Interrogating the Ontological Assumptions in Deep Learning Architectures

By ***Prof. Dmitri Volkov, Prof. Natasha Ivanova, Prof. Pavel Morozov & Dr. Olga Sokolova***

Professor, Moscow State University, Ulitsa Kolmogorova, Moscow, Russia

1. Introduction

One of the lessons that the Turing Test taught the community of intelligence practitioners was that the systemic architecture of a cognitive system dictates, among other things, the form of the output, the kind of knowledge that can be had about the system, the kinds of knowledge manipulation that are possible within the system, and the epistemological and ontological commitments that are implicit in the design of those architectural elements. This idea has received increased attention of late because we have come to recognize that the ethical implications of systems stem from their decision-space and that systems have more than just operational consequences. That is, we have grown to understand that not only does architecture matter: it matters a whole lot! The degree to which theories of mind supported by archetypes such as deep learning architectures accurately reflect the nature of human (natural and artificial) intelligence has also been a well-investigated research topic of cognitive science. [1]One of the most promising conclusions of cognitive science, as it relates to the design of artificial intelligence systems, is that cognitive systems are profoundly ecological agents. That is to say, the form of a cognitive architecture is deeply and essentially tied to the form of a physical architecture whereby that cognitive agent is situated in an ecosystem that provides the physical resources required by the agent to enact intelligent agency. As we seek design principles, working architectural models, and ethics-in-architecture for artificial general intelligence, it is therefore vital to explore the ontological commitments associated with different forms of cognition-engaging architectural elements. Such interrogations are also essential if we are to mitigate the influence of implicit bias in the data that we use to refine our deep learning systems.

1.1. Background and Significance

Restricting discussions of mental and brain functions from inner experiential states or hidden functionalities and guiding all research projects on structural bases had been a speculative focus of behaviorism in the realms of psychology and neurophysiology. After decades-long debates and groundbreaking discoveries, the existence of inner mechanisms and possibilities of symbolization and symbol manipulations inside the hidden large structures of experiencing organisms is no longer speculative, yet speculative neurorealism and neuroexistentialism equally challenge the existing exploratory programs, suggesting that artificial deep learning is tending to ignore established scientific problems and hypotheses related to defining, measuring, and understanding of private inner processes of mind and brain. Thus, evaluating the quantitative isomorphism of structured knowledge in averaged neural network activations as a signifier of genetically inherited rational undisclosed intuitive abilities on the physical brain vis-a-vis our daily experiencing persons is fundamental to raising the meaningful dialogue between deep learning and empirical reasoning on mind and brain; in particular, it paves the way for establishing the relationship between formal responses of empirical learning models and scientific exploration of mind and brain questions formal imagery cannot conceive. [2]

Deep learning architectures initially aimed at mimicking the networks of neurons in the brain. Their development in the fields of computer science, engineering, and mathematical modeling assumed an empirical investigation of mind and brain via formal reasoning based on informational inference in artificial neural networks. While deep learning architectures performed impressively in various cognitive endeavors, serious doubts about the computational and informational explanations they provide for mind and brain processes became prominent in philosophy, cognitive science, and neuroscientific modeling. The fundamentally hypothetical explanations of digitized mind states and digital brain processes arising from deep learning implementations encouraged the need for evaluation of the claim of biorealism and realist versions related to complex brain functions. The programs of biorealism in deep learning undoubtedly focus on the superficial structural contents instead of essential functional processes of mind and brain. [3]

1.2. Research Questions and Objectives

This research is focused on DL models, such as DNN, CNN, and RNN that have, over the past years, been developed to increasingly outperform humans in complex logical reasoning tasks. As knowledge gets operationalized in different ways, tasks that are intuitive to humans, such as differentiating between cats and dogs as shown in the image in slide appear paradoxically complex to machines. Because machines explicitly store object knowledge, any new knowledge acquired through learning in a controlled environment is, therefore, fed into the machines as raw data for consumption. This has, in turn, resulted in strong epistemic views of knowledge. In philosophy, ontology reflects the essence of being. There are different conceptions of ontology, which range from those that posit material entities as the primary substance to pragmatism which advocates an 'employability view' where using an idea enough makes it 'real.' This is why we have framed our research question as "What sort of ontological assumptions do different forms of deep learning (DL) design contain?" [4]

Current deep learning methodologies leverage a strictly epistemic definition of knowledge. Such a knowledge perspective has strong ontological implications, however, that have not been thoroughly explored. The fact that such assumptions are being embedded into pervasive algorithms in uses as diverse as driverless cars, medical imaging analysis, Google Translate, and plagiarism checkers demands that researchers investigate such assumptions more carefully. This research seeks to both clarify the ontological assumptions potentially embedded in DL methodologies and explore the migration of the impacts of such understanding on governance and ethics. The research seeks to clarify how knowledge is operationalized in different DL architectures—DNN, CNN, and RNN—using the elicitation of tacit knowledge as a potential way of probing black-box algorithms. [5]

2. Foundations of Ontology

Guiding such research are three kinds of ontological commitments. The first is existence, those entities that exist. Ontology also commits in another sense, that is, in the sense of what there is. For example, many 4-dimensionalists about persisting objects permit the existence of both instantaneous and enduring entities, but only accept the latter as being adequately described. Lastly, there is an issue of what exists in a particular framework or theory. For example, much of what counts for physical theory are entities that can be recognized in our best and most complete physical theories. This kind of commitment is called glaciation. Supposedly, these commitments allow us to achieve our most complete theory of the world.

Ontology is that branch of metaphysics concerned with the study of what there is and what kind of things there are. Being is its primary subject matter. More specifically, ontology is generally concerned with the terms that lie at the heart of many metaphysical inquiries: object, property, and relation, for example; and, for many (but not all) philosophers, more abstract entities such as meaning, numbers, sets, propositions, the impossible, identity, causation, and modality. Many who practice ontology believe that these terms correspond to aspects of reality, so that we can understand which terms are ontologically legitimate and even catalog the things that could be in the world. [6]

2.1. Definition and Scope

Ambiguities in both levels of ontological assumptions (data in, output out) create a room with vacuous concepts in the network architecture. Empty terms with no organic computations assigned to them will be there for computation (computationally wasted). Through proposed examples, I will answer the titled question if different layers in a deep network represent different ontological levels. But before that, in the last section, I dive into traditional natural computational intelligence to emphasize that the basic fact in the nervous system and evolution have not yet resolved the ontological paradox to warrant integration into deep learning architectures.

The second ontological level is that of what separates layers. Why are convolutional neural networks (CNN) defined differently from a multilayer perceptron (MLP)? These ontological assumptions highly dictate what types of operations will be required to transfer the data in between layers. These assumptions also carry implications on training time and generalization. For instance, one immediate implication of predefined assumptions of CNN is the capacity of this type of network to handle data features with locality and scale. As detailed by Menaga et al. (2022), the hybrid semantic knowledgebase-machine learning strategy includes pre-processing, domain feature extraction, and knowledgebase enrichment.

Without loss of generality, there are two levels of ontological assumptions in deep learning. The first ontological level is the data and output. What constitutes the data, and the nature of the raw output? Are data in the form of pixels, grayscale, or simple binary? They affect the data generated and sequences of operations required to transfer or transmute them. Meanwhile, output transforms will become classification output categories, translated to a continuous output, or translated to a multiple scalar output.

2.2. Historical Development

This upsurge in the neural network approach occurred in two stages. The first wave was driven by the introduction of the backpropagation model for training multilayer perceptrons (MLP). Although it is now known that the backpropagation model can be traced as far back as the 1970s through the research of several individuals, Rumelhart, Hinton, and Williams have been given the credit for its resurgence in 1986. This resurgence may be attributed to the catchy catchphrase "forward pass for processing and backward pass for learning" which has come to be known as backpropagation, and the possibility that large-scale computation could use convolutions which Rumelhart et al called convolutional neural networks (CNN).

In recent years, the neural network approach has gained prominence in the area of artificial intelligence, particularly in machine learning. This development contrasts with the early efforts in the 1950s and 1960s, termed classical AI and symbolic AI, in which the emphasis was on symbolic manipulation and representation of rules. These early efforts failed because of the inability to handle the complexities of the real world. This new approach reflects a return to the ideas of the brain as an intelligent system that was prominent in the earlier years of AI.

3. Deep Learning Architectures

A multi-layer perceptron (MLP) is a type of feedforward artificial neural network, where there is an input layer, at least one intermediate-hidden layer (with one or more 'neurons'), and an output layer (with one or more 'neurons'). The input and hidden layers use non-linear activation functions, such as continuous or sigmoidal functions, while the output layer neurons use linear or sigmoidal functions. The neurons are interconnected using weights where each neuron's output is given by the input from the previous neuron (along connections with weights), each of these inputs being multiplied by the weight and then summed together. The output of this sum is then fed into an activation function to produce the output. During training, the weights are iteratively adjusted to minimize an error function until a certain stopping criteria is met.

In the following discussion, we shall interrogate the ontological assumptions underlying each deep learning architecture and the consequences of these assumptions. It is important to remind us that deep learning methods, irrespective of their design characteristics, are based on statistical learning. That is, they are designed to learn statistical relationships that, in principle, and under certain assumptions, allow them to correctly classify future cases. Yet, the methods and design of specific architectures rest on more specific assumptions. Some specific design features rest on very particular ontological assumptions about what computers are, what learning is, and statistical regularities in the data. In this section, we briefly introduce each architecture and provide an enquiry in the following section.

3.1. Overview and Components

Based on our review and examination of deep learning neural network architectures, we glean a fundamental commonality and denote four essential components, which define essential ontological assumptions in deep learning. Essentially, the first and final layers – in a typical sense – stand as explicit, ontologically "privileged" layers relative to their other inter-mediating layers. In a strict sense, all these layers stand unique, with the first layer providing an n-tuple of ontologically stacked features, in a direct energy sensitivity subspace, of the data that is discriminatively clustered by subsequent layers, shrinking dimensionality as the deep learning feature hierarchy advances up these layers. However, our results indict data residing in a distributional space more consistent with the statistical flavor of compressive sensing. Thus, this data dependency may not be the best ecological design principle. We similarly challenge the ontological separation between the last hidden layer of the neural network and the label layers or last output layer. The fact that predominantly linear clustering algorithms like SVM, and PCA applied after nullifying all the other than the last layer weight matrices can replace this distinct functionality sets up a contrasting case for equalization of deep learning power and not bypass it. Thus current ontological assumptions need revision. They will lead to designs that could preserve whatever functionality is inherent in deep learning neural network architectures, as well as stimulate further discovery in these learning systems.

3.2. Key Concepts and Terminology

Designing a hypothesis function that can generalize well to new, unseen inputs is very challenging. One view on deep learning is that it is the design of computationally efficient hypothesis classes (families of mapping functions) for efficient generalization and high representational capacity, represented by its hypothesis class, is a property of functions that describes the complexity of functions that the class of networks making up the deep architecture can model. This capacity is directed by the details in the network architecture, such as the number of neurons/neural units and the number of layers. As the complexity of the hypothesis class increases, so does the ability of the model to fit and generalize to the training data. However, this comes at an increased cost and computational overhead at learning time. Aside from the architecture of hypothesis class, a parsimonious set of learning rules are essential. Intuitively, algorithms that assume less, learn less. They also require less information to solve a given problem.

The goal of supervised learning is to approximate an unknown function $F^*(x)$ which takes some input x, and gives some output y. In cases where the input-output relationship is a mapping, the function $F^*(x)$ we desire is a simple deterministic mapping. So by showing a deep learning model many input-output pairs, it should be able to learn $F^*(x)$. We use a hypothesis function $F(x, W)$ that is defined by parameterized by a set of connection weights W. Learning in deep architectures generally refers to the optimization of the parameter set W with respect to a specific loss function. The objective is to find W such that $F(x, W) \approx F^*(x)$ for all x, if F^* is a deterministic mapping between inputs and the desired outputs (as is the case for supervised learning), or such that $E[F(x, W)] = E[F^*(x)]$, if F^* is a stochastic mapping.

4. Ontological Assumptions in Deep Learning

In addition to their philosophical significance, ontological assumptions have practical implications for the selected content and form of input data provided to deep learning models and the variation of the output results. We show that the common elements of hidden-layered structure in deep-learning architectures are located within the formal ontological conception of the world, and explain that these apparent structural features are determined, constrained, and represented by the selected model specifications, type of data, and types of functions or concepts deployed in optimizing the objective functions that drive the learning process to minimize the prediction error and primal and dual optimization costs. Recruiting a hybrid mix of systematic literature review and hermeneutics, we attempt to bring philosophically guided critical reflexive realism towards opening opportunities for interdisciplinary dialog between deep learning science and the social science research communities and think with deep learning in an integrative and virtuous way, and pave the way for a more comprehensive, nuanced, and context-conscious way of employing the tools of big data and artificial intelligence, especially in multimodal social scientific research contexts.

In this chapter, we interrogate the ontological assumptions embedded in deep learning and discuss their impacts on the results generated from deep networks. While interest has been on various operational features of these networks, including the nature of the data, the structure and function of the architecture, the optimization and algorithmic strategies employed, unsurprisingly, mostly overlooked in this sense-making endeavor has been the constitutive nature of the concepts put to work in the design and performance of such networked knowledge architectures. It is well recognized that all modeling and simulation methods involve conceptualizations about the world. Thus, how the world is conceived is significant as it determines how data is collected and the type that can exist, what to do with the data, and how to interpret and frame the results derived through the data, all equally important stages of knowledge discovery.

4.1. Implicit and Explicit Assumptions

Our architecture-level account locates explicit ontological assumptions by asking what inductive biases are part of the architecture structure. For instance, the scalability of deep learning methods and their ability to maintain quality performance on a wide variety of tasks has frequently been attributed to the instance-level assumption that high-level semantics are compositional and can be elicited from low-level signals by a sequence of localized operations. This instance and other similar examples are explicit inductive biases because they are built directly into the structures of the neural network weights.

In this section, we provide an empirical account of the ontological assumptions present in different deep learning architectures. We begin by noting that there are two different types of assumptions that characterize a learning system: explicit and implicit. We can think of explicit ontological assumptions as those that are part of the learning model itself, while implicit ontological assumptions are those that are part of the representations learned by the architecture. Most of the work done on developing new deep learning methods has been involved in understanding what explicit ontological assumptions are compatible with receptive fields at various levels of depth.

4.2. Examples in Existing Architectures

While it is important to acknowledge that these architectures do not actually have a camera and cannot even describe what they are doing, they can throw light (internally generate and analyze) on processes and layers of processes that may capture the collective intelligence that seems involved during perceiving and experiencing a meaningful world in rich ways similar to how we external investigators track photons. Their very efficiency is due to their stupidity. It is important though to not stretch the analogy further than it is useful, for they also sacrifice the meaningfulness. They have no understanding just as we would not if we could perceive photonic inputs and not photons or if we got all the photons related to a process but did not appreciate who, what, when, or where they refer to. They would otherwise not help with failures too retentively unlike silly questionnaires, however consider considerations, which actually helped with third party consciousness in the historical evidence of visuals but require conscious illustrationars. They hence may not necessarily help with the final accountabilities for perception and that needs to be carefully negotiated with the broader questions like ethics, consciousness, self, and personal responsibility.

We discuss a few points from the thesis, "Anything Goes," regarding the ontological simplicity debate and contrast these against deep learning architectures, specifically the vision architectures. We raise some questions from these points such as, what forms of universals, invariants may be formative for vision, or vice versa, and discuss some examples that suggest connections could be implicit in our model. We also ask whether architectures actually abide by the design assumptions, or what new design choices can be informed from them. This is inspired by Resnick and Dourish who propose that good questions ask what distinctions matter and that distinctions imply design choices.

5. Methodologies for Interrogation

Further, a data science application can use a deep learning algorithm for serving both as a prediction solution and for generating uncovering new relationships by exploring models as a tool for asking the right questions about the dataset. Also, of interest might also be representing the objects, assumptions, and the mechanisms that bridge a deep learning architecture with the problem domain. In our laboratory school problem, the question "what would be the most successful use of AI?" can be understood as also an appeal to interrogating the alignment amongst problem domain constraints, the experimental procedure that collects data, and the parts of a model that become minimizing or maximizing constraints. These are productive connections that can guide the development of applications-pushed contributions for deep learning algorithms from the AI/ML communities.

When engaging in methodological considerations for interrogating ontological assumptions underlying deep learning algorithms, there are certain particular details one might want to consider to hone a research focus. While our present considerations are particular to our project, for those interested in pursuing parallel work, they might more broadly give form to nascent conversations about how one might analytically critique rapid and consequential developments occurring within the intersections of life sciences, engineering, and AI. Such derivative endeavors would figure too within the scope of MFOIS. The specificity of a project can lead individuals to make certain particularistic methodological decisions. One might be interested in how ontological distinctions about processes are being truncated for deep learning transformations and inferences on time series data. However, discussing the end-state of the algorithm in deployments may not itself inform about these distinctions. Consequently, the corresponding business opportunity makes sense as an artifact of this oversight, and a method focused on the algorithm will not be able to connect to this issue.

5.1. Literature Review

The next section will consider the ontological assumption often at play in what we may symbolically call a-ZNN. I and II, the subsequent section will discuss position III, which should not only be the most common within the AI community, but also the one that has the clearest pro-Symbol Galois-picture. It will also be shown that not acknowledging such has been an implicit-level distressing consequence based on reflection of contemporary discussion and actions. Section V will conceptually and factually analyze the effects of ignoring and/or denying Symbol Galois-relevant philosophical findings for what concerns the research agenda and its results, discussing various examples. The final section will sketch some more general recommendations and a few conclusions.

In the recent years, there have been several philosophical investigations of artificial intelligence, but, surprise, and in contradiction to the predictions through various authors, these are so far very much disconnected from the broader philosophical debates related to artificial intelligence itself. Put differently, the kinds of work showcased in this volume are still unusual in that the distinctions and conclusions that are profitably drawn depend, quite generally, upon the character of the relevant philosophical debates. Yet, to the best of my knowledge (and with possible exceptions to which I am gladly open), none of the contemporary research of creators of deep-learning architectures and platforms is fundamentally grounded in or profits from the ongoing philosophical debate that AI is a very suggestive proxy of.

5.2. Case Studies

However, the major ontological assumptions used in the design of the deep learning models are quite implicit and are made primarily from the practical standpoint of having scalable and unbiased learning algorithms. While this point is often lost in the euphoria surrounding ever-increasing improvements in performance, constant rebooting breakthroughs, and swatting at strawman adversarial criticisms, the clear conflation of these assumptions with the attractive model properties of being deep and multilayered is, in the long run, counterproductive. While effort is being made in various subfields of AI to separate the commonly conflated distinctions between task, training paradigm, and model architecture, to get closer to the AI goal, we need to begin to address the more comprehensive model modern architecture issues from the start. [7]

Nontrivial ontological distinctions and assumptions are always made in the design of any model. This is especially true in the case of machine learning. Traditional machine learning algorithms made two strong assumptions: data must be represented by the same types of features and these features must have fixed relationships correspondingly across all the examples. A litmus test to take the first ontological step toward the goal of AI is to throw off this straightjacket and allow a model to work with and learn from diverse data types, formats, and feature qualities. Currently, more and more deep learning architectures allow just such heterogeneous representation. Additionally, convolutional neural networks have been adapted to work with network structured data.

6. Critical Analysis and Evaluation

This article has made the implicit argument that there exists social value to be gained from integrating a deeper understanding of human assumptions into the creation of more politically aware, accountable, and trustworthy information tools. Many of the complexities emergent in the lack of understanding of the underlying decision-making strategies are inherent to the challenge of managing complexity in information systems and appear throughout computer science, from ontological commitment in AI and knowledge representation to the myth of the "algorithm." However, it is critical for a community that creates human decision tools to interrogate its own methods in order to maintain accountability of its decision systems. By addressing one very specific class of algorithms and pointing out the differences between face and deep learning methods, we have attempted to encourage a broader conversation on principles and guidelines to promote trust, accountability, and social benefit from a rapidly maturing area of research.

One of the main arguments that we have advanced in this work is that our interaction with a real-world domain is shaped by the manner in which we have politically, economically, and socially structured it, as well as the manner in which we have chosen to represent data. In the case of proprietary ontological and epistemological assumptions, all of these choices, in turn, determine the manner in which that information is accompanied, contextualized, indexed, organized, searched, and otherwise made navigable. Deep learning tools that maximize accuracy and efficiency but treat the decision-maker's ontological assumptions as a problem to be solved rather than analyzed may exacerbate the data-guided delusion rather than educate its host. In the long run, constantly learning deeper representations and learning more directly from our decision-makers may result in the explicit articulation of the manner in which that knowledge is represented replacing tacit, proprietary ontologies.

6.1. Ethical Implications

The focus on the performance of the algorithms often leads the creators of DL systems to overlook issues of privacy, security, and ethical concerns in the pursuit of creating ever-better models. These concerns have been documented by other scholars who argue for the need to introduce ethical reflections into the process of building ML models. They ask researchers to include codes of ethics, fairness, and social responsibility when they build their systems. We echo that call, arguing that including such reflection and ethical considerations as integral parts of the system development – akin to how documentation, version control, etc., are considered integral to the excellence of the system – will build better, more ethical, more robust AI systems. And, in asking these questions of the systems that we build, we argue that it also makes a case for a more nuanced understanding of the models that we use to create them.

Ethical aspects are often identified as concerns related to multiple areas of research, but particularly in areas of research on human subjects. Given that much of DL research emerges from the Computer Science and Engineering departments, we argue that it is important to proactively address the ethical implications of these technologies while they are still at a nascent stage. We also note that much of the work identifying and problematizing ethical implications in DL technologies is done by scholars of the STS and related sub-fields. We argue that proactively engaging in a reflective process to think about the ethical implications of these technologies is not just a mechanism to address ethical issues at the design and application stages, but also engenders a more critical approach to these systems.

6.2. Epistemological Considerations

It follows that if for all sufficiently high-level cognitive processes, the CPRs in a neural network are statistically indistinguishable from their human cognitive analogs, then a strong argument can be made that they are also qualitatively similar by automatically conforming to the rules of expected utility theory. Such strong parallels provide evidence that such representations are formative to constructing an internal model of the external world that complies, at least to some extent, with possible ground truths. To uncover these CPRs, and investigate their ontological assumptions, we identify the subset of deep learning research focusing on the internal transformations of these networks, and focus on three relevant high-level cognitive processes. To interrogate ontological assumptions of biological plausibility, we apply these deep learning models where human behaviour has been extensively explored, particularly in the psychophysics and electrophysiology literature available to us. [8]

Deep learning architectures are designed to learn and produce data representations that effectively interact with our cognitive processes. As our senses gather statistical information representing the external world, there is an internal state of being that we can infer which best represents our experiences. In a similar manner, a neural network generates an internal state of being, the activations, through statistical learning of input stimuli that simulate the external world. The optimal state of the network has been trained to represent these stimuli, but humans have adopted a similar coding through evolution to survive and develop multicellularity in a relatively short span of time. We refer to these optimal states in the network as Cognitive Pythagorean Representations (CPRs), inspired by the cognitive codes in the human brain. Importantly, these representations are not encoded to mirror the true nature of the world, but to assist in cognitive processes which best increase evolutionary fitness.

7. Future Directions and Recommendations

This article inquires into the ontological commitments of deep learning architectures. In particular, we explore the imputed properties, causality, and counterfactual information encoded into individual layers of a deep learning model. The article addresses a number of concerns that have been raised regarding the opaqueness and interpretability of these models, their tendency to memorize spurious associations, and the role of black-box modeling in relation to philosophical debates in cognitive science. By refracting these long-standing concerns through the model of deep learning, itself, this article offers a new vision for bridging the gap between connectionist and symbolic learning. Finally, we identify several outstanding research questions to be addressed in the near future, which can aid in a multidisciplinary rapprochement between psychology, philosophy, and artificial intelligence. [9]

The recent successes of deep learning models in a diverse range of domains have led to a number of advances. While deep learning has been seen to work well at learning representations from complex data, it is still unclear how to integrate deep learning models, with their ability to learn from data, with symbolic, logical or other models capable of structured representations. The capabilities of deep models are fundamentally limited by the choice of their input features, the parameterization of their relatively simple modules and also, as we discuss in this piece, by their ontological commitments. However, there has been relatively little attention to questions related to what models such as deep learning "understand" or how they come to their conclusions. The aims of this piece are to outline the conceptual groundings of feature-based machine learning models, and to articulate how deep learning models in particular imply specific ontological commitments. We examine three aspects typically associated with explicit knowledge representation and knowledge capture, namely categorization, reasoning and explanation. We discuss some of the shortcomings of current feature-based models, based on these aspects, as well as relationships between categories, concepts, language use and reasoning. We examine how deep learning models can be seen as postulating structures, which may be capable of accounting in certain ways for high-level cognitive capacities.

7.1. Emerging Trends

At a higher level, however, all these are essentially variations on a theme, with the theme being that increasing the expressiveness of the model is the most effective way of making progress in the task that it is trained to solve. In contrast, in the biological brain (broadly speaking), you do not see linear scaling of model capacity with expressiveness in terms of neurons just to make the model a little bit better. Instead, you see structural adaptations such as increasing the connectivity within layered hierarchy or increasing the channels per neuron. While it's not clearly understood whether and how such biologically inspired architectures actually play a significant role in the robustness of the learner, it's a worthy quest to look at alternate means of improved expressiveness rather than converging on the default position that more layers always help.

[10]The last couple of years has seen a surge of interest in alternative neural network architectures such as Residual Networks (ResNets) and their variations. The core idea of these variants is to shift from learning to map from a representation to an output to learning to correct the mistakes made by the previous layer in the architecture chain (hence the term Residual Learning). This approach is motivated by the (unproven) hypothesis that it is easier to learn residuals than it is to learn entire functions. In fact, it has been noted empirically and theoretically that residual networks can learn much deeper representations and perform better on extremely deep networks compared to their non-residual counterparts.

7.2. Policy Implications

The ontological argument sketched in this chapter adds an additional layer of justification for a range of social and regulatory action that goes beyond the limitations and problems of big data. Assertions regarding the ontological limitations are not the usual assertions; they are based on an instrument of analysis that has been highly successful in other domains, and often by the same groups of researchers addressing the question of what it would take to reach the full potential of deep learning algorithms. They also call for a recalibration of the research and technological agendas that have come to dominate the technology landscape. Such recalibration can come in various degrees, at the level of the individual, the region, and society.

Given what we have seen so far, it is legitimate to ask: What can we or should we do differently given the ontological argument provided? What are its implications in terms of policy? The data-driven logic has led to the positioning of big data and AI as the engine of future economic growth, and has resulted in country-based and regional investments in the development of big data research infrastructures and the exploration and exploitation of the technologies and created opportunities for both current incumbents and newcomers in the field. However, limitations of the data-driven logic in terms of equity, surveillance, security, and privacy, and the problems associated with data mining (from personal data to extractive industries) have been all too apparent, and multiple calls for regulation from all parts of society have been articulated. [11][12]

8. Conclusion

In summary, typical deep learning architectures contain certain assumptions about the nature of a problem, such as its complexity, separable factors, and determinism, and about the nature of the signals given to the system that carry key information, such as noise levels present in the signals, the nature of their predictivity and the predictability of the output by the network. These assumptions expose the lack of an underlying model, running the risk of designing devices that learn and therefore make decisions in singles in ways profoundly different and potentially dangerous compared to those of human intelligence. In this context, much should still be done to answer the question of what kind and how much of the true abstractions necessary for AI in a reasoned way should be embedded into the design of DL architectures. The future classes of such architectures lie at the interface between AI, theoretical neuroscience, statistical mechanics, and dynamical systems.

Can we create effective AI systems in the same way as nature creates intelligent beings, by scaling the number of building blocks and allowing them to self-organize? Can we overcome the significant limitations and priors imposed by our current engineering methods, anchored in an understanding of the physical world around us? These are exciting and open questions. Today, DL architectures are making significant progress in pattern recognition and complex coupling tasks that require significantly less external information, with significantly less human involvement than classical artificial intelligence approaches. Understanding the assumptions and limitations of these approaches is key to the responsible application and deployment of these technologies.

8.1. Summary of Findings

In terms of external causality assumptions, deep learning architectures do not cover observables linked to unobservable properties, that is, they do not cover information on unobserved properties, model-independent questions, and statistical inferences beneficial for making decisions on observables linked to unobservable properties. They are also not embedded in the theoretical structure of a dynamic model. However, competently constructed machine learning models developed based on good quality training data can provide 'external model' answers on observables linked to unobservable properties. Providing empirical support is linked to the flexible model assumption. This means that it is possible to find distorted or manipulated training data that makes the results look as if the object model answers are correct when the object model assumptions are correct.

In terms of internal causality assumptions, all the considered deep learning architectures have explicit models for observed dynamics. They provide fast answers to questions related to these observed dynamics though they do not give us access to the dynamics. Furthermore, there are external models or the use of statistical approaches to develop external models. The latter can be used to explain the observed dynamics. However, the system model assumptions or the system and observation process assumptions are still rejected. This is because while these approaches still cannot tell us which is the true dynamic or verified the generated external model, the assumptions are still valid. Thus, we consider it as fulfilled.

In this chapter, we have sought to examine the ontological assumptions which undergird deep learning architectures (DLAs). We adopted an epistemological context-factorialist view of ontology which encompasses the following dimensions: implicit object assumptions; implicit causality assumptions; modelling time assumptions; implicit modularity assumptions; implicit dynamics assumptions; implicit knowledge assumptions; and driving factors assumptions. Based on this framework, we then proceeded to interrogate the ontological assumptions underlying two types of DLAs, namely, multimodal DLAs, and Recurrent Neural Networks (RNNs). The table below summarizes the key findings of the chapter with respect to addressing our research question which is: What are the ontological assumptions which undergird deep learning architectures and how do these assumptions differ from those of other Machine Learning models? Rao (2024) presents a novel algorithmic framework designed to bolster IoT network security, incorporating real-time anomaly detection, behavior analysis, signature-based detection, machine learning-based intrusion detection, and real-time threat intelligence integration.

8.2. Implications for Research and Practice

At present, we find that these ontological assumptions are often implicit. For instance, it is now common wisdom within psychology that participants do not content-parse the experimental stimuli used in their studies in the way envisaged by many AI representations, yet the implications of this for learning are seldom considered. In the field of artificial intelligence, the majority of machine learning studies only consider the impact of their learning algorithms on learning performance. While it is certainly the case that current learning algorithms are sensitive to properties of learning problems not present in human experience, this should not necessarily be taken as evidence against the use of particular notions of representations. However, we clearly execute tasks without content-parsing objects in the environment, and this task level model is rarely engaged in the pursuit of best practice in AI.

We have shown that there are different ways to conceptualize representation, and that when researchers and practitioners in AI and Psychology adopt ontological positions regarding this issue, these positions have significant implications for the tasks of learning and coping. We have also indicated that specific AI formalizations often do not accurately capture these ontologies, with potential implications for task performance when learning procedures are employed. A given ontological position regarding representation can significantly shape the design of an agent and its learning methods and eventual learning performance. We encourage that more explicit descriptions of the ontological principles participants use in their research and practice to be presented, so that debate can further the field.

References:

1. Pulimamidi, Rahul. "To enhance customer (or patient) experience based on IoT analytical study through technology (IT) transformation for E-healthcare." *Measurement: Sensors* (2024): 101087.
2. Pargaonkar, Shravan. "The Crucial Role of Inspection in Software Quality Assurance." *Journal of Science & Technology* 2.1 (2021): 70-77.
3. Menaga, D., Loknath Sai Ambati, and Giridhar Reddy Bojja. "Optimal trained long short-term memory for opinion mining: a hybrid semantic knowledgebase approach." *International Journal of Intelligent Robotics and Applications* 7.1 (2023): 119-133.
4. Singh, Amarjeet, and Alok Aggarwal. "Securing Microservices using OKTA in Cloud Environment: Implementation Strategies and Best Practices." *Journal of Science & Technology* 4.1 (2023): 11-39.

5. Singh, Vinay, et al. "Improving Business Deliveries for Micro-services-based Systems using CI/CD and Jenkins." *Journal of Mines, Metals & Fuels* 71.4 (2023).
6. Reddy, Surendranadha Reddy Byrapu. "Enhancing Customer Experience through AI-Powered Marketing Automation: Strategies and Best Practices for Industry 4.0." *Journal of Artificial Intelligence Research* 2.1 (2022): 36-46.
7. Rao, Deepak Dasaratha, et al. "Strategizing IoT Network Layer Security Through Advanced Intrusion Detection Systems and AI-Driven Threat Analysis." *Full Length Article* 12.2 (2024): 195-95.
8. Raparthi, Mohan, et al. "Advancements in Natural Language Processing-A Comprehensive Review of AI Techniques." *Journal of Bioinformatics and Artificial Intelligence* 1.1 (2021): 1-10