
NATURE-INSPIRED AI (NIA)

A NOVEL APPROACH TO SOLVE FOR GENERALIZED AUTONOMOUS AGENTS

Gagandeep Reehal
Minus Zero Robotics Pvt Ltd, India

ABSTRACT

Existing autonomous vehicle systems face formidable obstacles in reliably navigating diverse driving environments, resulting in unpredictable behavior, disengagements, and even crashes. This paper critically examines the limitations of current approaches, where the inherent fragility due to the infinite variability of real-world scenarios and objects poses challenges that cannot be fully addressed through traditional training methods. This paper proposes a different approach to the underlying AI methodologies, rather than trying to adapt inadequate approach to address a safety-critical use case. We introduce nature-inspired AI as a new paradigm for developing generalized autonomous agents.

1 Introduction

While numerous companies are working on autonomous vehicles, most are pursuing iterative variations of the robotics paradigm - a traditional sense-plan-act workflow. This approach relies on neural network based understanding of the environment, combined with high-definition (HD) map localization. Neural networks are trained using billions of images to detect and identify drivable areas, road actors, etc. in a scene, enabling the system to reason about their positions and predict their behavior. This information is then matched in real-time with an HD map of the world to navigate.

However, this approach is inherently fragile because it is nearly impossible to have well representative data of every possible variation of the environments and non-ego agents encountered on the real world roads to train a network. To address this, many companies resort to equipping vehicles with an increasing number of costlier sensors like LiDARs, etc. While additional sensory input is presumed to be beneficial, it introduces conflicts in data and compounds problems like calibration errors, integration bottlenecks and hardware failures.

The inherent brittleness of the current AI methodologies used transcend into the following outcomes 1) lack of generalization across geographic areas and environment conditions, 2) difficult scalability across vehicle form factors or use-cases, and 3) extensive complexity and cost of infrastructure and sensors.

This paper explores the limitations of current approaches with the common analogy to AI, the human mind, and proposes the necessary shifts in school of thought to build generalized autonomous agents.

2 The Current Challenges

2.1 The Broken Loop

In a dynamic environment for decision making agents like AVs, the agents work in a closed loop system correcting their decision through continuous feedback to account for the variations in environment.

For the flow in traditional autonomous agents shown in Fig. 1 (a), the actor (vehicle) participates in a loop with its state estimation variables as feedback through LiDARs, IMU, etc. The participation of actors (vehicle) makes the agent very prone to differences in the vehicle dynamics like form factor, sensor orientations, etc. and from variations due to wear and tear of the vehicle. The necessity to re-tune the algorithms for such variation hinders it from being safe and actor-agnostic (vehicle-agnostic).

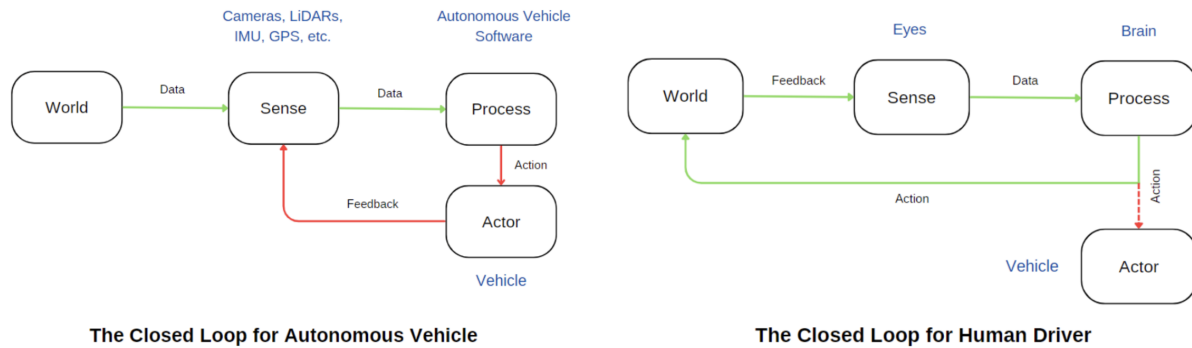


Figure 1: The closed loop for (a) autonomous vehicles and (b) a human driver

A similar analogy of the human mind, as shown in Fig. 1 (b), tends to have a more generalized approach to decision making, perceiving the feedback from the world instead of the actor. It tends to understand the variations in visual features over time to get feedback. By being able to utilise visual feedback without relying on actor, it easily adapts its actions to any kind of actor like a new vehicle or even a racing PC game.

2.2 How much data is unbiased data?

The underlying nature of current AI methodologies to be data-driven requires a brute force approach to learning with the most representative training data for the operating environment.

Relying on the number of training images or miles is a weak measure of completeness, as training on general events cannot compensate for missing rare ones. Since a neural network outputs confidence probabilities for every outcome encapsulated from data distribution, data bias arises from under-represented data samples. This distorts the representation of real-world frequencies, skew relevant features, over-generalize characteristics, and reflect prejudices, affecting the accuracy and fairness of decision-making.

On the contrary, human perception is driven by reasoning allowing it to develop better correlations and predict possible variations. For example, a human mind doesn't need a million instances of a bird to help it identify a bird whenever it encounters one.

2.3 The Barriers of Proportionality - Size, Compute and Data

The data bias being a major bottleneck, the general consensus to improve model performance is to, 1) get more training data, 2) increase the size of the model and 3) put in more GPUs training hours.

Collecting and labeling large-scale datasets is time-consuming and expensive. Large networks require extensive compute-resources for both training and inference, making them impractical for resource-constrained devices or real-time applications. On the opposite side, smaller models often do not have enough trainable parameters to ingest insights from the bigger datasets or make more generalized correlations, often causing overfitting or a significant performance drop.

2.4 Taxonomical Limitations of Representation

Are the current models able to extract sufficient context from the raw data for decision making? Reducing sensor data to symbolic definitions or labels created by and for human comprehension often does not generalize well. For example, a model might misinterpret a cone-shaped ice-cream truck to be an ice cream cone, within the hand-crafted taxonomy due to classic data distribution bias. Often a new variation may not confidently fit within restrictive identification classes of a supervised learner.

A generalized model requires a lean generalizable representation to account for large dissimilarities in visual appearance and motion of the same non-ego agent. A human mind is extremely adept at avoiding misclassification by being able

to use both semantic labels and knowledge-based reasoning to perceive situations and threats. Relative speed, size, randomness, density, scene awareness and proximity all alert us to physical risks.

Ecological perception is more than just sensing and classification. It also requires context, judgment, and reasoning.

2.5 Complicated vs. Complex Systems

The traditional approaches seem not to have paid attention to differences between complex and complicated systems. In particular, complicated systems: have many variables, but they operate in patterned ways. It allows accurate predictions with statistical modeling - a specific input will produce a predictable output. This happens in environments with a single intelligent agent, where every non-ego agent follows a certain set of rules.

The roads — encompassing the ever-changing interrelationships of vehicles, pedestrians and environment — is complex. Complex systems tend to be multi-agent environments where each agent is an intelligent subsystem making independent decisions with their interactions with other agents and environments. Every possible interaction across its physical and behavioral properties between N such agents over a time duration t is a new variable with its own uncertainty, hampering predictability of a system. The conflict of any two intelligent agents competing for an exhaustive resource in the environment impacts the non-participating agents.

Existing machine learning techniques based on brute-force training do not solve the reality of ambiguous, uncertain, or unpredictable scenarios common on public roads. An attempt to have perfect measurements of every factor in a system does not perfectly predict outcomes of that system statistically.

On a similar line a human mind at all stages of its evolution is excellent in developing ‘social contracts’ capable of understanding magnitude of threat allowing safe response to both tangible and intangible stimuli with context-aware reasoning. This opens a critical opportunity to redefine how we understand, represent and predict the environment.

2.6 The Faulty Frame of Reference

The decision-making is often a tricky or computationally heavy task. Traditional agents simplify real-time decisions by transferring part of the representation offline with clean, curated data like high-fidelity maps. To enable this, the agent requires its state estimation in the global frame of reference, and also represents every other agent relatively variable to it in the global coordinate frame or absolute metric units.

While this brings maintenance complexity as the world changes, any error in ego and non-ego state estimation due to calibration issues, noisy data or environmental factors gets accumulated in the global frame with little opportunity to adapt or correct. To prevent this, one adds costlier sensors to decrease the noise and uncertainty in the fused data.

The human mind on the contrary perceives and decides in an ego-centric frame of reference. Having to not rely on allocentric awareness, this allows it to plan in a relative spatial and temporal context, and make corrections even with lack of absolute metric measurements. This helps it adapt across different environments with little prior context of the new surroundings.

3 The New Paradigm - Nature-inspired AI

The Beauty of Human Mind

Our critical examinations for developing generalized autonomous agents reveal a big gap between the traditional approaches and the human mind. It becomes essential to explore the cognitive side of decision-making.

In circumstances which are unpredictable, uncertain, or unknown, a human mind is a more capable decision-maker than current AI algorithms. There has been an interesting incident demonstrating this when *Capt. Chesley “Sully” Sullenberger* landed Flight 1549 on the Hudson River following dual engine failure resulting from bird strike. It would have been beyond any probabilistic computer algorithm.

We develop a “social contract” - physics-informed and context-aware algorithmic relationships, inspired from neuronal computations inside the brain contributing to attention, perception, episodic memory, and predictive decision-making.

What might seem a general judgment, reasoning, and uncertainty for an average human currently leaves a robotic agent confused.

How do you define a generalized autonomous agent?

Any suite of algorithms exhibiting properties of a generalized autonomous agent has to be:




Figure 2: An image drawn (a) by recalling from memory and (b) by referring to the object physically

- **Resource aware:** minimizing resource usage (like computations) as an objective function
- **Data aware:** representing the associations and correlations beyond limited taxonomy
- **Context aware:** reasoning with understanding of domain it is operating in
- **Interaction aware:** interacting with other intelligent agents and being aware of outcomes of potential interactions
- **Spatio-temporal aware:** understanding the behavior of the environment in relative space and time.
- **Outcome aware:** optimizing its understanding from deviations in expected outcome

3.1 Representation and Association

3.1.1 Physics-aware, Context-aware Algorithms

In general, DNNs can approximate any high-dimensional function given that sufficient training samples are available. But with the lack of physical characteristics underlying the problem, the level of approximation accuracy is still heavily dependent on boundaries of the problem geometry. An algorithm in a critical use-case like AVs should be able to express generalizability by learning coherent representations of their world and interpretable explanations of its dynamics like motion, entropy and semantic context.

We propose physics-aware system design where a learning algorithm with some knowledge about the physical characteristics of the problem and even sparse training data would be able to find an optimal solution with high fidelity. The prior awareness of physical laws and the semantic context during training will increase the correctness of the function approximation, even with a low amount of training examples. This would show superior generalizability and interpretability compared to larger black-box learning systems.

3.1.2 Loosely Supervised Networks with Sparse Representations

If you are asked to draw a dollar note by recalling from memory, you draw something similar to Fig. 2 (a), but if you had a physical currency note to refer to, you would draw something on the lines of Fig. 2 (b).

The difference between the two pictures visualizing something is far less accurate or sparse than seeing something in its presence. This is why we're much better at recognising than recalling. Our reception mechanisms don't create accurate representations of visual stimuli.

We propose loosely supervised networks, guided by general semantic scenes of the environments, that can understand (without explicit identification) any agent or environment variable with unsupervised representations of threats, without prior experience/training of that particular variation. Instead of environment information being limited to distinct clusters in a data, the learnable parameters should allow contextual correlations between clusters that could be dynamically representable, and generalized across all domains.

A generalized architecture should be extremely sparse and multi-modal to allow for learning diverse variations and quick associations for identification.

3.1.3 Network-of-Networks with Contrasting Bias

For a task driven model, the neural network is often supervised so the entire network has to be trained in entirety. This is necessary for weights to converge to a pre-defined set of outcomes, so every part of the network is affected by bias.

To avoid this, we further propose a network-of-networks approach. With sparse physics-based and context-based representations and a loosely supervised network, the multi-modality of data would allow us to split the network into multiple sub-networks which could be trained separately which wasn't feasible earlier. Having that ability could be utilized to train sub-networks across different data distributions having varied bias which can accommodate for the discrepancies of the other. This would allow one to have a better control over reflection of the overall bias in the decisions of the network as a whole.

In other words, a model could outsource a task to another model, and various features representations like classification, scene understanding, motion awareness, depth awareness, etc. could be covered by different networks, and new layers could emerge as the "glue" stitching the entire network together.

3.2 Predictive Awareness

In an environment, any decision d (like the planned path) is essentially a non-linear optimization function that maximizes or minimizes an output y from m input variables that exert their relative influence on the output given n constraints or costs. Many decisions require an estimate of the future state of the environment, which is complex due to the dynamic scene. It becomes important to have a predictive awareness of the environment and actors to optimal decision making.

3.2.1 Space-Time Awareness

One of the unique capabilities for the human mind is computation of complex visual motion patterns as we move through the environment, resulting in spatio-temporally deviations of features points that can be corresponded with the semantic representations of the environment. The brain constructs structured sequences to infer trajectories through the world with exploratory experience. It should be noted that exploratory experience is not a passive representation or perception of space but an active process of construction of relationships contributing to episodic memory.

However, traditionally practiced map-based navigation based on semantic understanding does not represent human decision making which accounts for both semantic and episodic memory. While current agents have limited context windows due to only spatial representations of the environment, we propose generalized representation of the world in a spatio-temporal context that afford the ability to correlate past experiences to potential future interactions, and ultimately use it to direct action in the present.

3.2.2 Multi-agent Environment Interactions

Most dominant learning algorithms rely on single-agent optimization-based paradigms that search for explanatory insights into the collective behavior of agents (which don't necessarily need to be "intelligent") obeying simple patterns. But the real world is a multi-agent environment. We propose designing systems that can account for the intrinsic interactions between its multiple components.

While multi-agent environments vary across dynamics, discreteness, episodicity, and dimensionality, the agents in such a system exhibit these characteristics:

- **Autonomy:** agents at least partially independent, self-aware, autonomous
- **Local views:** no agent has a full global view, or the system is too complex for an agent to exploit such knowledge
- **Decentralization:** no agent is designated as controlling

This allows a multi-agent system to manifest complex behaviors for any interaction involving more than two agents, which impacts potential outcomes of the ego agent.

For an autonomous agent to generalize across dynamics of complex systems, it would require a careful design of the correlations of behavioral properties like determinism, entropy, size, etc. The traditional agents tend to smooth out fluctuations by modeling the average behavior across all the agents. However, under certain conditions, the fluctuations can be amplified when the system becomes unstable to larger perturbations. Therefore, simulating a population of intelligent agents using dynamic parameters computed from averaging over the behavior of all agents is an oversimplification. The behavior of each individual is nonlinear and highly dependent on the dynamics of its

neighbors, which are non-deterministic due to the randomness in the behavior of each agent. This would serve as dynamic constraints for optimizing the functions for safe action given the state of the environment.

3.3 Adaptation

Infinite Correction Loop - a truly autonomous agent doesn't make decision but it corrects them in an infinite loop to adapt or survive.

3.3.1 Local Frame of Reference

Intelligent autonomous agents are capable of auto-adapting their next action w.r.t perceived duration between visual stimuli, governed by a localized temporal rate due to the spatially invariant temporal statistics of natural scenes. So we propose an ego-centric frame of reference as an effective strategy to generalize across different environments with little prior context of the new environment. This is vital for making the closed feedback loop of an autonomous agent independent of the actor as shown in Fig. 1 (a).

3.3.2 Self-adaptive control systems

While learned vehicle control agents often show great performance in offline testing and simulations, this considerably degrades during real world driving. By eliminating the actor as a part of closed loop, we propose vision-only control systems that can exploit both the temporal nature and physics-informed context-aware representation of environment, enabling them to auto-adapt to disturbances in the actor's dynamics. This will allow autonomous agent scale to any kind of vehicle (actor).

3.3.3 Continual Learning - Reward predictions rather than outcomes

The inherent nature of traditional policy based learning methods to improve predictions by maximizing reward is adept for end-to-end black box models, but it is often limited by hand-crafted rewards for outcomes which may not optimise the right segment of the network.

We propose a continual learning method similar to a cognitive architecture that instead penalises magnitude of the deviations between the predicted and the actual outcome. For decisions governed by certain set of constraints, the system can reverse engineer the representative variables that contributed to the said constraints, allowing it to optimally tune the weights with explainability.

4 The Enablers

The Pseudo-data Approach The classical problems of data bias can be solved by using generative data, where multi-modal prompting can allow creating diverse variations and better simulation. Earlier synthetic data poised problems in domain adaption due to gap in realism, but physics-awareness, contextual understanding, and sparse generalized representations would limit the admissible information allowing it to generalize well.

The Self Supervised Way Ahead Driving is sufficiently complex that hand-crafted abstraction layers and features are unable to adequately model the problem. Often with task-focused networks, the model tries to correlate a feature in a brute-force manner to a given set of outcomes specific to a particular problem. In doing so, the backbones which are capable of finding more correlations are often restricted to be biased towards the certain hard-crafted outcomes. This hinders it from adapting its knowledge to a different environment. We believe that self supervised or loosely supervised training method will be vital to develop better associations and correlations guided by context awareness.

Teaching for Edge Building for edge will require physics-informed and contextual reasoning based models to be able to fit enough information in the said form-factor. We emphasize student-teacher methods as a part of the pipeline - with loosely supervised networks for extracting learnable representations, and self supervised physics-aware or context-aware networks learning distilled knowledge for faster convergence and diverse correlations.

5 Conclusion

While the notion of "more sensors, more data, and more powerful models" may seem appealing, the reality is that increased hardware and software complexity can undermine safety and reliability. Moreover, the cost of the system

has a significant impact on its feasibility and applicability to every use-case, which is crucial for achieving substantial improvements in societal safety.

We position nature-inspired AI as the scalable approach to generalized autonomous agents. With design derived from neuronal computations inside the brain contributing to attention, association, threat response, episodic memory, and predictive decision-making, it will help learning algorithms self-adapt across use-case and operational domains.

We propose this shift in paradigm not just for newer research problems in this domain, but emphasize defining justified metrics that will allow verification and validations for such AI-driven design working on abstract representations like our nervous systems, rather than the metric ground truth.

References

- [1] David Burr and Concetta Morrone. Time perception: Space–time in the brain. *Current Biology*, 16(5):R171–R173, 2006.
- [2] Soumya Sudhakar, Vivienne Sze, and Sertac Karaman. Data centers on wheels: Emissions from computing onboard autonomous vehicles. *IEEE Micro*, 43(1):29–39, 2023.
- [3] Robert Epstein. The empty brain. *Aeon*, May 2016.
- [4] Rachel Cooke Chesley Sullenberger: an old-fashioned kind of hero *The Gaurdian*, 2016
- [5] Mohamed Akrouf, Amal Feriani, and Bob McLeod. Dynamic noises of multi-agent environments can improve generalization: Agent-based models meets reinforcement learning, 2022.
- [6] Priyamvada Jadaun. Adaptive cognition implemented with a context-aware and flexible neuron for next-generation artificial intelligence. *PNAS nexus*, 1(5), 2022.
- [7] Rishi Bommasani et al. On the opportunities and risks of foundation models, 2022.
- [8] Patrick Fraser, Ricard Solé, and Gemma De las Cuevas. Why can the brain (and not a computer) make sense of the liar paradox? *Frontiers in Ecology and Evolution*, 9, 2021.
- [9] Ramin Hasani, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. Liquid time-constant networks, 2020.
- [10] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Comput. Surv.*, 54(6), jul 2021.
- [11] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [12] Zhongkai Hao, Songming Liu, Yichi Zhang, Chengyang Ying, Yao Feng, Hang Su, and Jun Zhu. Physics-informed machine learning: A survey on problems, methods and applications, 2023.
- [13] Kevrekidis I.G. Lu L. Karniadakis, G.E. et al. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.
- [14] Ashish Vaswani, Noam Shazeer, et al. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [15] Harald Atmanspacher. Quantum Approaches to Consciousness. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2020 edition, 2020.
- [16] Anil Ananthaswamy. Artificial neural nets finally yield clues to how brains learn. *Quanta Magazine*, 2021.
- [17] Brady Zhou, Philipp Krähenbühl, and Vladlen Koltun. Does computer vision matter for action? *Science Robotics*, 4(30):eaaw6661, 2019.
- [18] Brad Plumer. 5 big challenges that self-driving cars still have to overcome. *Vox*, 2016.
- [19] Colin Barnden. Autonomous vehicle: Complex or complicated? *EET Asia*, 2020.
- [20] Tom Brown, Benjamin Mann, et al. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [21] Brian Paden, Michal Čáp, Sze Zheng Yong, Dmitry Yershov, and Emilio Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1):33–55, 2016.

-
- [22] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2255–2264, Los Alamitos, CA, USA, jun 2018. IEEE Computer Society.
 - [23] Alexey Dosovitskiy, Lucas Beyer, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
 - [24] Constantin Hubmann, Marvin Becker, et al. Decision making for autonomous driving considering interaction and uncertain prediction of surrounding vehicles. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1671–1678, 2017.