Respect the Unstable, Scale and Constraints in the Era of Artificial Intelligence

Evangelos A. Theodorou
School of Aerospace Engineering
Georgia Institute of Technology

Date: May 27, 2025

1 Motivation

As we enter the second quarter of the 21st century, there is enormous enthusiasm about Artificial Intelligence (AI) and its role in our societies. The expectation is that AI will influence all aspects of our lives by providing solutions to problems that have not been solved before, using enormous amounts of data and expensive compute. And while there is overwhelming excitement for AI and its impact, there are concerns about creating machines that will be smarter than humans. So, where is the truth in all this?

In this article, we argue that the current dialog on AI is limited to problems and use cases where AI is destined to succeed. And these are problems where there are enormous amounts of data (big data) and available compute. But even in the cases where both data and compute are available, ensuring safety, trustworthiness, and generalization is a difficult task, especially for use cases such as decision making. We assert that all the discussion regarding Artificial General Intelligence (AGI) is naive and disregards the need to respect the unstable, scale, and constraints of many real-world systems and networks of systems in which it is important to have some form of autonomy. There exist different forms of complexity in the real world and these forms are overlooked when we talk about AI and its capabilities in decision making. So, what are these different forms of complexity in the real world? Why they are important to our societies? And why will there not be a *one-architecture-fits-all* approach when developing decision-making systems using AI? What other fields and disciplines are needed to bring AI to solid ground and make it trustworthy?

To support our arguments in this article, we look back at the history of control systems society and identify moments of reflection and skepticism regarding scientific practices. We chose control theory because it is one of the closest disciplines to current efforts and future promises to use AI for decision making. My goal is to revisit epistemological debates of scientific disciplines that had a great societal impact, such as control theory, and to create a relevant context within which we can have meaningful discussion that goes beyond the current hype about AI.

In discussing these issues, we will use the terminology of *AI-Driven Predictive Decision Making* or *AI-Driven Planning* systems to refer to decision-making architectures that use machine learning algorithms to represent different modules within the decision-making architecture, including objective functions, dynamics, and/or policies. In addition, we will use the term *Architecture* in two settings. The first to describe the internal structure of *AI-Driven Predictive Decision Making* systems which may consist of computational modules with specific functionalities such as planning, perception, low-level control etc. The second is to describe the structure of neural network representations themselves.

If you do not have time to read this article just read the text in boxes ☺.

2 Looking Back: What happened to Control Systems Society?

Back in 1989 at the international Conference on Decision and Control (CDC), Gunter Stein delivered the first Bode Lecture in the history of Control Systems Society [1]. The Bode lecture and the corresponding Bode award is the most prestigious recognition for a scientist working in areas of control and dynamical system theory. The title of Gunter Stein's lecture is "Respect the Unstable" and the topic is motivated by the two major trends of that era, namely the need to work with more dangerous systems, and the trend to worship more formal mathematics for control. At that time the control systems society had reached its peak with four prior decades of significant contributions in areas of optimal control, state estimation, robust and adaptive control. As Gunter Stein put it bluntly:

"Control group is a very succesful group. Think about control systems ... they are in our houses, they are in our cars and transportation systems. They are in our factories and military systems ... they are everywhere ... they work and work and work we hardly ever notice. So we are enormously successful and we should be proud of that. But what issues do we have?"

Gunter Stein continues by expressing his concern regarding the detachment of mathematics from the real-world physical consequences. His focus is on methodological approaches used for the design of feedback control systems in domains and applications characterized as safety critical. As he explains:

''We certainly place a lot of value on mathematics, and I think we are doing it at the expense of the real practical physical consequences of some of the things that mathematics describes. If someone accepts that, then these two trends are in opposition. .... We, at the moment, are very respected society. Society around us let us control things that are dangerous ... more and more of them. They trust us. Our

technology is well respected and no-one questions it basically. But if these two trends continue that might not be true. We may in fact join the other areas of technology in terms of how trustworthy we are."

Gunter Stein's lecture is a testament to the maturity of the control systems society of that time. He not only emphasized the importance of connecting the mathematics of control design to physics, but also provided a path forward by proposing that we can worship mathematics and bring advanced methods for as long as there is connection with the real physical consequences that these mathematics describe. In his lecture as well as in his paper, Gunter Stein discusses famous accidents in the domains of aerospace engineering, nuclear energy, including the Chernobyl accident in April of 1986. The core concept in Gunter Stein's message is the concept of stability in dynamical systems and the fact that unstable systems are much harder to control than systems with inherently stable dynamics. His paper is definitely worth reading even for scientists outside the area of control theory.

Mathematical fanciness should not come at the expense of losing the connection with the actual physical consequences that mathematics describes. This loss can have serious implications.

3 Moving Forward: The Era of Artificial Intelligence

When listening to Gunter Stein's lecture [1] and reading his article [2], the first thing that comes into my mind is whether the AI community today is at the point of where the control systems society was back in 1989. Whether the dialog about AI addresses the true epistemological issues of the field. Does AI bring us closer to the real, physical consequences that mathematics and AI representations describe, or does it abstract us further? Are there trends today that are in opposition? Can we have a *one-architecture-fits-all* approach or do we need architectures that are tailored to the decision-making task being considered? If the answer is yes, how can we talk about AGI when architectures need to be tailored to the decision-making task to ensure safety, guardrail satisfaction, and scalability? Should we focus only on Autonomous Vehicles when we talk about applications of AI to decision making? Next, we will try to address all these questions starting from the last one and going backwards.

3.1 Autonomous Vehicles (AV): A popular application of AI for decision-making.

AVs, such as self-driving cars and terrestrial vehicles, are one of the most popular use cases of AI for decision making. Existing technologies are just at the beginning, with few promising efforts in demonstrating autonomous driving in specific environments and cities. There is a lot of work to do in scaling autonomy architectures to different environments and achieving generalization and safety. It is one thing to demonstrate a proof of concept, and it is another to make sure that this proof of concept scales in different environments and meets the safety and generalization requirements. And, as usually happens in challenging engineering projects, this last step takes 80% of the total time and effort. Therefore, I am not underestimating the difficulty of AV. I am just trying to say that this should not be the only use case when we talk about AI as there are other forms and notions of complexity that go beyond the complexity of self-driving cars and single-agent robotics problems. In addition, there is the area of Urban Aerial Mobility, which adds an entire new dimension to the open problems we have in autonomy and transportation.

In general, I think many of the questions related to safety and generalization will boil down to what is the best decision-making architecture, that is, how to bring perception, planning, and control/adaptation together. Should we design architectures that are end-to-end and map raw multimodal sensing data to control decisions? Should we impose a nominal architecture consisting of perception, planning, and low-level control modules? How are uncertainty representations embedded? Which architecture is better suited for generalization to different environments? There is a difference when vehicles have to operate in partially known environments and in an environment for which there is no prior information. How can adaptation improve generalization and which modules of the perceptual decision-making architecture should we adapt? Although there is a lot of work on this space, it is important to test the proposed architectures in the real physical world and not just rely on simulations. For this reason, there is an issue with those who try to predict the future of AI for decision making without showing demonstrations with experiments in the real world or when they do not test their research ideas on challenging autonomy tasks with real hardware. It is one thing to propose a decision-making architecture for autonomous driving, it is another thing to actually do it.

Safety in autonomous vehicles will not come by just imitating experts. We need to be thinking about optimization-based decision-making architectures and be obsessed with rigorous real world testing.

3.2 The dialog on AI should include large-scale decision-making problems

The mainstream discussion today about the capabilities of AI includes use cases such as self-driving cars, recommendation systems, applications of AI to health care, biology, information retrieval, weather forecasting, and general applications of AI to science. Although this is a fairly broad list of areas, there are challenging large-scale decision-making problems that have a high socioeconomic impact and are crucial to our society. An example of large-scale decision-making problems is resource allocation for general supply-demand systems.

In resource allocation problems, resources are allocated to meet demand requirements and achieve certain economic objectives in the presence of constraints and uncertainty. What makes resource allocation problems challenging is that they are typically large-scale, with dynamics, constraints, and decisions operating at different spatio-temporal scales. These scales may vary from seconds to days or weeks. Another characteristic of resource allocation problems, and in general of large-scale decision-making problems, is that there are no human data that one could use to learn optimal policies. Although humans are good at performing low-dimensional decision-making tasks, such as driving a car or manipulating objects, they really struggle to handle large-scale decision-making tasks. This means that the only option available to find optimal policies is optimization with all the scientific and engineering challenges that this option entails in terms of scalability, feasibility, and robustness to uncertainty. And this is where the actual challenges lie in dealing with networks, the complexity and scale of which increases day by day, month by month, year by year. An enthusiastic advocate for AI may say "*I do not care how data are generated if they are generated by a human or an existing decision-making system. For as long as the data exist, I can always learn from them*". Well, I wish it were so easy. So, yes you can learn from the data, but the question is whether you can guarantee that you will not violate the physics, guardrails, and constraints of the problem in consideration and in particular as the scale of problem increases. I will address this issue later in this article when discussing model-free versus model-based neural network representations. But for now the main point I am making is the following.

From self-driving cars to resource allocation problems in general supply-demand and network control systems, there exist different notions of complexity. These notions start from the complexity of perception-action-adaptation loops at the single agent level and expand to the complexity of scale in networks and large-scale decision-making problems (see Fig. 1). The dialog on AI should include the entire spectrum of complexity and go beyond the popular use-cases of content (text, audio, video) generation, self-driving cars, and applications of AI to robotic systems which, with the exception of swarms, are for the most part low-dimensional.

Of course, the public and the media are excited about the capabilities of Generative AI especially as they relate to use cases such as content generation and information retrieval. Today, every professional can use AI tools to support work and increase productivity, and this is great. However, as scientists, we have the obligation to talk about problems that are difficult not only to solve but also to communicate to the public. In the context of large-scale networks and general supply demand systems, we take for granted our ability to deal with the increasing complexity, scale, and the different sources of uncertainty. In many cases, such systems are part of the critical infrastructure of our societies, and their continuous operation is vital to basic

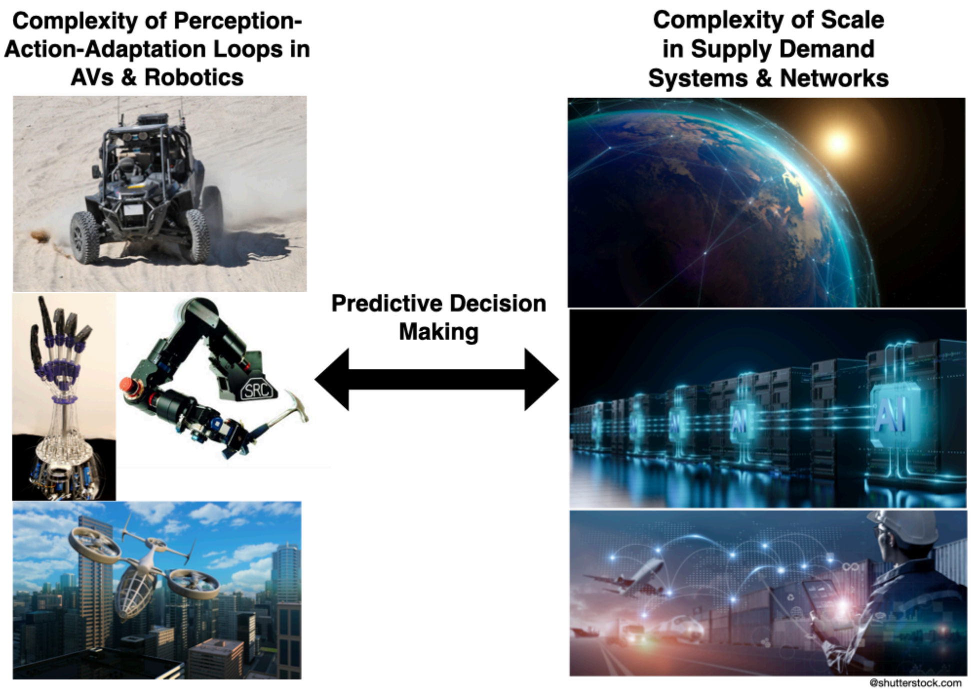

Fig. 1: The different notions of complexity in predictive decision making: From autonomous vehicles and robotics to networks and supply demand systems such as satellite mega-constellations, data center networks and transportation/logistics.

operations, economic activities, and needs in the areas of transportation, telecommunication, and information networks as well as in power and energy management.

It is myopic to focus only on autonomous vehicles when discussing AI and decision making. There exist different notions of complexity that go beyond the complexity of perception-action loops.

3.3 Why are resource allocation problems important?

Next, I would like to discuss three examples of resource allocation problems that will have a significant impact on our societies. These examples include resource allocation problems in data center networks, satellite mega-constellations, and transportation/air traffic management. All examples share the same characteristics, that is, they include networks of increasing complexity, scalability, and demand. In addition, if not already, these networks will be part of the critical infrastructure of our societies, and thus their non-disrupted and safe operation is essential.

Starting with the data center management use case, there is a clear trend in increasing energy consumption due to the growth of AI-related technologies and the need for more compute. Recent reports [3] on data center energy usage predict a demand increase of 4.4% of total electricity consumption in the US in 2023, to forecasts that range from 6.7% to 12% by 2028. Given this surge in computing demand and the need for more energy, there is an ongoing discussion in terms of what the best technologies are to support the energy needs of data center networks. These include clean energy solutions, nuclear energy and technologies such as Small Modular Reactors (SMRs) [4, 5]. From a decision-making perspective, the data center management is a large-scale network control problem with decisions, dynamic processes, and constraints at different spatio-temporal scales [6], and different sources of uncertainty including forecasting demand, stochasticity in energy production and weather conditions. As the scale and complexity of these networks increases, there is a need for distributed optimization and scalable decision-making technologies that can efficiently plan for resource allocation - how to allocate compute resources - to meet increasing demand requirements while satisfying capacity, energy constraints, and safety guardrails in the presence of uncertainty. Some relevant academic work on this space includes references [7, 8, 9, 10, 11, 12, 13, 14, 15].

In the area of satellite networks, there are several studies that forecast a growth of the total market economy [16, 17]. Although forecasts have their own assumptions, there is a clear trend in terms of an economy that will increase its size due to the decrease in launch cost, several technological innovations, and a diversified pool of investments that includes the government and the private sector. This rise will have a significant impact on a wide range of industries, including automotive and manufacturing, aviation, information technology, supply chain and transportation, and many others, by providing new capabilities and services. And just to understand where we are headed, projections on the scale of satellite networks tell us that by 2030 there will be $63,000$ satellites orbiting Earth [18] from 5,500 in 2022. In this context of satellite constellations on a large scale, we are talking about moving supply demand networks with a topology that varies in time in which there are several large-scale decision-making and resource allocation problems that need to be considered and solved in a coordinated fashion [19]. As the scale of such networks increases in terms of number of satellites, users, and services, the complexity of corresponding decision-making problems will also increase.

Finally, in the domain of aviation, in a recent 2024 report by Federal Aviation Administration (FAA) [20] there is a clear trend of increasing demand in air transportation. As the report indicates, the total number of passengers to/from the US from American and foreign flag carriers will increase from 250M to 475M by 2044. There are several large-scale decision-making and resource allocation problems to be solved in air traffic management, such as slot allocation, capacity management, and route planning to ensure minimum delays, reduce airspace congestion, and perform optimal scheduling [21, 22]. The situation in transportation and air traffic management is becoming more complex when considering emerging technologies in the domain of Urban Aerial Mobility with the use cases of air taxis and air-cargo vehicles as well as drone delivery services [23, 24, 25]. Are we going to double the size of our airports and the number of aircrafts to meet the increasing demand? How are we going to handle the increasing traffic in the airspace? Although it is difficult to answer these questions, one thing is certain: we will need reliable and scalable technologies for large-scale resource allocation that can handle uncertainty, constraints, and increasing demand for transportation networks. Recent efforts [26, 27] on the design of market mechanisms are steps in the right direction; however, this is only the beginning, as more emphasis should be placed on how to handle sources of uncertainty and how to scale distributed optimization as networks increase their complexity and size.

I hope that by this point the reader sees a common theme in these three examples. This common theme says that:

There are resource allocation problems with a rapidly increasing complexity due to a number of socio-economic and technological reasons. Although AI technologies can be used to support functionalities such as forecasting or model learning within predictive decision-making architectures designed for these resource allocation problems, the core challenges lie in scaling dynamic optimization and having a principled methodology to take into account the different sources of uncertainty and constraints. The shortcut of learning from a human does not apply here. In addition, the discussion on comparison with human intelligence does not apply here either. These architectures must be interpretable so that they can be audited when this is necessary.

We should appreciate the complexity of scale in resource allocation and networked control problems. These problems have high socio-economic impact and their complexity rapidly increases with time.

3.4 Scale makes things harder

Satisfying guardrails and handling uncertainty becomes harder and harder as the scale increases. The scale of resource allocation problems is large due to the fact that they typically have many decisions to be optimized, and as we have already discussed, these decisions operate at different spatio-temporal scales.

An argument that could be an antilogos to our message here is: How is it possible that we can optimize neural network representations with billions of parameters but we struggle with resource allocation problems. The answer to this question is that resource allocation problems are large-scale optimization problems that are constrained by the physics of processes, as well as constraints that represent guardrails and safety requirements. From an optimization perspective, there is a profound difference between training a neural network for supervised learning and performing optimization in a large-scale network to allocate resources in the presence of constraints and uncertainty. In the former case, we are looking for a point in the state space that optimizes an objective function, while in the latter case, we are looking for a point in the state space that, besides optimizing the objective, has also to satisfy guardrails and constraints. In addition, another reason for why the two paradigms are different is that the economics of training deep neural networks will not work in the case of large-scale resource allocation problems. Very often re-optimization needs to take place in large-scale networked control systems as they operate in changing environments. How often depends on the levels of non-stationarity of the environment. For this reason, optimization in resource allocation problems can not be expensive.

It is also important here to appreciate the relationship between scale and uncertainty. In layman terms, in an ideal world that is completely deterministic, resource allocation problems in supply demand systems have already an enormous scale for the reasons that we have already explained. The addition of uncertainty in the decision-making process further increases this initial scale in a multiplicative way. Essentially, if you want to consider uncertainty in the decision-making process, the price that you have to pay is the increase in scale and, therefore, in computation. Due to this multiplicative effect that uncertainty has on the initial scale, handling uncertainty in large-scale decision-making problems is much harder than in problems with a small scale. In general, my main point here is this.

Starting with Gunter Stein's presentation back in 1989, it is important not only to respect the unstable but also to respect the increasing scale and constraints of networked systems. These systems have high socioeconomic impact, and their safe and non-disrupted operation are essential to our society. In the same spirit of how hard it is to control unstable systems vs. stable systems, it is important to acknowledge that it is hard to control and optimize large-scale systems. So, scale makes things harder even if the world was purely deterministic. The consideration of uncertainty in the decision making process adds more complexity by increasing the initial scale in a multiplicative fashion.

Scale and uncertainty make resource allocation problems very hard to solve. It is a lot harder to deal with uncertainty as the scale of a decision-making problem increases.

3.5 Removing the opacity of neural network representations.

AI-driven predictive decision making systems, especially those relying on black-box (model-free) representations, are not configured to guarantee safe operation that is

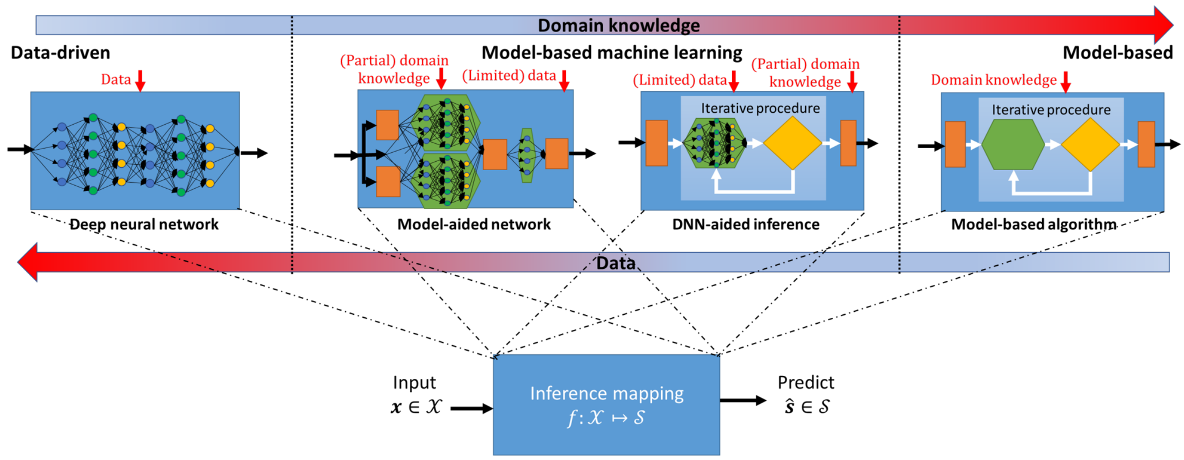

Fig. 2: Deep Unfolding and Iterative Optimization Algorithms. Figure taken from [32] with the author's permission

interpretable and generalizable. Safe operation means constraint satisfaction in the presence of uncertainty at all times and in all environments and for the case of resource allocation problems at all scales. It is hard to believe that because neural networks have learning rules such as backpropagation or because they can be trained in a distributed fashion that they will be able to solve resource allocation problems. There are other characteristics of neural network representations that are relevant to constraint satisfaction and safety, and from my point of view these characteristics are related to the underlying architecture and structure of neural networks. The good news is that there are ways to design the underlying architecture on the basis of the physics of the processes to which neural networks are trained to represent. The majority of the work on this space is related to Physics-Informed Neural Networks. These are essentially neural networks that are designed to represent systems with temporal or spatiotemporal dynamics [28, 29, 30, 31].

In a similar vein with the physics-informed neural networks, there have been efforts to design neural network representations based on the dynamics of iterative optimization algorithms. The resulting neural networks are used to learn the solutions of optimization algorithms with each layer of the neural network architecture corresponding to one iteration of the iterative optimization counterpart. The learnable parameters in each layer correspond to the open parameters of the corresponding optimization algorithm (see Fig.2). This area is known by the names *Algorithmic Unrolling* (AU), *Learning-To-Optimize* (L2O), or *Deep Unfolding* (DU) and the first attempt is reference [33]. Applications include the areas of signal and image processing, compressed sensing, and state estimation [34, 32, 35, 36], game theory [37], and optimization problems where decision variables are integers [38]. DU provides a range of architectural designs for neural network representations that can be customized based on the levels of domain knowledge incorporated into them (see Fig.2).

One of the most promising areas for DU is distributed optimization for large-scale decision making, training and quadratic programming [39, 40, 41, 42]. The result so far demonstrated that the deep unfolded architecture surpasses the performance of the iterative optimization counter parts, they can generalize well to similar problems and satisfy constraints as these are embedded in the architecture, they are interpretable since each layer corresponds to an iteration of the corresponding optimization algorithm, and they provide a functional representation of the optimization algorithms. This last characteristic paves the way towards certifying the aforementioned neural network representations and then deploying them to perform decision-making tasks in safety critical domains. I think this is where there is a lot of room for new ideas and research on neural network architectures that are tailored to the decision-making problem being considered. We will call this relatively new research direction: *From Optimization Algorithms in Decision Making to Neural Network Representations.*

In another line of research, existing neural network representations, such as transformers, have been derived based on optimization unrolling [43, 44, 45] using specific energy functions. These energy functions are designed such that when performing algorithmic unrolling the structure of known neural networks architectures emerge. Besides seeding light into how known representations can be designed in principled way, the connections to energy functions and algorithmic unrolling offer opportunities for novel neural network architectures. We will call this research direction: *From Neural Network Representations to Unrolling Optimization and Backward.* The logical steps of this line of research are in the opposite direction to the one described in the previous paragraph, in which there is no prior neural network representation in mind and the unrolling process starts with pre-existing optimization algorithms. In the context of decision-making and planning tasks, these algorithms correspond to constrained dynamic optimization and optimal control with objective functions comprise of actual physical energy objectives, penalties that encode safety guardrails, and terms that have some economic meaning and interpretability. Thus, it should be clear that the two lines of research are substantially different as they serve different purposes. So, my main point here is this:

The idea of AGI using general-purpose black-box neural network architectures is far from the need to design AI-driven predictive decision making systems that are interpretable and trustworthy. Only when neural network architectures are designed using the dynamics of the physical process or the dynamics of iterative optimization algorithm they are trained to represent can we talk about safe, reliable and interpretable operation.

The above message may appear to be restrictive for the main-stream research on AI. One may argue that there are abstract *planning* tasks that are related to higher and more strategic levels of cognition for which the corresponding states do not directly map into the physics of the real world. In fact, there are philosophical considerations on the topic of opacity of neural networks that are willing to accept opacity when neural networks are used as guides to explore or draw attention to promising avenues in scientific discovery [46] in contrast to previous epistemological efforts to define a framework for transparency in complex computational systems [47]. Although I am cautiously open to accepting this proposal for tasks such as scientific discovery, we should never give up on seeking to reduce opacity in deep learning. Even for high-level cognitive tasks such as scientific discovery, strategic decision making in management and economics [48], or planning in diplomacy and geopolitics [49, 50], there should be abstract representations of a state space in which high-level planning and search occurs. Either by defining a planning task as an optimization algorithm and then designing a neural network architecture using algorithmic unrolling, or by first heuristically designing a black-box neural network architecture and then seeking to understand how its structure emerges from properly designed energy objectives and optimization algorithms, there is a path towards reducing opacity even for high-level cognitive tasks. This path will require multidisciplinary work at the intersection of optimization, control, and statistics.

Optimization-informed neural network architectures is the way to go towards safe and reliable AI-driven decision-making. These architectures are interpretable and tailored to decision-making problems being considered.

3.5.1 Explainability in AI-driven Planning Systems

Although in the previous section we discuss the important issue of opacity in general neural network representations, the elephant in the room is on the explainability of decisions produced by AI-driven predictive decision-making systems. With physics- and optimization-informed neural network representations, we can reduce opacity in such systems. But what other sources of opacity exist? The answer to this question is related to the need to represent Value Functions (VFs) and the crucial role that these VFs play in decision making. VFs encode information about the future beyond the look-ahead horizon used in planning systems. Such systems typically optimize objectives that are split into short- and long-term, while the overall optimization is constrained by the physics and guardrails of the processes for which they are designed to make decisions. Without getting technical, let us consider a process that has states s and actions a with look ahead horizon T and $\dot{s}$ denoting the time derivative of state. The objective to be optimized has the typical form:

$$\text{Objective} = \underbrace{\int_{t_0}^{t_0+T} \ell\Big(s(\tau), \alpha(\tau)\Big) \mathrm{d}\tau}_{\text{Short Term}} + \underbrace{V\Big(s(t_o+T), t_o+T\Big)}_{\text{Long Term}} \tag{1}$$

Subject To :

$$\text{Dynamics}(s, \alpha, \dot{s}) = 0 \tag{2}$$

$$\text{Guardrails}(s, \alpha) \leq 0 \tag{3}$$

where V is the corresponding VF and ℓ is cost or reward accumulated over the look ahead horizon. Based on the mathematical representation in (1), (2) and (3) we shall think of VFs as a way to guide planning systems to proactively act in presence of future events and environment changes. One way to reduce the sensitivity of decision-making to the VF representation is to increase the look-ahead horizon T. Unfortunately, this choice comes at the cost of increasing the computational complexity, and it is typically avoided when decisions have to be real time or when scale is large. Essentially, knowing the VF is a very hard problem given the fact that VF should satisfy the Dynamic Programming Principle [51]. From a practical perspective VFs can be manually designed; however, in complex, and large-scale dynamical systems and planning tasks VFs should be learned. The ability to learn VFs depends on the levels of nonstationarity of the environment. This is a particular important point of consideration for systems that operate in the real physical world such as robotics, AVs, swarms as well as general supply demand and network controlled systems.

So, my overall point here is this:

When thinking about the opacity of AI-driven planning systems and the explainability of their decisions, this will certainly depend on the complexity and interpretability of the underlying VF representations. But even in this case, the resulting decision-making architecture would be so much easier to interpret compared to an architecture in which dynamics, constraints, objectives and optimization algorithms are all together represented by a single giant black-box neural network.

Value Functions (VFs) are important for guiding AI-driven decision-making systems. The explainability of decisions of such systems depends on our ability to interpret VFs and understand their interaction with optimization. This interaction is the source of *intelligent* behavior.

4 Concluding Remarks: The Lack of *Mesotes*[1]

In closing this article, it is worth discussing whether there exist two trends today that are in opposition. The answer is yes. The first trend is to develop general-purpose AI systems towards AGI. The second is the request to deploy AI systems to more and more cases in the real world and have the expectation that these systems will be trustworthy. Physics-informed and optimization-informed neural network architectures reduce the gap between mathematical representations and the real physical

[1] This concept is central in Aristole's Nicomachean Ethics. Mesotes ($M\epsilon\sigma\acute{o}\tau\eta\varsigma$) the virtue of avoiding extremes by integrating opposite views towards a golden mean ($\chi\rho\upsilon\sigma\acute{\eta}$ $\tau o\mu\acute{\eta}$).

consequences that these representations describe. Such architectures pave the way towards the development of trustworthy predictive AI decision-making systems and allow their deployment to safety critical applications, starting from self-driving cars and UAM vehicles to large-scale decision-making and resource allocation problems in networks and supply-demand systems.

In general, I would like to be optimistic and say that there is a great future for AI and that the overall impact on our societies will be positive. However, to develop trustworthy AI technologies, we will have to rely on disciplines such as optimization theory, control and statistics that are canonical and provide ways to design predictive AI decision-making systems that respect the unstable, scale and constraints of real physical as well as socio-economic systems and are transparent. In this sense, they will be safe for deployment in the real world. Therefore, I believe that there is a path forward for ethical, responsible, and trustworthy AI. However, this path is epistemologically in the opposite direction to where the mainstream AI ecosystem is moving today. In the post-modern era that we live in, I think it is important to have a balanced view between tradition and modernism, between the history of science and its future. Not everything that comes from tradition is great, but at the same time we should be skeptical with overpromising statements about the capabilities of certain technologies such as AI. Recognizing the existence of conservation laws, constraints, and scale makes us humble and allows us to make careful predictions about the future.

Thanks for reading!

Disclaimer: The positions and claims in this article represent only the personal opinions and views of the author.

Author BIO: Evangelos Theodorou is an Associate Professor with the School of Aerospace Engineering at Georgia Institute of Technology and is the Director of the Autonomous Control and Decision Systems Laboratory. Prof. Theodorou is an affiliate of the Institute of Robotics and Intelligent Machines (IRIM) and the Centers for Machine Learning Research. He is also an Amazon Scholar. His academic and industry research interests span the areas of Control Systems, Optimization, Autonomy, Robotics, Operational Research and Supply Chain Logistics.

References

[1] G. Stein, "Respect the unstable," *Hendrik W. Bode Lecture at the IEEE Conference on Decision and Control in Tampa, Florida*, 1989. [Online]. Available: https://www.youtube.com/watch?v=9Lhu31X94V4t=1638s

[2] ——, "Respect the unstable," *IEEE Control Systems Magazine*, vol. 23, no. 4, pp. 12–25, 2003.

[3] A. Shehabi, S. Smith, A. Hubbard, A. Newkirk, N. Lei, J. M. E. Siddik, M.A.B.and Holecek B. Koomey, and D. Sartor, "United states data center energy usage report," *2024 . Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637.*

[4] B. Poudel, K. Joshi, and R. Gokaraju, "A dynamic model of small modular reactor based nuclear plant for power system studies," *IEEE Transactions on Energy Conversion*, vol. 35, no. 2, pp. 977–985, 2020.

[5] Y. Wang, W. Chen, L. Zhang, X. Zhao, Y. Gao, and V. Dinavahi, "Small modular reactors: An overview of modeling, control, simulation, and applications," *IEEE Access*, vol. 12, pp. 39 628–39 650, 2024.

[6] X. Guo, Y. Che, Z. Zheng, and J. Sun, "Multi-timescale optimization scheduling of interconnected data centers based on model predictive control," *Frontiers in Energy*, vol. 18, pp. 1–14, 12 2023.

[7] S.-P. Huang, D. Yan, and Y.-C. Chen, "An online algorithm for combined computing workload and energy coordination within a regional data center cluster," *International Journal of Electrical Power & Energy Systems*, 2023. [Online]. Available: https://api.semanticscholar.org/CorpusID:259137341

[8] L. Yu, T. Jiang, and Y. Zou, "Distributed real-time energy management in data center microgrids," *IEEE Transactions on Smart Grid*, vol. 9, no. 4, pp. 3748–3762, 2018.

[9] S. Zhang, J. Lyu, W. Jin, H. Cheng, C. Li, and X. Wang, "Coordinated planning of multiple energy hubs considering the spatiotemporal load regulation of data centers," *IEEE Transactions on Power Systems*, vol. 39, no. 2, pp. 4193–4207, 2024.

[10] X. Xie, Q. Ling, P. Lu, W. Xu, and Z. Zhu, "Evacuate before too late: Distributed backup in inter-dc networks with progressive disasters," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 5, pp. 1058–1074, 2018.

[11] J. Sun, Y. Yan, Z. Wang, J. Ma, and Y. Wang, "Privacy-preserving coordinated operation of cross-enterprise data centers," *Applied Energy*, vol. 383, p. 125343, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626192500073X

[12] Z. Shi and A. Eryilmaz, "A zeroth-order admm algorithm for stochastic optimization over distributed processing networks," in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020, pp. 726–735.

[13] C. Feng, H. Xu, and B. Li, "An alternating direction method approach to cloud traffic management," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 8, pp. 2145–2158, 2017.

[14] G. Zhang, S. Zhang, W. Zhang, Z. Shen, and L. Wang, "Distributed energy management for multiple data centers with renewable resources and energy storages," *IEEE Transactions on Cloud Computing*, vol. 10, no. 4, pp. 2469–2480, 2022.

[15] D. Shen, J. Luo, F. Dong, X. Guo, C. Chen, K. Wang, and J. C. S. Lui, "Enabling distributed and optimal rdma resource sharing in large-scale data center networks: Modeling, analysis, and implementation," *IEEE/ACM Transactions on Networking*, vol. 31, no. 6, pp. 2745–2760, 2023.

[16] *Measuring the Space Economy: Estimating the Value of Economic Activities in and for Space*, Institute for Defense Analyses (IDA), 2020.

[17] *Space: The $1.8 Trillion Opportunity for Global Economic Growth*, World Economic Forum, Report April 2024.

[18] "Large constellations of satellites," *TECHNOLOGY ASSESSMENT. Report to Congressional Addressees. United States Government Accountability Office*, 2022.

[19] N. Pachler, J. J. Garau-Luis, E. F. Crawley, and B. G. Cameron, "A unified resource allocation framework and impact evaluation for ngso satellite constellations," *International Journal of Satellite Communications and Networking*, 2025. [Online]. Available: https://api.semanticscholar.org/CorpusID:275517650

[20] *Aerospace Forecasts Fiscal Years 2024-2044*, Federal Administration Avation (FAA), Report 2024.

[21] V. Aditya, D. Aswin, S. Dhaneesh, S. Chakravarthy, B. Kumar, and M. Venkadavarahan, "A review on air traffic flow management optimization: trends, challenges, and future directions," *Discover Sustainability*, vol. 5, 12 2024.

[22] Y. Chen, Y. Zhao, and Y. Wu, "Recent progress in air traffic flow management: A review," *Journal of Air Transport Management*, vol. 116, p. 102573, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0969699724000383

[23] "Automated air cargo operations market research and forecast," *NASA Report–20210015228*, April, 2021.

[24] A. P. Cohen, S. A. Shaheen, and E. M. Farrar, "Urban air mobility: History, ecosystem, market potential, and challenges," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 9, pp. 6074–6087, 2021.

[25] S. Shahriar Ahmed, G. Fountas, V. Lurkin, P. C. Anastasopoulos, Y. Zhang, M. Bierlaire, and F. Mannering, "The state of urban air mobility research: An assessment of challenges and opportunities," *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 2, pp. 1375–1394, 2025.

[26] C. Maheshwari, M. G. Mendoza, V. Tuck, P.-Y. Su, V. L. Qin, S. A. Seshia, H. Balakrishnan, and S. Sastry, “Privacy preserving mechanisms for coordinating airspace usage in advanced air mobility,” *ArXiv*, vol. abs/2411.03582, 2024. [Online]. Available: https://api.semanticscholar.org/CorpusID:273850432

[27] P.-Y. Su, C. Maheshwari, V. M. Tuck, and S. Sastry, “Incentive-compatible vertiport reservation in advanced air mobility: An auction-based approach,” in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 7720–7727.

[28] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021. [Online]. Available: https://doi.org/10.1038/s42254-021-00314-5

[29] S. Cuomo, V. S. D. Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, “Scientific machine learning through physics–informed neural networks: Where we are and what’s next,” *Journal of Scientific Computing*, vol. 92, 2022. [Online]. Available: https://api.semanticscholar.org/CorpusID:246015481

[30] M. Raissi, P. Perdikaris, and G. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0021999118307125

[31] M. Raissi, A. Yazdani, and G. E. Karniadakis, “Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations,” *Science*, vol. 367, no. 6481, pp. 1026–1030, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.aaw4741

[32] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, “Model-based deep learning,” *Proceedings of the IEEE*, vol. 111, pp. 465–499, 2020. [Online]. Available: https://api.semanticscholar.org/CorpusID:229180899

[33] K. Gregor and Y. LeCun, “Learning fast approximations of sparse coding.” in *ICML*, J. Fürnkranz and T. Joachims, Eds. Omnipress, 2010, pp. 399–406. [Online]. Available: http://dblp.uni-trier.de/db/conf/icml/icml2010.htmlGregorL10

[34] V. Monga, Y. Li, and Y. C. Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.

[35] F. S. Heaton, H., “Explainable ai via learning to optimize,” in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, vol. 13, no. 10103, 2023.

[36] X. Chen, J. Liu, and W. Yin, “Learning to optimize: A tutorial for continuous and mixed-integer optimization,” *Science China Mathematics*, vol. 67, no. 6, pp. 1191–1262, 2024. [Online]. Available: https://doi.org/10.1007/s11425-023-2293-3

[37] D. McKenzie, H. Heaton, Q. Li, S. Wu Fung, S. Osher, and W. Yin, "Three-operator splitting for learning to predict equilibria in convex games," *SIAM Journal on Mathematics of Data Science*, vol. 6, no. 3, pp. 627–648, 2024. [Online]. Available: https://doi.org/10.1137/22M1544531

[38] B. Tang, E. B. Khalil, and J. Drgoňa, "Learning to optimize for mixed-integer non-linear programming," 2024. [Online]. Available: https://arxiv.org/abs/2410.11061

[39] A. D. Saravanos, H. Kuperman, A. Oshin, A. T. Abdul, V. Pacelli, and E. Theodorou, "Deep distributed optimization for large-scale quadratic programming," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=hzuumhfYSO

[40] Y. Noah and N. Shlezinger, "Distributed learn-to-optimize: Limited communications optimization over networks via deep unfolded distributed admm," *IEEE Transactions on Mobile Computing*, vol. 24, no. 4, pp. 3012–3024, 2025.

[41] Y. He, Q. Shang, X. Huang, J. Liu, and K. Yuan, "A mathematics-inspired learning-to-optimize framework for decentralized optimization," 2024. [Online]. Available: https://arxiv.org/abs/2410.01700

[42] R. Sambharya and B. Stellato, "Learning algorithm hyperparameters for fast parametric convex optimization," 2024. [Online]. Available: https://arxiv.org/abs/2411.15717

[43] Y. Yang, Z. Huang, and D. Wipf, "Transformers from an optimization perspective," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, ser. NIPS '22. Red Hook, NY, USA: Curran Associates Inc., 2022.

[44] B. Hoover, Y. Liang, B. Pham, R. Panda, H. Strobelt, D. H. Chau, M. J. Zaki, and D. Krotov, "Energy transformer," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., 2023.

[45] Z. Fang, S. Du, Z. Cai, S. Lan, C. Wu, Y. Tan, and S. Wang, "Representation learning meets optimization-derived networks: From single-view to multi-view," *IEEE Transactions on Multimedia*, vol. 26, pp. 8889–8901, 2024.

[46] E. Duede, "Deep learning opacity in scientific discovery," *Philosophy of Science*, vol. 90, no. 5, p. 1089–1099, 2023.

[47] K. A. Creel, "Transparency in complex computational systems," *Philosophy of Science*, vol. 87, no. 4, p. 568–589, 2020.

[48] F. A. Csaszar, H. Ketkar, and H. Kim, "Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors,"

Strategy Science, vol. 9, no. 4, pp. 322–345, 2024. [Online]. Available: https://doi.org/10.1287/stsc.2024.0190

[49] E. Schmidt, H. Kissinger, and C. Mundie, *Genesis: Artificial Intelligence, Hope, and the Human Spirit.* John Murray Press, 2024.

[50] H. Kissinger, *Diplomacy*, ser. A Touchstone Book. Simon & Schuster, 1994.

[51] R. Bellman, *Dynamic Programming*, 1st ed. Princeton, NJ, USA: Princeton University Press, 1957.