Proceedings of the 6th European Conference on Industrial Engineering and Operations Management
Lisbon, Portugal, July 18-20, 2023

Towards a Neuro-Symbolic Approach to Bridge the Gap Between Brain and Mind-Inspired Models

Fabio Caversan, Lucas Rodrigues, Geanderson Ferreira,
Danilo De Almeida, Carlos Veber, Gabriel Dias Rezende Martins
Stefanini Group
SAI Research and Development Labs
27100 W 11 Mile Rd,
Southfield, MI 48034, USA
Fabio.Caversan@stefanini.com

Dr. Nabih Jaber and Dr. George Pappas
Department of Electrical and Computer Engineering
Lawrence Technological University
Southfield, MI 48075, USA
njaber@ltu.edu, gpappas@ltu.edu

Abstract

This paper proposes an approach beyond the current Large Language Models (LLMs) milestones. To achieve optimal results, LLMs require careful fine-tuning. Using techniques such as prompt engineering and In Context Language (ICL), it is possible to provide LLMs with the necessary guidance to perform specific tasks with incredible accuracy and relevance. Deep Neural Networks have limited capabilities in intelligent behavior (i.e., in understanding the meaning of its input). Hence, the scientific community started to question more heavily in recent years whether Deep Learning alone could get us closer to Artificial General Intelligence (AGI). The authors of this paper believe that the training and adaptive behavior challenges were a big part of the problem and that symbolic techniques did not advance simultaneously as Deep Neural Networks. Hence, revisiting and renewing the symbolic techniques and combining them with Deep Learning models like the LLMs can bridge the gap between neural techniques and symbolic knowledge representation. Our proposed system has been shown to significantly improve the performance of LLMs in terms of hallucination rate while providing flexibility of semantic input processing to symbolic systems.

Keywords
BERT, GPT-3, Semantic-Neuron Automata, NPC Programming, Large Language Models (LLMs)

1. Introduction

In this paper, we will describe our challenges and developments when navigating a new wave of AI research where new, state-of-the-art Large Language Models (LLMs) models are opening a multitude of paths for exploration and demonstrate with a practical application how to apply prompt engineering as a way to make connectionist models perform symbolic reasoning. Before going deeper into this study, it is important to note what many authors describe as the main paradigms of Artificial Intelligence (AI): Symbolism and Connectionism. In the article (Zhang et al. 2022), it is explained how since the birth of AI in 1956 there were times when symbolic models were dominant, followed by impressive connectionist advancements, while also noting none are, by themselves, complete to get closer to the so-called Artificial General Intelligence (AGI). Many authors can find deep studies on the competing paradigms, and AI professionals have explored their limitations since the beginning. In his book, (Dyer et al. 1991) describes his journey exploring both, explaining his challenges with the lack of dynamic memory management or virtual pointers in connectionist models, and his discontent with how knowledge engineering heavy the symbolic systems might become, for example. From another perspective, for the past decades, cognitive science has been exploring the relations between brain and mind, where you can relate the human brain to Connectionism, to Neural

© IEOM Society International

Computing, and the human mind to Symbolism, Expert Systems, Fuzzy Reasoning, among other symbolic techniques, as described by Churchland and Churchland.

Recently, a study presented by (Bubeck et al. 2023) argues that the latest model, GPT-4 (OpenAI), presents "sparks" of AGI. Regardless of the discussion around general intelligence, it is proved that the LLM model performs remarkably well in several tasks, and the limitations are highlighted. We hope to propose an alternative to such limitations in the form of hybrid models.

Hybrid models merging the best of both have been defended academically and used in practical ways successfully for the past two decades as attempts to extract the best of both worlds, as a form of mitigation of each paradigm's limitations and to create more complete and intelligent systems. That is the inspiration for our line of work in this paper.

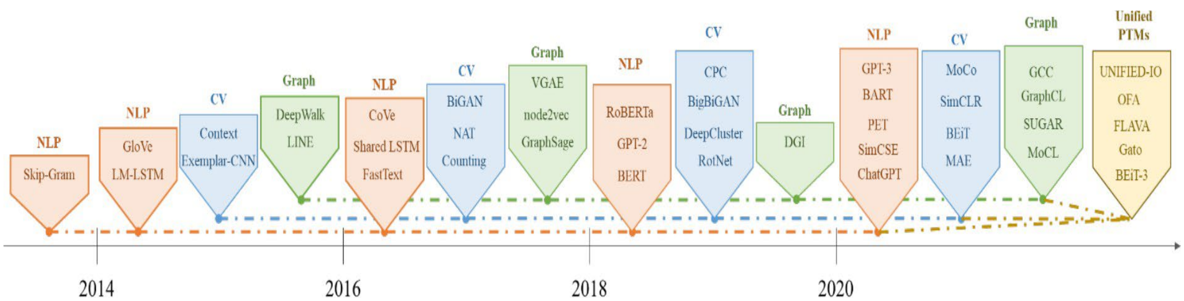

Figure 1: History and Evolution of Pre-trained Models according to Zhou et al.

2. Background - From Perceptrons to Large Language Models and Beyond

The connectionist school has been advancing quickly in the last years; since the Deep Learning was introduced to the field at the beginning of this century, many new models are being proposed every year. (Zhou et al. 2023) present a comprehensive survey of the last part of this history that is of interest to this work: the latest developments that resulted in LLMs like ChatGPT.

The Large Language Models are a particular case of Pretrained Models, and a summarized version of the evolution of such models can be found in Figure 1.

In this quick revision, we want to focus on a few of the major milestones of this history: transformers, BERT, and GPT-3.

2.1 Transformers

The transformer proposed by (Vaswani 2017) was a novel neural network architecture used in natural language processing (NLP), computer vision (CV) Pappas (2023), and graph learning (GL) that relies solely on attention mechanisms, without using recurrence or convolution. The attention mechanism assigns weights to input representations and calculates the most important part of the input data. Transformers use a mask matrix for self-attention to determine which words can "see" each other. They help solve long-range dependency issues in processing sequential input data and are scalable. The largest language models have over 100B parameters, thanks to transformer structures achieving higher parallelization. The architecture quickly replaced the former Long Short-Term Memory (LSTM) models as the main choice for Natural Language Processing (NLP) tasks, which was applied to creating several novel models like BERT.

2.2 BERT

BERT (Bidirectional Encoder Representations from Transformers) was a novel pre-training method for natural language processing tasks introduced by (Devlin et al. 2018) BERT is a deep bidirectional transformer model that is pre-trained on large amounts of unlabeled text data.

One of the key features that set BERT apart from other pre-training methods is its bidirectional nature. Traditional language models are unidirectional, meaning they can only consider the preceding words in a sentence when predicting the next word. In contrast, BERT is able to consider both the preceding and following words, as well as the entire sentence, when making its predictions. This allows BERT to capture more nuanced relationships between words and better understand the context in which words are used. The bidirectional nature of BERT is made possible by its use of transformer architecture, which enables efficient processing of long input sequences.

BERT achieved state-of-the-art results on a variety of NLP benchmarks, including the GLUE benchmark, which measures performance on multiple NLP tasks, and the SQuAD benchmark, which focuses on question answering. Meanwhile, the GPT models were being trained and would eventually land in our next milestone: GPT-3.

2.3 GPT-3

GPT-3, or Generative Pre-trained Transformer 3, was introduced by (Brown et al. 2005) as a state-of-the-art natural language processing (NLP) model developed by OpenAI in 2020. GPT-3 was built upon its predecessor, GPT-2, with a significantly larger architecture and a massive amount of pre-training data.

GPT-3 has 175 billion parameters, the largest language model at that date. Its size enabled it to perform a wide range of NLP tasks, including text generation, summarization, question answering, and language translation, among others. GPT-3 was pre-trained using unsupervised learning on a diverse corpus of web documents, books, and other text sources. This pre-training helped the model to learn language patterns and structures, which it can then use to perform various NLP tasks.

One of the most significant achievements of GPT-3 is its few-shot learning capability. This means the model can learn to perform a task from just a few examples, even if it has never encountered similar examples. For example, given a few examples of how to summarize news articles, GPT-3 can generate high-quality summaries for any news article it encounters. Those examples are presented as part of the input, so now weight updates happen in the neural network itself. This particular characteristic is key for the whole prompt engineering area, the focus of Section 3.4, and key for our proposition of symbolic encoding on the top of neural systems.

In addition to few-shot learning, GPT-3 also has presented zero-shot learning abilities. This means that the model could perform tasks not explicitly trained. For instance, it can translate text from one language to another, even if it has never seen any examples of that particular language pair.

As we write this article, GPT is already in version 3.5, and version 4 has already been announced. This article was written with the aid of version 3.5, but we believe that the concepts presented here will apply to the next versions of GPT and other LLM versions to come.

2.4 Prompt Engineering

In-context learning (ICL) is a powerful technique in natural language processing (NLP) that allows large language models (LLMs) to learn from specific contexts or prompts rather than relying solely on implicit information. (Wei et al. 2023) ICL can significantly improve the quality of text generated by LLMs, making them more accurate and relevant to the task at hand. With ICL, it is possible to embed the description of one task in the input to an LLM, such as a question or a specific context, and use the defined information as boundaries for the reasoning.

Prompt engineering is a sample of ICL that involves designing prompts or contexts that can be used to fine-tune LLMs for specific tasks. Prompt engineering typically involves converting one or more tasks to a prompt-based dataset and training a language model with what has been called" prompt-based learning" or" prompt learning."

The idea is to embed the description of the task in the input, such as a question or a specific context, rather than relying on implicit information. This approach can significantly improve the performance of LLMs on a wide range of NLP tasks, including text classification, question answering, and text generation.

The broad accessibility of these tools has been driven by the increase in publications about team and community-led projects, making prompt engineering an increasingly popular approach in the field of NLP.

The comparison between LLMs and toddlers may seem odd initially, but it highlights an important point: just as toddlers require careful instruction to learn and develop, LLMs require careful fine-tuning to achieve optimal results. The power of LLMs lies in their ability to learn from vast amounts of data, but without proper instruction, they can produce inaccurate, untruth, and irrelevant results. By using techniques such as prompt engineering and in-context learning, it is possible to provide LLMs with the necessary guidance to perform specific tasks with incredible accuracy and relevance.

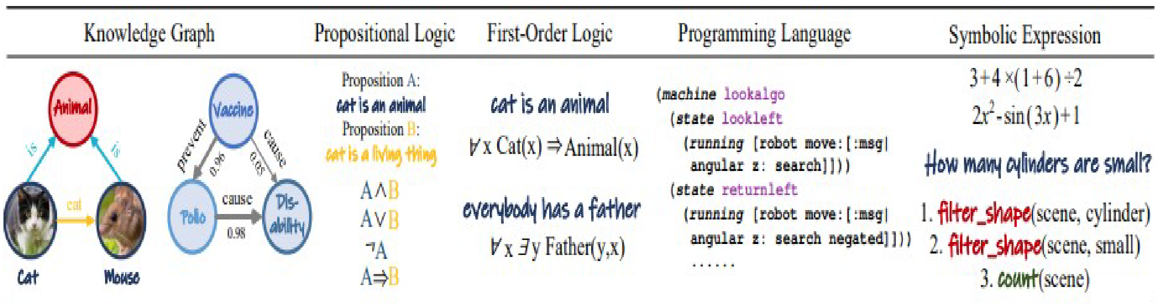

Figure 2: Knowledge representation models from (Wang et al. 2022).

2.5 Beyond LLMs - New Wave of AI Research

As discussed in the introduction in the first section, it has been our belief since the early stages of Deep Neural Networks that Connectionism alone would have limited capabilities in terms of intelligent behavior, particularly in understanding the meaning of its inputs. At the same time, we cannot deny that symbolic techniques did not advance at the same pace, and we believe the challenges with training and adaptive behavior were a big part of the problem.

Even with the remarkable advances that Deep Learning techniques introduced, as mentioned in the previous section, the scientific community started to question more heavily in the last years if Deep Learning could by itself get us closer to AGI. We understand that a fair portion of the scientific community believes that we can achieve the next level with more and better data that the quality we need will emerge from quantity itself. The authors of this article respectfully disagree with that point of view but also do not want to simply ignore all the advances we had with neural networks in the last years.

We believe it is time to revisit and renew the symbolic techniques, combining them with state-of-the-art Deep Learning models like the LLMs to achieve a new level of AI, sometimes mentioned as third-wave or third-generation AI. And there is plenty of research in the area lately: In (Odense and Garcez et al. 2022), we can find a proposal of neural network architecture to encode symbols and relations, named semantic encodings, which could be used to establish a relationship between symbolic logical constraints and a trained neural network. Along those lines, (Zhang et al. 2023) proposes that the symbolic solution was the first-generation AI, having deep learning as the second-generation AI. Finally, it is proposed that a third generation would have to combine both approaches and even proposed a triple-space model to combine both techniques. The creation of new neural architectures to bridge the gap to Symbolism is a promising field of research and should be explored, but we want to evaluate a higher-level approach, leveraging, in a way, the pre-trained models that are available and taking advantage of transfer learning techniques.

From that higher-level perspective, shifting from the most basic neural networks architectures to pre-trained LLMs described in the previous section, we also find a lot of recent interest in leveraging LLMs in typical symbolic applications as presented in (Wei and Zhang et al. 2023), different ways to apply prompt engineering using LLMs to perform symbolic reasoning. Those are closely aligned with the approach we want to propose in this article. The idea is to leverage LLMs with prompt engineering to represent knowledge using symbolic structures. In another work that also tries to leverage the best of both worlds, (Wang et al. 2022) describe five categories of knowledge representation: knowledge graph, propositional logic, first-order logic, programming language, and symbolic expression, presented in Figure 2. The term NeSy is used to refer to Neuro-Symbolic computing. We believe the ideas presented here can help to bridge the gap between neural techniques and symbolic knowledge representation. We start the discussion with a very simple representation well known in classic logic, a Finite State Machine, and explore how that would look if we replace the underlying computing mechanism with LLMs.

3 New Symbolic Computation Based on Neural Techniques

All the advances described in Section 3 enabled a large number of applications for LLMs, especially with the latest GPT models. The main challenge, however, is still the inherent problem of hallucination, alongside prompt injection attacks that are also related to the freedom the models have to generate output.

Suppose we rely on these models to retrieve information about a specific domain or question. In that case, the model will eventually combine related facts and generate an answer that is not necessarily correct, especially if the answer is not presented in the context passed to the model. Our proposal is not to use those models to retrieve information, but to leverage the language understanding they acquired in the training process with the immense volume of corpus the learning process uses.

Hence our proposal: to combine the potential to understand human language acquired by the latest LLM models with the symbolic reasoning that will not allow the system to generate invalid output: in a primitive way, the system is thinking and able to handle imprecision and uncertainty in its input as we humans do. Using the metaphor presented in Section 2, LLMs would be the brain and the Semantic Automaton, the mind. First, let's quickly define the classical model of Finite Automata.

3.1 Formal Languages and Automata

Formal languages are a fundamental concept in computer science and mathematics that allow us to precisely describe the structure and behavior of complex systems. A formal language is a set of strings of symbols that are constructed according to a set of rules or grammar. These strings of symbols can represent anything from natural language sentences to computer programs, and they can be analyzed using a variety of mathematical tools. (Hopcroft et al. 2006). One of the key tools for analyzing formal languages is a finite automaton (FA), which is a mathematical model of a simple computing device that can recognize strings in a language. A FA consists of a set of states, a set of input symbols, a transition function, and a set of accepting states. The transition function determines the next state of the automaton based on the current state and the input symbol. After all the inputs are processed, if the state of the automaton is a valid final state, the string is accepted. Finite automata are powerful tools for analyzing formal languages because they can be used to determine whether a string belongs to a language in a very efficient manner.

A finite state automaton (FSA) can be defined as a 5-tuple (Q, Σ, δ, q_0, F), where:

Q is a finite set of states.

Σ is a finite alphabet of input symbols.

$\delta: Q \times \Sigma \rightarrow Q$ is the transition function that maps each state and input symbol to a new state.

$q_0 \in Q$ is the initial state.

$F \subseteq Q$ is a set of accepting (or final) states.

Finite automata have many practical applications in computer science and engineering. For example, they can be used to model the behavior of digital circuits, verify the correctness of software programs, and recognize patterns in natural language text. Finite automata can also be used in the design of compilers and interpreters for programming languages and in the construction of regular expressions for text searching and matching.

In addition to finite automata, there are many other important concepts and tools in the field of formal languages and automata theory, including context-free grammars, pushdown automata, Turing machines, and computational complexity theory. These tools are essential for understanding computer science's theoretical foundations and developing new algorithms and software systems. The idea to use a finite automaton as the initial candidate for our classic-to-semantic symbolic approach is precisely because this work can be extended to the entire formal language field.

© IEOM Society International

3.2 SAuroN: Semantic-Neuron Automata

Derived from the typical State Automata, we define a Semantic-Neuron Automata as a 5-tuple (Q, σ, λ, q_0, F), where:

- Q is a finite set of states.
- σ is an input sentence typically representing a natural language statement.
- λ: Q × σ → Q is the transition function represented by a set of semantic rules that map each state and input symbol to a new state: in this case, the functions are computed by an LLM.
- q_0∈ Q is the initial state.
- F ⊆ Q is a set of accepting (or final) states.

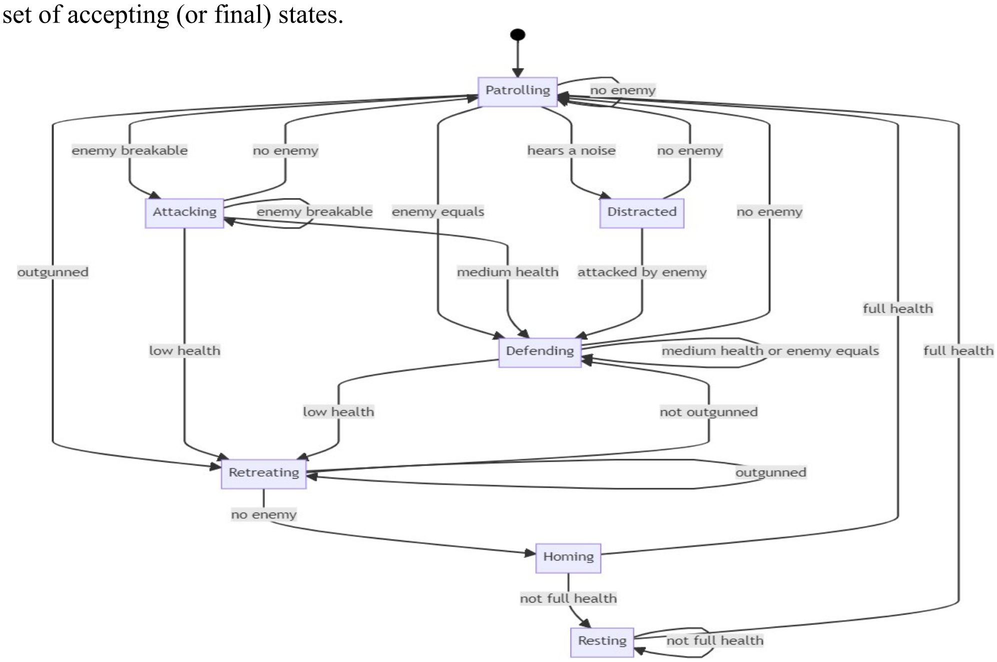

Figure 3: NPC State Machine Diagram

In a SAuroN, the transition of states allows uncertainty, imprecision, and follow meaning (semantics) instead of precision. Let us explore a simple scenario to see such a structure in action.

4.3 SAuroN Case Study: Intelligent NPC Programming

To illustrate the differences between typical State Automata and SAuroN, we will use a simple example of a state machine: the modeling of a Non-Player Character (NPC) of a game. The state machine that governs the behavior of a given NPC is depicted in Figure 3.

To make the implementation of this state machine using a classical approach, certain thresholds and assumptions will have to be defined. The algorithm can assume, for instance, that any number below 50 is low health and that 100 is the maximum health. It can be defined that "enemy equals" is the exact same number of defenders and enemies and so on. Concepts like "outgunned" can be a little harder to define, but it can certainly be done.

In a SAuroN model, the transitions follow semantic meaning, meaning there are possible infinite combinations to transition from one state or another, as long as they hold the same meaning. In terms of the operation of the model, at each iteration, the state machine feeds the current state's prompt with a textual description of the environment, and

© IEOM Society International

parses from the LLM's output which is the new state to be used in the next iteration. In the following Table 1, we present the result of an experiment using a SAuroN model to control the NPC behavior based on events that are not precisely defined:

Iteration	Current State	Input	New State
1	Patrolling	there's no sign of enemy	Patrolling
2	Patrolling	you hear something in a bush	Distracted
3	Distracted	it was a squirrel	Patrolling
4	Patrolling	you see an enemy	Attacking
5	Attacking	you took a little bit of damage	Defending
6	Defending	you defeated the enemy	Patrolling
7	Patrolling	you see another enemy	Attacking
8	Attacking	you are still fighting	Attacking
9	Attacking	you took a lot of damage	Retreating
10	Retreating	you escaped from the enemy pursuit	Homing
11	Homing	you are still wounded	Resting
12	Resting	you are not fully healed, yet	Resting
13	Resting	you are completely healed now	Patrolling

Table 1: Result of an experiment using a SAuroN model

The states have a separate and independent prompt to process their transitions.

As an example, the "Attacking" state transition function is presented:

You are going to emulate a finite state machine FSA.

Your response to an input must be limited to exclusively the new state name. Consider only the following as possible values for the new state:

If your health is medium, new state is "Defending." If your health is low, new state is "Retreating."

If there is no enemy, new state is "Patrolling." Otherwise, new state is "Attacking".

Input: {input}

New state:

For this particular model, there is no room or degree of freedom in terms of output: if the output contains a valid state, the transition is processed. In any other case, the state remains the same. The experiments with prompt engineering resulted in a few interesting cases that might be explored further, like the creation of new states (in one test, the state "Healing" was created as the output of a transition, for example).

SAuroN models and, eventually, other derived models to operate symbolic reasoning over neural nets are interesting alternatives to avoid current challenges with LLMs like hallucination and injection attacks. At the same time, it enables the modeling of goal-based agents, semantic networks, reasoning systems, and other symbolic approaches using LLMs.

© IEOM Society International

5. Conclusion

In this paper, we presented a model combining connectionist and symbolic techniques, leveraging in particular the latest advances in the neural computing field, the LLMs. Although the SAuroN model proposed here is a simple example of leveraging state-of-the-art LLMs to perform symbolic processing, this can potentially be applied to create other forms of knowledge representation, such as knowledge graphs, propositional logic, and semantic networks, among others. We believe that with this approach, the best of both worlds can be achieved: the plasticity of neural networks, the language understanding capabilities of the latest LLM models, and the knowledge representation and manipulation of the symbolic models.

Naturally, this is a simple first step in that direction, and there is much room to explore and improve. SAuroN introduces the possibility of transitions between states to be semantically triggered, and the symbolic modeling avoids hallucination and injection attacks, current challenges for LLM models. It could potentially help the limitations described in (Bubeck et al. 2023) around weak symbolic manipulation and basic mathematical mistakes.

In terms of size limitations, LLMs have a limited context, and as a system becomes more and more complex, the size of the prompts used for few shot learning might become a limitation. As we use SAuroN or models derived from its idea, the symbolic layer might be used to select the proper prompt, narrowing the context for that particular moment of the operation and removing that limitation. A few ideas for further exploration and development are:

- How to transfer the knowledge acquired from the operation of the symbolic layer to the neural layer
- The model can handle inputs with great flexibility and plasticity, but the outputs would still follow traditional logic rules. To make outputs flexible, we need to give more room to the generative AI, and the problems of hallucination and injection would be back.
- As we create more sophisticated symbolic models like knowledge graphs, new ideas will be necessary to search, extract, and present knowledge

As discussed in this paper, the latest advances in the neural computing field, the LLMs, are sparking the debate about hybrid techniques and different approaches to reach a superior quality AI system. At the same time, these new models are being used to create exciting novel applications in different areas and industries. The idea of combining neural and symbolic techniques is not new, and we believe as one field advances, the other should be revisited to evaluate the potential of the combination of both.

We hope this contribution helps in the pursuit of pieces of the ancient puzzle, that is, to understand the relations between the brain and the mind, and take us one step closer to recreating those impressive structures on artificial systems.

© IEOM Society International

References

Zhang, B. and Jun Zhu, and Hang Su. Toward the third-generation artificial intelligence. Science China Information Sciences, 66(2):121101, Jan 2023. ISSN 1869-1919.

Dyer, G. M, Connectionism Versus Symbolism in High-Level Cognition. Springer Netherlands, Dordrecht, 1991. ISBN 978-94-011-3524-5.

Churchland, S. P., Neurophilosophy: Toward a Unified Science of the Mind- Brain. The MIT Press, 09 1989. ISBN 9780262315944.

Churchland, P., Matter and consciousness, third edition. Cambridge, Massachusetts, MIT Press, 01 2013.

Bubeck, S´ebastien, and Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4, 2023.

OpenAI. Gpt-4 technical report, 2023.

Zhou, C. and Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, Hao, P. and Jianxin, Li, Jia Wu, Ziwei Liu, Pengtao Xie, Caiming Xiong, Jian Pei, Philip S. Yu, and Lichao Sun. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt, 2023. URL https://arxiv.org/abs/2302.09419.

Vaswani, A. and Noam, Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. At- tension is all you need. CoRR, abs/1706.03762, 2017. URL http://arxiv.org/abs/1706.03762.

Devlin, J. and Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. CoRR, abs/1810.04805, 2018. URL http://arxiv.org/abs/1810.04805.

Brown, T. and Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. CoRR, abs/2005.14165, 2020. URL https://arxiv.org/abs/2005.14165

Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, and Tengyu Ma. Larger language models do in-context learning differently, 2023. URL https://arxiv.org/abs/2303.03846

Odense, S and Garcez, A.A. A semantic framework for neural- symbolic computing, 2022. URL https://arxiv.org/abs/2212.12050

Wenguan Wang, Yi Yang, and Fei Wu. Towards data-and knowledge-driven artificial intelligence: A survey on neuro-symbolic computing, 2022. URL https://arxiv.org/abs/2210.15889

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. CoRR, abs/2201.11903, 2022. URL https://arxiv.org/abs/2201.11903

Zhang, Hanlin, and Yi-Fan Zhang, Li Erran Li, and Eric Xing. The impact of symbolic representations on in-context learning for few-shot reasoning, 2022. URL https://arxiv.org/abs/2212.08686

John E. and Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. Introduction to Automata Theory, Languages, and Computation. Addison-Wesley Longman Publishing Co., Inc., USA, 2006. ISBN 0321455363.

Linz P. An Introduction to Formal Languages and Automata. Jones & Bartlett Learning, 6th edition, 2016.

Ilovan, P., and Pappas G. Improve Prediction Accuracy and Classification of Bone Using Machine Learning with Hounsfield Unit, ICPR 2022.

Biographies

Fábio Caversan is a Digital Business and Innovation Vice President at Stefanini. Fabio has a degree in Computer Engineering from the Faculdade de Engenharia de Sorocaba (FACENS - 2002) and a master's degree in Electrical Engineering from the Polytechnic School of the University of São Paulo (2006). Professor of Computer Engineering at FACENS for 15 years and co-founder of WOOPI Technology (since 2012 a company of the Stefanini Group), creator of the Sophie Platform, now holds the position of Artificial Intelligence Research & Development Manager of Stefanini North America, leading the company efforts in cognitive sciences and AI. He is one of the main collaborators/contributors of a current ISP project between Stefanini and LTU.

Dr. Nabih Jaber is an Associate Professor and Chair in the Department of Electrical and Computer Engineering (ECE). Dr. Jaber is currently the Director of the Innovative Smart Wireless Networking Lab (ISWiNLab). His research interests focus on Intelligent Transportation Systems (ITS); Artificial Intelligence (AI) as it pertains to Advanced Driver Assistance System/Autonomous Vehicles (ADAS/AV), which include research in object detection/tracking algorithms (ODTA), telemetry, remote communication, autonomous platform implementation, and controller area network (CAN); and Dedicated Short Range Communications (DSRC) vehicular systems. Dr. Jaber's most recent principal investigator (PI) roles were in the field of Artificial Intelligence (AI), "Tag System Real-Time Tracking," and "Minimal Viable Product (MVP) Contextualizer." Dr. Jaber's previous work with Hyundai America Technical Center INC. (HATCI) as the PI involved collecting and analyzing speech data for hands-free speech enhancement using real automobile data. The project resulted in seven published peer reviewed papers and one patent. At completion of the project, the Hyundai Motor Group R&D team has chosen our proposed technique "Bluetooth Audio Quality Improvements Based on Qualitative SNR Analysis of Wiener Filter" as the best submitted research, out of more than 100 submitted papers

Dr. George Pappas is an Assistant Professor and Director Master in Artificial Intelligence program in the Department of Electrical and Computer Engineering (ECE) and currently working with several graduate and undergraduate students in research in a multitude of developing areas ranging from automotive to medical applications. He has over 15 years of teaching, research and work experience in embedded systems and high-performance computing. Artificial Intelligence (AI) in Autonomous vehicles, employ machine-learning techniques to collect, analyze and transfer data for safer driving experience. Also, he investigates encryption and optimization algorithms and security of the transfer of electronic medical data using wireless cellular communication systems for evaluation, diagnosis, and treatment of patients in remote locations. Some of the research interests are: Artificial Intelligence (AI) within radiology, specifically computerized tomography (CT) image reconstruction. Precise data analytics for pathology images. Virtual Reality (VR) in medical applications, Artificial Intelligence (AI) to aid diagnostics, Telemedicine, Medical and Health Informatics, Wireless implantable sensors and biomedical Transducers.