The rest of this paper is organized as follows. Sec. 2 presents the origins and notable cases of this classical approach to AI. Sec. 3 refers to theoretical conceptions in Machine Learning. Sec. 4 presents the idea of the need for suitable explanations offered by these systems. Sec. 5 digs deep into why AI's future should contain more traceable and interpretable models. Sec. 6 holds the idea of merging both symbolic and subsymbolic approaches. Sec. 7 highlights the well-known Artificial Neural Networks' relevance in connectionist computing. Sec. 8 introduces a paradigm aiming to benefit from symbolic and subsymbolic AI. Last, Sec. 9 serves as a reflection and to understand the need for new and more AI models that are solid computationally and transparent to human understanding.

2. SYMBOLIC AI

Symbolic AI, or "Good Old-Fashioned Artificial Intelligence," refers to a branch of AI research and development emphasizing symbolic representations of problems, logic, and search. This approach to AI relies on manipulating symbols and expressions to perform tasks, solve problems, and model the world. The following report delves into symbolic AI's origins, notable case studies, advantages, and disadvantages.

2.1. Discovering the Roots of Symbolic AI

The roots of Symbolic AI can be traced back to the mid-20th century, with foundational work by influential figures such as Alan Turing, John McCarthy, and Marvin Minsky. Alan Turing's conceptualization of the Turing machine and the Turing test laid the groundwork for thinking about machines that could simulate human intelligence. In the 1950s and 1960s, John McCarthy, often considered one of the fathers of AI, coined the term "artificial intelligence" and introduced the concept of using symbolic logic to represent and solve problems. Marvin Minsky's work on frames and knowledge representation further advanced the development of symbolic AI. The period from the 1950s to the late 1980s is often considered the golden age of symbolic AI, during which researchers focused on developing systems that could reason about the world using symbolic logic. This era saw the creation of expert systems, among the first commercial applications of AI. These systems used rules and databases of knowledge to make inferences and provide advice in specialized domains such as medicine and engineering.

2.2. Relevant Case Studies

- MYCIN: Developed in the early 1970s at Stanford University, MYCIN was an expert system designed to diagnose bacterial infections and recommend antibiotics. It was one of the first successful demonstrations of symbolic AI in medicine, using a rule-based system to make decisions.
- SHRDLU: Created by Terry Winograd in the 1970s, SHRDLU was a natural language understanding system that could interact with a user in English to move blocks around a virtual world. It demonstrated the potential of symbolic AI for understanding and manipulating language and objects in a constrained environment.
- Deep Blue: Although primarily known for its chess-playing ability, IBM's Deep Blue represents a blend of symbolic AI (in terms of chess strategy and positions represented symbolically) and brute-force computation. In 1997, Deep Blue famously defeated world chess champion Garry Kasparov, showcasing the potential of AI in complex decision-making.

2.3. Advantages

- Explainability: One of the primary advantages of symbolic AI is its inherent explainability. Because decisions are made through explicit logical rules, it is easier to understand and trace symbolic AI systems' reasoning processes than more opaque models like deep neural networks.
- Efficiency in Domain-Specific Knowledge: Symbolic AI systems excel in domains where knowledge can be clearly defined and encoded in rules. This makes them particularly useful for expert medicine, law, and engineering systems.
- Handling Logical Reasoning and Complex Problems: Symbolic AI is well-suited for tasks that involve complex problem-solving and logical reasoning, where clear rules and relationships can be established.

2.4. Limitations

- Knowledge Acquisition Bottleneck: One of the major challenges of symbolic AI is the knowledge acquisition bottleneck. Encoding expert knowledge into rules and symbols is time-consuming and requires significant expertise. This makes scaling symbolic AI systems difficult.
- Lack of Flexibility: Symbolic AI systems are often criticized for lacking flexibility and adaptability. They struggle with handling uncertainty, learning from new data, and performing in unstructured environments.
- Limited Perception and Learning: Unlike their machine learning counterparts, symbolic AI systems have limited abilities to learn from data or perceive complex patterns without explicitly programmed knowledge. This limits their applicability in tasks that require significant generalization or data-driven learning.

2.5. Summary

Symbolic AI has been a critical player in the development of AI. Its focus on logic, explicit knowledge representation, and symbolic reasoning has led to significant progress in understanding and imitating human intelligence. However, the limitations of symbolic AI, especially in terms of scalability, flexibility, and learning, have given rise to alternative approaches such as machine learning and neural networks. Despite these challenges, the advantages of symbolic AI, particularly its explainability and effectiveness in specific domains, remain essential for study and practical application. Hybrid approaches that combine the strengths of symbolic AI with machine learning are emerging as a promising way to address the limitations of both paradigms. As AI progresses, the principles of symbolic AI are likely to continue shaping the development of intelligent systems, contributing to our understanding and implementation of AI.

3. SUB-SYMBOLIC AI

Sub-symbolic AI is a paradigm in AI research different from the traditional symbolic approach. Unlike symbolic AI, which relies on clearly defined symbols and rules to process and convey knowledge, sub-symbolic AI focuses on the underlying intelligence mechanisms. This approach aims to model the processes and patterns of thought that occur below the level of conscious, symbolic thought, drawing inspiration from the functioning of the human brain and biological systems. This report delves into the origins, notable case studies, and the advantages and disadvantages of sub-symbolic AI.

3.1. Discovering the Roots of Sub-symbolic AI

Sub-symbolic AI has its roots in the early days of AI research, but it gained significant momentum in the 1980s with the resurgence of neural networks and the development of algorithms that could learn from data. The limitations of symbolic AI, such as its inability to handle ambiguous or incomplete information and to learn from raw data, inspired researchers to investigate alternative models that could imitate the brain's ability to understand and generalize from experiences. The rise of connectionism, which focuses on the role of neural networks and parallel distributed processing in cognitive functions, marked a crucial shift towards sub-symbolic AI.

3.2. Relevant Case Studies

- Deep Learning for Image Recognition: Convolutional Neural Networks (CNNs), a class of deep neural networks, have revolutionized image recognition. A landmark moment was when AlexNet, a CNN designed by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, won the ImageNet Large Scale Visual Recognition Challenge in 2012, significantly outperforming traditional image recognition methods.
- Natural Language Processing (NLP): Sub-symbolic AI has dramatically improved the ability of machines to understand and generate human language. Google's BERT (Bidirectional Encoder Representations from Transformers) and OpenAI's GPT (Generative Pre-trained Transformer) series are prime examples of how deep learning models can grasp complex language patterns, enabling breakthroughs in translation, summarization, and question-answering systems.
- AlphaGo: Developed by DeepMind, AlphaGo is a program that defeated the world champion Go player in 2016. It used deep neural networks and reinforcement learning to master a game known for its complexity and strategic depth, a feat previously thought to be decades away [5].

3.3. Advantages

- Learning from Data: One of the most significant advantages of sub-symbolic AI is its ability to learn directly from data without explicit programming. This makes it incredibly powerful in handling complex, high-dimensional data such as images, speech, and text.
- Generalization: Sub-symbolic AI models, particularly deep learning networks, can generalize, meaning they can perform well on unseen data after training on a sufficiently large and representative dataset. This ability to generalize from examples is closer to human learning and is a key strength of sub-symbolic AI.
- Handling Ambiguity and Uncertainty: Unlike symbolic AI, sub-symbolic AI is adept at dealing with ambiguity and incomplete information. Neural networks, for instance, can make probabilistic predictions and decisions even in uncertain or incomplete data [6].

3.4. Limitations

- Opacity (Black-Box Problem): A significant drawback of sub-symbolic AI, especially deep neural networks, is its lack of transparency. These models are often described as "black boxes" because it is difficult to understand how they arrive at specific decisions or predictions, complicating efforts to debug or explain their behavior [7].
- Data and Computational Requirements: Training sub-symbolic AI models, particularly deep learning networks, requires vast data and significant computational resources. This can make cutting-edge AI research and applications inaccessible to organizations with limited resources [8].
- Overfitting and Generalization Issues: While sub-symbolic AI models are good at generalizing from data, they can also be prone to overfitting, where they perform well on training data but poorly on new, unseen data. To mitigate this risk, careful design, regularization techniques, and validation strategies are required [9].

3.5. Summary

Sub-symbolic AI has become a powerful approach, offering capabilities surpassing traditional symbolic methods in many areas, such as incredibly complex pattern recognition, learning from data, and generalization. The success of deep learning and neural networks has highlighted the

potential of sub-symbolic AI to solve previously unsolvable problems. However, challenges related to interpretability, data, resource requirements, and the risk of overfitting emphasize the need for ongoing research and development. The future of AI is likely to involve a hybrid approach that combines the strengths of symbolic and sub-symbolic AI, utilizing the transparency and structured knowledge representation of symbolic systems along with the learning capabilities and adaptability of sub-symbolic models.

4. Explainable AI

Explainable AI (XAI) refers to methods and techniques that make the output of AI systems transparent and understandable to humans. XAI aims to create a suite of machine learning techniques that produce more explainable models while maintaining high learning performance (accuracy), enabling human users to understand, trust, and effectively manage the emerging generation of artificially intelligent partners. This report covers the origins, notable case studies, advantages and disadvantages of explainable AI. As mentioned, the symbolic AI paradigm provides quickly interpretable models with satisfactory reasoning capabilities, allowing us to trace the reasoning for a particular outcome easily. Yet, expressing the entire relation structure, even in a specific domain, is difficult [10]. Symbolic AI models are not able to capture all possibilities without significant effort. In contrast, sub-symbolic AI models are highly successful and can be designed and trained with less effort while delivering impressive accuracy. However, one of the significant drawbacks of sub-symbolic models is the lack of explainability in the decision-making process. This becomes particularly important in sensitive domains such as court rulings, military actions, and loan applications, where understanding the reasoning behind the outcome is crucial. Therefore, high-performing but opaque models cannot be relied upon in these contexts.

4.1. The need for Explainable AI

The idea of explainable AI is not new and has been around since the early days of AI research. However, the focus on explainability has grown in recent years due to the increasing use of complex machine learning models, like deep learning, which are often perceived as "black boxes" because of their opaque decision-making processes. The need for explainability stems from concerns about accountability, fairness, transparency, and compliance with regulatory requirements, such as the European Union's General Data Protection Regulation, which includes a right to explanation. In the past, AI systems were more understandable as they heavily relied on symbolic AI approaches, like rule-based systems, where the reasoning behind decisions could be easily traced and understood. As the field shifted towards more powerful but less interpretable models, the demand for techniques to make these models explainable grew [11].

4.2. Relevant Case Studies

1. **Healthcare Diagnosis:** AI models are increasingly used to diagnose diseases from medical imaging. Researchers have developed XAI systems that can identify specific features in imaging data that lead to their diagnosis, providing doctors with insights into why the AI system made a particular diagnosis. This not only aids in validating the AI's conclusions but also enhances the doctor's understanding and trust in the tool.
2. **Financial Services for Loan Approval:** AI models evaluate loan applications in the financial sector. XAI can be crucial in explaining why a loan was approved or denied, ensuring compliance with regulations against discriminatory practices, and helping applicants understand what factors influenced the decision.
3. **Criminal Justice Risk Assessment Tools:** Tools like COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) have been used to assess the

likelihood of reoffending. XAI methods can help uncover, explain, and correct biases in such predictive models, ensuring fair and transparent decision-making.

4.3. Advantages

- Increased Trust and Confidence: Explainability builds trust among users and stakeholders by transparentizing decision-making. When users understand how an AI system arrives at its conclusions, they are more likely to trust it.
- Improved Model Debugging and Validation: XAI techniques enable developers to identify and correct errors or biases in AI models. Developers can make targeted adjustments to improve performance and fairness by understanding the factors influencing model decisions.
- Regulatory Compliance: Many industries are subject to regulations that require decisions made by automated systems to be explainable. XAI facilitates compliance with such regulations, enabling AI solutions deployed in highly regulated sectors like finance and healthcare.
- Ethical and Fair Decision-Making: Explainable AI can help identify and mitigate biases in AI models, promoting more ethical and fair decision-making processes. This is particularly important in applications with significant social implications, such as criminal justice and employment.

4.4. Limitations

- Potential Reduction in Model Performance: In some cases, making a model more explainable may require simplifying its architecture or using less complex algorithms, which can reduce accuracy or performance [12].
- Complexity and Resource Requirements: Developing explainable AI models can be more complex and resource-intensive than traditional models. It requires additional efforts in design, implementation, and validation to ensure that explanations are meaningful and accurate.
- Risk of Oversimplification: There is a risk that the explanations provided by XAI systems might oversimplify the underlying processes, potentially leading to misunderstandings or misplaced trust in the AI system's capabilities.
- Security and Privacy Concerns: Explaining how AI systems work might inadvertently reveal sensitive information about the data or the model itself, posing security and privacy risks.

4.5. Summary

Explainable AI is a critical advancement addressing AI systems' need for transparency, trust, and understanding. As AI becomes increasingly integrated into essential sectors of society, the importance of explainability will only grow. The challenge lies in balancing the demand for complex, high-performing AI models with the need for transparency and comprehensibility. Symbolic AI models are explainable by design, while sub-symbolic AI models are usually not. Two fields deal with creating high-performing AI models with reasoning capabilities, usually requiring combining components from symbolic and subsymbolic paradigms. While XAI aims to ensure model explainability by developing models that are easier to understand for their users, NSC focuses on finding ways to combine subsymbolic learning algorithms with symbolic reasoning techniques. Future developments in XAI will likely focus on innovative approaches to maintaining or enhancing model performance while providing clear, accurate, and helpful explanations [13]. As the field evolves, it will also be essential to develop standardized metrics

for explainability and ensure that explanations are accessible and understandable to all users, regardless of their technical background. Ultimately, the success of explainable AI will depend on its ability to foster trust and collaboration between humans and machines, enable more informed decision making, and ensure that AI systems align with societal values and ethical principles [14].

5. INTERPRETABLE AI

Interpretable AI focuses on developing models and algorithms that are easy for humans to understand. This means that users can comprehend and trace back the decisions, predictions, or classifications an AI system makes. Interpretable AI is crucial for applications in sensitive and critical domains where understanding the reasoning behind AI decisions is essential for trust, compliance, and improvement. This report delves into the origins, notable case studies, advantages, and disadvantages of interpretable AI.

5.1. What is Interpretable AI

The concept of interpretable AI can be traced back to the early days of AI when rule-based systems were more common. These systems were inherently interpretable, allowing users to follow the logical steps of AI to reach a decision. However, as AI research progressed, especially with the development of complex models like deep neural networks, the focus shifted towards improving performance, often at the expense of interpretability. The increasing use of AI systems in crucial areas such as healthcare, finance, and criminal justice has revitalized the importance of interpretability. Stakeholders in these fields require AI systems to make decisions and provide explanations that humans can understand. This need has led to developing new techniques and research to make even the most complex models interpretable.

5.2. Relevant Case Studies

- **Healthcare Diagnosis and Treatment:** AI systems are increasingly used to diagnose diseases and recommend treatments. For instance, models that predict cardiovascular diseases based on patient data must be interpretable so that health care providers can understand the reasoning behind the predictions. This ensures trust and allows healthcare professionals to make informed decisions.
- **Financial Services Compliance and Decision-Making:** In finance, AI models are used for credit scoring, fraud detection, and automated trading. Interpretability in these models helps users understand the factors influencing decisions, ensuring compliance with regulatory standards and building customer trust.
- **Criminal Justice and Bail Decisions:** AI is used to assess the risk of recidivism and inform bail and sentencing decisions. Using interpretable AI models in this context is crucial for fairness, transparency, and accountability, allowing for scrutinizing decisions that significantly impact individuals' lives.

5.3. Advantages

- **Trust and Transparency:** Interpretable AI fosters trust from users by making the decision-making process transparent. When stakeholders understand how decisions are made, they are more likely to trust and accept AI solutions.
- **Improved Decision-Making:** Interpretability allows users to verify the correctness of the AI's reasoning, leading to more informed and better decision-making. This is especially important in domains where decisions have significant consequences.

- Regulatory Compliance: Many industries are subject to regulations that require decisions to be explainable. Interpretable AI facilitates compliance with such regulations, avoiding potential legal and financial penalties.
- Error Detection and Model Improvement: By understanding how an AI system makes decisions, developers and users can identify errors or biases in the model, leading to continuous improvement of AI systems.
- Ethical Considerations: Interpretable AI can help identify and mitigate biases in AI systems, promoting fairness and ethical decision-making.

5.4. Limitations

- Potential Trade-off Between Interpretability and Performance: Sometimes, making a model more interpretable may require simplifying its architecture or using less complex algorithms, potentially leading to decreased accuracy or performance.
- Complexity in Interpretation: Achieving true interpretability can be challenging for complex models. Even when interpretations are provided, they may be difficult for non-experts to understand, limiting their usefulness.
- Risk of Misinterpretation: There's a risk that interpretations provided by AI systems might be misunderstood by users, leading to incorrect conclusions or decisions based on those interpretations.
- Time and Resource Intensive: Developing interpretable AI models can require additional time and resources. Designing models that balance interpretability and performance involves extra effort in model selection, development, and validation [15].

5.5. Summary

Interpretable AI is essential for the responsible deployment of AI, particularly in sensitive and high-stakes areas. It ensures transparency, trust, and ethical considerations in AI systems. As AI becomes more prevalent in society, the demand for interpretable models will likely grow, driving further research and development. Future advancements in interpretable AI will strive to balance performance and interpretability, establish standardized measures for interpretability, and provide more user-friendly explanations. This will help ensure that AI systems are robust, effective, and aligned with societal values and ethical standards, leading to greater acceptance and integration of AI technologies across different sectors.

6. THE CONVERGENCE IS POSSIBLE

The intersection between symbolic and sub-symbolic AI represents a fascinating and promising area of research within AI. This interests both worlds: symbolic AI's explicit reasoning and interpretability with the learning capabilities and adaptability of sub-symbolic AI, particularly neural networks. This hybrid approach aims to overcome the limitations inherent in each approach when used in isolation, enabling the development of AI systems that are both powerful and understandable. This report explores the origins, notable case studies, advantages, and disadvantages of the intersection between symbolic and subsymbolic AI.

6.1. Origins

The division between symbolic and sub-symbolic AI dates back to the early days of AI research. Symbolic AI, predominant in the initial stages of AI development, focuses on logic and rule-based systems. On the other hand, sub-symbolic AI, which gained prominence with the emergence of machine learning and neural networks, emphasizes learning from data and pattern recognition. Combining these two approaches stemmed from recognizing their complementary strengths and weaknesses. Symbolic AI excels in handling complex reasoning and explicit knowledge representation, while sub-symbolic AI is adept at working with raw data and learning from experience. This presented a strong case for integrating the two approaches.

6.2. Notable Case Studies

- Neuro-Symbolic AI for Visual Question Answering (VQA): Research projects have combined neural networks with symbolic reasoning to improve VQA systems, which answer questions about images. These hybrid systems use neural networks to interpret visual data and symbolic systems to reason about the content, enabling more accurate and interpretable answers.
- Commonsense Reasoning: Projects like OpenAI's GPT-3 have integrated symbolic reasoning to enhance the model's ability to perform commonsense reasoning tasks. These systems can better understand and generate human-like responses by embedding symbolic representations within a neural frame work.
- Robotics and Planning: Combining symbolic AI for high-level planning and decision making with sub-symbolic AI for perception and motion control has led to more versatile and efficient robots. This approach allows robots to navigate and interact with their environment in a more human-like manner, adapting to new tasks and environments through learning.

6.3. Advantages

- Enhanced Reasoning and Generalization: Integrating symbolic and sub-symbolic AI can lead to systems that learn from data and apply logical reasoning to generalize beyond their training data. This results in more flexible and capable AI systems [16].
- Improved Interpretability and Transparency: Symbolic components can provide clear explanations for the decisions made by sub-symbolic models, addressing one of the major drawbacks of purely sub-symbolic AI systems.
- Efficient Learning and Knowledge Representation: Symbolic AI can encode domain knowledge that guides the learning process of sub-symbolic models, making them more efficient and effective in learning from data. Conversely, sub-symbolic models can discover patterns and relationships that can be formalized into symbolic knowledge.
- Flexibility and Adaptability: Hybrid systems can adapt to new tasks and environments more readily by leveraging the learning capabilities of sub-symbolic AI with the structured knowledge representation of symbolic AI.

6.4. Limitations

- Complexity in Integration: Combining symbolic and sub-symbolic AI involves significant challenges, including integrating disparate representations and reasoning mechanisms. This complexity can make the development of hybrid systems more challenging and resource-intensive.
- Scalability Issues: The scalability of hybrid AI systems can be limited by the symbolic

component, which may not easily handle the vast amounts of data that sub-symbolic models can process [17].

- Limited Understanding of Integration Mechanisms: The field is still exploring the most effective way to integrate symbolic and sub-symbolic AI. This includes challenges in combining learning and reasoning, representing knowledge, and ensuring that the systems are robust and reliable.

6.5. Summary

The combination of symbolic and sub-symbolic AI shows great promise for the future of AI. By merging the strengths of both approaches, researchers and practitioners aim to create robust AI systems capable of learning from large amounts of data while also being able to reason, generalize, and explain their decisions in a way that is understandable to humans. This hybrid approach represents a step towards more sophisticated, versatile, and trustworthy AI systems that can be effectively applied in various domains, from healthcare and finance to autonomous systems [6]. Realizing the full potential of this intersection requires overcoming significant challenges. These challenges include integrating different AI paradigms, ensuring scalability, and developing effective mechanisms for combining learning and reasoning. Continued research and experimentation in this area are crucial for advancing the state of the art and for achieving the goal of creating AI systems that are both intelligent and interpretable. As the field evolves, it is expected that the integration of symbolic and sub-symbolic AI will play a key role in developing next-generation AI systems capable of addressing complex problems with unprecedented efficiency and effectiveness.

7. FROM ARTIFICIAL NEURAL NETWORKS (SUB-SYMBOLIC) TO RULES (SYMBOLIC)

Extracting rules from Artificial Neural Networks (ANNs) is crucial in demystifying these models' "black-box" nature, allowing their decisions to become understandable and interpretable to humans. This process involves translating the intricate, non-linear relationships learned by the network into a set of rules that humans can easily understand. To illustrate this process, we'll explore a detailed example of how rules can be extracted from an ANN trained on a simplified dataset for predicting loan approval based on applicant features.

7.1. Background

Let us use the example of a fictional financial institution that has created an ANN to evaluate loan applications. The ANN considers various applicant features such as Age, Income, Credit Score, and Employment Status and provides a binary decision: Approve or Deny. Despite the ANN's high accuracy, the decision-making process is not transparent. This makes it challenging for loan officers to explain decisions to applicants or to ensure compliance with regulations. The institution aims to derive understandable rules from the ANN to address this.

7.2. ANN Architecture

The ANN in this example is a simple feedforward network with one hidden layer. The input layer has four neurons corresponding to the applicant features. The hidden layer has a few neurons (say five for simplicity) using ReLU (Rectified Linear Unit) as the activation function [10]. The output layer has one neuron and uses a sigmoid activation function to output a probability of loan approval.

7.3. Rule Extraction Process

The rule extraction process involves several steps designed to translate the ANN's learned weights and biases into a set of if-then rules that replicate the network's decision-making process as closely as possible:

- Simplification: The first step involves simplifying the ANN to make the rule extraction more manageable. This could include pruning insignificant weights (shallow values) and neurons that have little impact on the output based on sensitivity analysis.
- Discretization: Since ANNs deal with continuous inputs and hidden layer activations, a discretization process is applied to convert these continuous values into categorical ranges. For instance, age might be categorized into 'Young', 'Middle-aged', and 'Old'; Income into 'Low', 'Medium', and 'High'; Credit Score into 'Poor', 'Fair', 'Good', and 'Excellent'; and Employment Status into 'Unemployed' and 'Employed'.
- Activation Pattern Analysis: Next, the activation patterns of the neurons in the hidden layer are analyzed for each input pattern. This involves feeding various combinations of the discretized input variables into the simplified network and observing which neurons in the hidden layer are activated for each combination. An activation threshold is defined to determine whether a neuron is considered activated.
- Rule Generation: Based on the activation patterns observed, rules are generated to replicate the ANN's decision process. Each rule corresponds to a path from the input layer through the activated hidden neurons to the output decision. For example:
 - If (Age is Young) and (Income is High) and (Credit Score is Good) and (Employment Status is Employed), then Approve Loan.
 - If (Age is Middle-aged) and (Credit Score is Poor), then Deny Loan. This step involves identifying which combinations of input features and hidden neuron activations lead to loan approval or denial, effectively translating the ANN's complex decision boundaries into more interpretable formats.
- Rule Refinement and Validation: The initial set of rules may be too complex or too numerous for practical use. Rule refinement techniques simplify and consolidate the rules without significantly reducing their accuracy in replicating the ANN's decisions. The refined rules are then validated against a test dataset to reflect the ANN's behavior accurately. This may involve adjusting the rules based on misclassifications or applying techniques to handle exceptions and edge cases.

After applying the rule extraction process to our hypothetical ANN, we might end up with a set of simplified, human-readable rules such as:

- Rule 1: If (Income is High) and (Credit Score is Excellent), then Approve Loan.
- Rule 2: If (Employment Status is Unemployed) and (Credit Score is Poor or Fair), then Deny Loan.
- Rule 3: If (Age is Old) and (Income is Low) and (Employment Status is Employed), then Deny Loan. These rules provide clear criteria derived from the ANN's learned patterns, making the decision-making process transparent and justifiable.

7.4. Advantages and Challenges

Some advantages include:

- Transparency: The extracted rules make the ANN's decisions transparent and understandable to humans.
- Compliance: Clear rules can help ensure compliance with regulatory requirements for explainable AI.
- Trust: Understanding how decisions are made can increase user trust in the AI system. Some

challenges are:

- Complexity: The rule extraction process can be complex, especially for deep or highly non-linear networks [18].
- Approximation: The extracted rules approximate the ANN's decision process and may not capture all nuances.
- Scalability: Extracting rules from large, deep neural networks with many inputs and hidden layers can be challenging and may result in many complex rules [19].

7.5. Summary

Extracting rules from ANNs allows AI decision-making to be transparent, understandable, and justifiable. Although there are challenges, especially with complex networks, this process is crucial for responsible and ethical AI use. By making AI systems more interpretable, we can establish trust with users, ensure compliance with regulations, and gain valuable insights into decision-making.

8. FUZZY COGNITIVE MAPS

The pendulum in AI is swinging back from purely statistical approaches toward integrating structured knowledge. FCMs are powerful cognitive tools for modeling and simulating complex systems. They blend elements from artificial neural networks, graph theory, and semantic nets to offer a unique approach to understanding and predicting system behavior. FCMs incorporate the concept of fuzziness from fuzzy logic, enabling them to handle ambiguity and uncertainty inherent in real-world scenarios. This extensive report delves into the origins of FCMs, provides illustrative case studies, and discusses their advantages and disadvantages, with references to their similarities to artificial neural networks, graphs, and semantic nets [20].

8.1. Origins

Bart Kosko introduced the concept of FCMs in the 1980s as an extension of cognitive maps. Cognitive maps, developed by Axelrod, were diagrams that represented beliefs and their interconnections. Kosko's introduction of fuzziness to these maps allowed for the representation of causal reasoning with degrees of truth rather than binary true/false values, thus capturing the uncertain and imprecise nature of human knowledge and decision-making processes. FCMs combine elements from fuzzy logic, introduced by Lotfi A. Zadeh, with the structure of cognitive maps to model complex systems.

8.2. Structure and Functionality

FCMs are graph-based representations where nodes represent concepts or entities within a system, and directed edges depict the causal relationships between these concepts. Each edge is assigned a weight that indicates the relationship's strength and direction (positive or negative). This structure closely mirrors that of artificial neural networks, particularly in how information flows through the network and how activation levels of concepts are updated based on the input they receive, akin to the weighted connections between neurons in neural networks [21].

However, unlike typical neural networks that learn from data through back propagation or other learning algorithms, the weights in FCMs are often determined by experts or derived from data using specific algorithms designed for FCMs. The concepts in FCMs can be activated like neurons, with their states updated based on fuzzy causal relations, allowing for dynamic modeling of system behavior over time. Integrating structured knowledge graphs with distributed

neural network representations offers a promising path to augmented intelligence. We get the flexible statistical power of neural networks that predict, classify, and generate based on patterns—combined with the formalized curated knowledge encoding facts, logic, and semantics via knowledge graphs [22].

8.3. The Inherent Reasoning Mechanism

The primary function of the reasoning rule in FCM models is to update the activation values of concepts iteratively, starting from initial conditions and continuing until a stopping criterion is satisfied. During each iteration, the reasoning rule utilizes three primary components to conduct these calculations: the weight matrix, which signifies the connections between concepts; the activation values of concepts from the previous iteration; and the activation function.

Eq.(1) shows a general rule commonly found in FCMs-related papers:

$$a_i^{(t)} = f \left(\sum_{j=1, i \neq j}^N a_j^{(t-1)} w_{ji} \right), \quad (1)$$

Recently, N'apoles et al. [23] proposed an updated quasi-nonlinear reasoning rule depicted in Eq. (2):

$$a_i^{(t)} = \underbrace{\phi \cdot f \left(\sum_{j=1}^N a_j^{(t-1)} w_{ji} \right)}_{\text{nonlinear component}} + \underbrace{(1 - \phi) \cdot a_i^{(0)}}_{\text{linear component}}, \quad (2)$$

such that $0 \leq \phi \leq 1$ is the nonlinearity coefficient. When $\phi = 1$, the concept's activation value depends on the activation values of connected concepts in the previous iteration. When $0 < \phi < 1$, we add a linear component to the reasoning rule devoted to preserving the initial activation values of concepts. When $\phi = 0$, the model narrows down to a linear regression where the initial activation values of concepts act as regressors. In their paper, N'apoles et al. [23] used the quasi-nonlinear reasoning rule to quantify implicit bias in pattern classification datasets. In contrast, the authors in [10] resorted to this rule to develop a recurrence-aware FCM-based classifier.

8.4. How Activation Functions Work

The activation function $f : \mathbb{R} \rightarrow I$ is an essential component in the reasoning rule of FCMbased models. This monotonically non-decreasing function keeps the activation value of each concept within the desired image set I , which can be discrete (a finite set) or continuous (a numeric-valued interval). It should be mentioned that I must be bounded; otherwise, the reasoning rule could explode due to the successive additions and multiplications when updating concepts' activation values during reasoning. Table 1 portrays relevant activation functions found in the literature.

Table 1. Some popular activation functions (both discrete and continuous).

Equation	Description
$f(x) = \begin{cases} 1, & x > 0 \\ 0, & x \leq 0 \end{cases}$	The bivalent activation function is a discrete activator to produce either 0 or 1. It allows modeling high-level qualitative states with concepts being either active or not. Moreover, $f(0) = 0$.
$f(x) = \begin{cases} 0, & x \leq v \\ 1, & x > v \end{cases}$	The step activation function is another discrete activator extending the bivalent activation function. It allows modeling qualitative states denoting whether a concept is activated or not, based on a threshold value v .
$f(x) = \begin{cases} -1, & x < 0 \\ 0, & x = 0 \\ 1, & x > 0 \end{cases}$	The trivalent activation function is another case of a discrete activator; it produces three states (e.g., -1, 0, 1). This approach is common when modeling qualitative states denoting whether a concept is negatively activated, inactive, or positively activated. In this function, $f(0) = 0$.
$f(x) = \tanh(x) = \frac{e^{2x}-1}{e^{2x}+1}$	The hyperbolic tangent function is a continuous activator that produces infinite states within the $[-1, 1]^N$ hypercube. This function can be used to model both qualitative and quantitative scenarios. Moreover, it holds that $f(0) = 0$.
$f(x) = \frac{1}{1+e^{-\lambda(x-h)}}$	The sigmoid activation function is another continuous activator that produces infinite states within the $[0, 1]^N$ hypercube. $\lambda > 0$ and $h \in \mathbb{R}$ are parameters controlling the function slope and offset. This function can be used for modeling both qualitative and quantitative scenarios.

8.5. Relevant Case Studies

For illustration purposes, Figure 1 shows an example of an FCM created to model a case of autism [24]. FCMs have been applied across various domains, demonstrating their versatility and effectiveness as a hybrid AI tool:

- Decision Support Systems: FCMs model complex decision-making processes, integrating expert knowledge and data-driven insights to support decisions in healthcare, environmental management, and business strategy.
- Predictive Modeling: In healthcare, FCMs model the progression of diseases or the impact of treatments, incorporating medical expertise and patient data to predict outcomes and support personalized medicine [25].
- System Analysis and Design: FCMs help analyze and design complex systems, such as socio-economic systems or ecosystems, by modeling the interactions between various factors and predicting the impact of changes or interventions.
- Healthcare Management: FCMs have been employed to model and predict patient outcomes in healthcare settings. For example, an FCM can be developed to understand the complex interplay between patient symptoms, treatment options, and possible outcomes, aiding medical professionals in decision-making [26].
- Environmental and Ecological Systems: In environmental studies, FCMs have been used to model the impact of human activities on ecosystems, allowing for the simulation of various scenarios based on different policies or interventions. This application showcases the strength of FCMs in handling systems where data may be scarce or imprecise [27].
- Business and Strategic Planning: FCMs assist in strategic planning and decision making within business contexts by modeling the relationships between market forces, company policies, and financial outcomes, offering a tool for scenario analysis and strategy development [28].

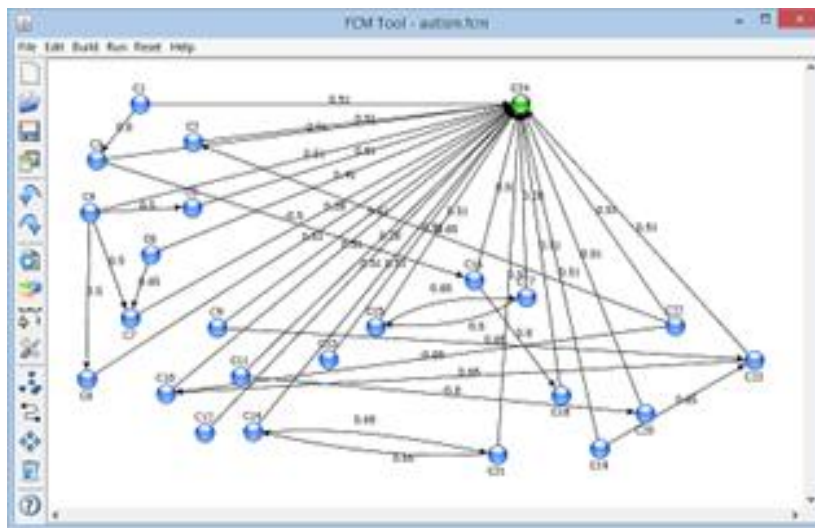


Fig. 1. Real example created with FCM Tool.

8.6. Advantages

The hybrid nature of FCMs offers several advantages:

- Interpretability and Transparency: The symbolic representation of concepts and causal relationships in FCMs provides clarity and understandability, facilitating communication with experts and stakeholders and supporting explainable AI.
- Flexibility and Adaptability: FCMs can be easily updated with new knowledge or data, allowing them to adapt to changing conditions or insights. This makes them particularly valuable in fields where knowledge evolves rapidly. – Handling of Uncertainty: Using fuzzy values to represent causal strengths enables FCMs to deal effectively with uncertainty and ambiguity, providing more nuanced and realistic modeling of complex systems [29].
- Integration of Expert Knowledge and Data-Driven Insights: FCMs uniquely combine expert domain knowledge with learning from data, bridging the gap between purely knowledge-driven and purely data-driven approaches.
- Interpretability: The graphical representation of FCMs, similar to semantic nets, allows for straightforward interpretation and understanding of the modeled system, making it accessible to experts and stakeholders without deep technical knowledge of AI.
- Flexibility: FCMs can incorporate quantitative and qualitative data, effectively handling uncertainty and imprecision through fuzzy logic. This flexibility makes them suitable for a wide range of applications.
- Dynamic Modeling Capability: FCMs can simulate the dynamic behavior of systems over time, providing valuable insights into potential future states based on different inputs or changes in the system [30].

8.7. Limitations

Despite their advantages, FCMs also face several challenges:

- Complexity with Large Maps: As the number of concepts and relationships in an FCM increases, the map can become complex and challenging to manage, analyze, and interpret [15].
- Learning and Optimization: While FCMs can learn from data, adjusting the fuzzy values of causal relationships can be computationally intensive and require sophisticated optimization techniques, especially for large and complex maps [31].
- Quantification of Expert Knowledge: Translating expert knowledge into precise fuzzy values for causal relationships can be challenging and may introduce subjectivity, requiring careful validation and sensitivity analysis [32].
- Subjectivity in Model Construction: The reliance on expert knowledge for constructing FCMs can introduce subjectivity, especially in determining the strength and direction of causal relationships between concepts.
- Complexity with Large Maps: As the number of concepts increases, the FCM can become complex and challenging to manage and interpret, potentially requiring sophisticated computational tools for simulation and analysis.
- Limited Learning Capability: While FCMs can be adjusted or trained based on data to some extent, they lack the deep learning capabilities of more advanced neural networks, which can autonomously learn complex patterns from large datasets [33].

8.8. References to ANNs, Graphs, and Semantic Nets

FCMs share several similarities with artificial neural networks, graphs, and semantic nets:

- Artificial Neural Networks: Like neural networks, FCMs consist of nodes (concepts) and weighted edges (causal relationships), where the state of each concept is updated based on the inputs it receives, akin to the activation of neurons. However, FCMs use fuzzy logic to handle the degrees of truth, whereas neural networks typically use continuous activation functions.
- Graphs: FCMs are directed graphs with weighted edges, employing graph theory concepts to represent and analyze the causal relationships between concepts. This graphical structure facilitates the visualization and analysis of complex systems [34].
- Semantic Nets: FCMs resemble semantic nets using nodes representing entities or concepts and edges representing relationships. However, FCMs focus on causal relationships and use fuzzy logic to capture the uncertainty and vagueness inherent in real-world systems [35].

8.9. FCMs as the Needed Hybrid Approach

There is momentum toward hybridizing connectionism and symbolic approaches to AI to unlock potential opportunities for an intelligent system to make decisions. This hybrid approach is gaining ground; FCMs embody a hybrid AI approach through their integration of symbolic and sub-symbolic elements:

- Symbolic Components: The concepts and causal connections in FCMs are symbolic, explicitly representing entities and their interrelations. This aligns with the symbolic AI paradigm, where knowledge is structured and interpretable, allowing for reasoning and inference based on explicit rules and relationships [36].

- Sub-symbolic Components: The strengths of the causal relationships in FCMs are represented by fuzzy values, which are learned and adjusted based on data or expert input, much like the weights in neural networks. This learning capability and the use of fuzzy logic to handle uncertainty and ambiguity mirror the characteristics of sub-symbolic AI, which learns from patterns in data without requiring explicit programming.

FCMs offer a compelling hybrid approach to AI, combining the symbolic representation of knowledge with sub-symbolic learning and reasoning; they bridge a crucial gap between symbolic AI's interpretability and structured knowledge representation and the adaptability and data-driven learning of sub-symbolic AI. Their applications across diverse domains underscore their versatility and potential to address complex problems by integrating human-like reasoning with machine learning [37]. The challenges FCMs face, including complexity management and the quantification of expert knowledge, highlight areas for further research and development. As AI continues to evolve towards more integrated and versatile models, FCMs stand as a testament to the potential of hybrid approaches to combine the strengths of symbolic and sub-symbolic AI, offering a pathway to more intelligent, understandable, and adaptable AI systems [38].

8.10. Summary

FCMs offer a robust framework for modeling and analyzing complex systems, blending the best symbolic and sub-symbolic AI by integrating fuzzy logic, graph theory, and neural network-like dynamics. While FCMs provide a powerful tool for understanding system behaviors and decision-making processes, their effectiveness is contingent upon accurately representing causal relationships and managing map complexity. Future developments in FCMs aim to enhance their learning capabilities, reduce subjectivity in their construction, and improve scalability, further solidifying their role as a valuable tool in complex system analysis and decision support across various domains [39].

9. CONCLUSIONS AND FUTURE WORK

Both symbolic AI and sub-symbolic AI are effective in solving complex problems. Symbolic AI best suits expert/knowledge systems that require human input and domain-specific knowledge. At the same time, sub-symbolic AI is ideal for applications requiring continuous learning, such as natural language processing, speech recognition, and image recognition tasks. Therefore, it is essential to consider the problem requirements and constraints before deciding which approach to use. The success of an AI application ultimately depends on selecting the appropriate strategy that best suits the requirements of the problem. The debate between symbolic AI and sub-symbolic AI is ongoing, with proponents on both sides. Symbolic AI proponents argue it is the only way to achieve accurate intelligence and understanding, as it relies on human-like cognitive processes such as reasoning and logic. They view sub-symbolic AI as limited in its ability to produce truly intelligent behavior, as it is primarily based on statistical algorithms and cannot reason abstractly. On the other hand, sub-symbolic AI enthusiasts argue that it offers a more flexible and powerful means of achieving intelligence. By mimicking how the brain processes information, it can better handle the complexity and variability of real-world situations. Moreover, it is less dependent on hand-coding and can learn from experience, achieving greater accuracy and adaptability over time. Ultimately, the debate between symbolic and sub-symbolic AI cannot be quickly resolved. Both approaches have their strengths and weaknesses, and the relative importance of each will depend on the specific application at hand. Nevertheless, understanding the debate and the merits of each approach can help drive progress toward the development of more advanced and effective AI systems.

It's essential to consider both symbolic and sub-symbolic approaches in AI development. Symbolic AI is helpful for logical reasoning and problem-solving, which can be expressed using rules and symbols. On the other hand, sub-symbolic AI excels in areas such as pattern recognition and learning from experience, like speech recognition, image recognition, and natural language processing. Combining both approaches can create hybrid models that are more effective in solving complex problems. For example, sub-symbolic AI can be used for feature extraction in image recognition, while symbolic AI can be used for classification. Understanding the strengths and limitations of each approach and applying them complementarily can lead to more comprehensive and intelligent AI systems that overcome the challenges and limitations of individual approaches. Both symbolic AI and sub-symbolic AI have their strengths and weaknesses, and they have different applications. Symbolic AI, with its rule-based system, works well when the problem-solving process requires many rules and is well-defined. On the other hand, sub-symbolic AI, which focuses on learning, is better suited for dealing with situations where the problem is not well-defined, and data can be used to generate new insights. Symbolic AI requires expert knowledge to create well-defined rules, while sub-symbolic AI only needs raw data to learn from. However, symbolic AI has a more deterministic and transparent approach, allowing developers to understand how the AI model reaches its conclusions. In comparison, subsymbolic AI is more of a black box, making it difficult to know how the model generates its results. Ultimately, the choice between symbolic AI and sub-symbolic AI depends on the specific application and the project's goals.

REFERENCES

- [1] G. N'apoles, F. Hoitsma, A. Knobon, A. Jastrzebska, and M. Leon, "Prolog-based agnostic explanation module for structured pattern classification," *Information Sciences*, vol. 622, p. 1196–1227, Apr. 2023.
- [2] W. Saeed and C. Omlin, "Explainable ai (xai): A systematic meta-survey of current challenges and future opportunities," *Knowledge-Based Systems*, vol. 263, p. 110273, 2023.
- [3] C. Borrero-Domínguez and T. Escobar-Rodríguez, "Decision support systems in crowdfunding: A fuzzy cognitive maps (fcm) approach," *Decision Support Systems*, vol. 173, p. 114000, 2023.
- [4] O. Emir and S,ule Onsel Ekici, "An integrated assessment of food waste model through intuitionistic fuzzy cognitive maps," *Journal of Cleaner Production*, vol. 418, p. 138061, 2023.
- [5] E. Struble, M. Leon, and E. Skordilis, "Intelligent prevention of ddos attacks using reinforcement learning and smart contracts," *The International FLAIRS Conference Proceedings*, vol. 37, May 2024.
- [6] G. N'apoles, M. L. Espinosa, I. Grau, K. Vanhoof, and R. Bello, *Fuzzy cognitive maps based models for pattern classification: Advances and challenges*, vol. 360, pp. 83–98. Springer Verlag, 2018.
- [7] M. Zhang, J. Wu, A. Belatreche, Z. Pan, X. Xie, Y. Chua, G. Li, H. Qu, and H. Li, "Supervised learning in spiking neural networks with synaptic delay-weight plasticity," *Neurocomputing*, vol. 409, pp. 103–118, 2020.
- [8] Y. Xia, J. Wang, Z. Zhang, D. Wei, and L. Yin, "Short-term pv power forecasting based on time series expansion and high-order fuzzy cognitive maps," *Applied Soft Computing*, vol. 135, p. 110037, 2023.
- [9] H. Wang, Y. Wu, Y. Liu, and W. Liu, "A cross-layer framework for lpwan management based on fuzzy cognitive maps with adaptive glowworm swarm optimization," in *2023 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1–6, 2023.
- [10] G. N'apoles, Y. Salgueiro, I. Grau, and M. Leon, "Recurrence-aware long-term cognitive network for explainable pattern classification," *IEEE Transactions on Cybernetics*, vol. 53, no. 10, pp. 6083 - 6094, 2023.
- [11] S. S Band, A. Yarahmadi, C.-C. Hsu, M. Biyari, M. Sookhak, R. Ameri, I. Dehzangi, A. T. Chronopoulos, and H.-W. Liang, "Application of explainable artificial intelligence in medical health: A systematic review of interpretability methods," *Informatics in Medicine Unlocked*, vol. 40, p. 101286, 2023.
- [12] J. Wang, Z. Peng, X. Wang, C. Li, and J. Wu, "Deep fuzzy cognitive maps for interpretable

- multivariate time series prediction,” *IEEE Transactions on Fuzzy Systems*, vol. 29, no. 9, pp. 2647–2660, 2021.
- [13] H. DeSimone and M. Leon, “Explainable ai: The quest for transparency in business and beyond,” in 2024 7th International Conference on Information and Computer Technologies (ICICT), IEEE, Mar.2024.
- [14] A. B. Arrieta, N. D’íaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. GilLópez, D. Molina, R. Benjamins, et al., “Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai,” *Information fusion*, vol. 58, pp. 82–115, 2020.
- [15] R. Schuerkamp and P. J. Giabbanelli, “Extensions of fuzzy cognitive maps: A systematic review,” *ACM Comput. Surv.*, vol. 56, no. 2, 2023.
- [16] Y. Xixi, D. Fengqian, and L. Chao, “Time series prediction based on high-order intuitionistic fuzzy cognitive maps with variational mode decomposition,” *Soft Computing*, vol. 26, pp. 189–201, 2022.
- [17] M. Wrzesień, M. Wrzesień, and W. Homenda, “Time series processing with cognitive maps. the case of general forecast modeling for time series of similar nature,” in 2022 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), pp. 1–8, 2022.
- [18] Z. Yang, J. Liu, and K. Wu, “Learning of boosting fuzzy cognitive maps using a real-coded genetic algorithm,” in 2019 IEEE Congress on Evolutionary Computation (CEC), pp. 966–973, 2019.
- [19] W. Liang, Y. Zhang, X. Liu, H. Yin, J. Wang, and Y. Yang, “Towards improved multifactorial particle swarm optimization learning of fuzzy cognitive maps: A case study on air quality prediction,” *Applied Soft Computing*, vol. 130, p. 109708, 2022.
- [20] Y. Hu, Y. Guo, and R. Fu, “A novel wind speed forecasting combined model using variational mode decomposition, sparse auto-encoder and optimized fuzzy cognitive mapping network,” *Energy*, vol. 278, p. 127926, 2023.
- [21] W. Hoyos, J. Aguilar, and M. Toro, “A clinical decision-support system for dengue based on fuzzy cognitive maps,” *Health Care Management Science*, vol. 25, no. 4, pp. 666–681, 2022.
- [22] W. Hoyos, J. Aguilar, and M. Toro, “Prv-fcm: An extension of fuzzy cognitive maps for prescriptive modeling,” *Expert Systems with Applications*, vol. 231, p. 120729, 2023.
- [23] G. N’apoles, I. Grau, L. Concepci’ón, L. Koutsoviti Koumeri, and J. P. Papa, “Modeling implicit bias with fuzzy cognitive maps,” *Neurocomputing*, vol. 481, pp. 33–45, 2022.
- [24] M. Leon Espinosa and G. Napoles Ruiz, “Modeling and experimentation framework for fuzzy cognitive maps,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, Mar. 2016.
- [25] K. Poczeta and E. I. Papageorgiou, “Energy use forecasting with the use of a nested structure based on fuzzy cognitive maps and artificial neural networks,” *Energies*, vol. 15, no. 20, 2022.
- [26] G. D. Karatzinis, N. A. Apostolikas, Y. S. Boutalis, and G. A. Papakostas, “Fuzzy cognitive networks in diverse applications using hybrid representative structures,” *International Journal of Fuzzy Systems*, vol. 25, no. 7, pp. 2534–2554, 2023.
- [27] O. Orang, P. C. de Lima e Silva, and F. G. Guimarães, “Time series forecasting using fuzzy cognitive maps: a survey,” *Artificial Intelligence Review*, vol. 56, p. 7733–7794, 2023.
- [28] M. Leon, “Business technology and innovation through problem-based learning,” in Canada International Conference on Education (CICE-2023) and World Congress on Education (WCE-2023), CICE2023, Infonomics Society, July 2023.
- [29] M. Leon, “Aggregating procedure for fuzzy cognitive maps,” *The International FLAIRS Conference Proceedings*, vol. 36, no. 1, 2023.
- [30] E. Jiya, O. Georgina, and A. O., “A review of fuzzy cognitive maps extensions and learning,” *Journal of Information Systems and Informatics*, vol. 5, no. 1, pp. 300–323, 2023.
- [31] S. Yang and J. Liu, “Time-series forecasting based on high-order fuzzy cognitive maps and wavelet transform,” *IEEE Transactions on Fuzzy Systems*, vol. 26, no. 6, pp. 3391–3402, 2018.
- [32] T. Koutsellis, G. Xexakis, K. Koasidis, N. Frilingou, A. Karamaneas, A. Nikas, and H. Doukas, “Incognitive: A web-based python application for fuzzy cognitive map design, simulation, and uncertainty analysis based on the monte carlo method,” *SoftwareX*, vol. 23, 2023.
- [33] D. Qin, Z. Peng, and L. Wu, “Deep attention fuzzy cognitive maps for interpretable multivariate time series prediction,” *Knowledge-Based Systems*, vol. 275, p. 110700, 2023.
- [34] B. Qiao, J. Liu, P. Wu, and Y. Teng, “Wind power forecasting based on variational mode decomposition and high-order fuzzy cognitive maps,” *Applied Soft Computing*, vol. 129, p. 109586, 2022.
- [35] Z. Xuan and Z. Shengen, “Research of task complexity decision system for manual assembly tasks

- using fuzzy cognitive maps to construct bayesian networks,” in Volume 9: Mechanics of Solids, Structures, and Fluids; Micro- and Nano-Systems Engineering and Packaging; Safety Engineering, Risk, and Reliability Analysis; Research Posters, American Society of Mechanical Engineers, 2022.
- [36] M. Leon, L. Mkrtchyan, B. Depaire, D. Ruan, and K. Vanhoof, “Learning and clustering of fuzzy cognitive maps for travel behaviour analysis,” *Knowledge and Information Systems*, vol. 39, no. 2, pp. 435–462, 2013.
- [37] M. Leon, “Fuzzy cognitive maps as a bridge between symbolic and sub-symbolic artificial intelligence,” *International Journal on Cybernetics & Informatics (IJCI)*. 3rd International Conference on Artificial Intelligence Advances (AIAD 2024), vol. 13, pp. 57–75, Aug. 2024.
- [38] G. N’apoles, J. L. Salmeron, W. Froelich, R. Falcon, M. Leon, F. Vanhoenshoven, R. Bello, and K. Vanhoof, *Fuzzy Cognitive Modeling: Theoretical and Practical Considerations*, p. 77–87. Springer Singapore, July 2019.
- [39] G. N’apoles, M. Leon, I. Grau, and K. Vanhoof, “FCM expert: Software tool for scenario analysis and pattern classification based on fuzzy cognitive maps,” *International Journal on Artificial Intelligence Tools*, vol. 27, no. 07, p. 1860010, 2018.

AUTHOR

Dr. Maikel Leon is interested in applying AI/ML techniques to modeling real-world problems using knowledge engineering, representation, and data mining methods. His recent research focuses on XAI and has been featured in high-impact factor journals, such as *Information Sciences* and *IEEE Transactions on Cybernetics*. Dr. Leon is a reviewer for the *International Journal of Knowledge and Information Systems*, *Journal of Experimental and Theoretical Artificial Intelligence*, *Soft Computing*, and *IEEE Transactions on Fuzzy Systems*. He is a Committee Member of the Florida Artificial Intelligence Research Society. He is a frequent contributor on technology topics for CNN en Espa˜nol TV and the winner of the Cuban Academy of Sciences National Award for the Most Relevant Research in Computer Science. Dr. Leon obtained his PhD in Computer Science at Hasselt University, Belgium, previously having studied computation (Master of Science and Bachelor of Science) at Central University of Las Villas, Cuba.