NeuroSAN+NeuroAI: AI-assisted Decision-making through a Synergy of Technologies

Risto Miikkulainen[1,2], **Dan Fink**[1], **Olivier Francon**[1], **Babak Hodjat**[1],
Noravee Kanchanavatee[1], **Elliot Meyerson**[1], **Xin Qiu**[1], **Darren Sargent**[1],
Hormoz Shahrzad[1,2], **Deepak Singh**[1] **Jean Celestin Yamegni Noubeyo**[1] **Daniel Young**[1,2],
[1]Cognizant AI Lab, San Francisco, CA
[2]The University of Texas at Austin

Cognizant AI Lab Technical Report 2025-01

Abstract

Several technologies have emerged recently to support decision-making in organizations including: LLMS that allow defining opportunities, obtaining data, and building user interfaces; machine learning methods for building predictors with a variety of data types; population-based search methods that discover good decision strategies; methods for estimating uncertainty in the predictions and in LLM output; and systems for coordinating multiple agents to integrate knowledge into comprehensive answers. This paper reviews a system called NeuroSAN+NeuroAI that brings them together into a synergetic approach, with applications in business, engineering, healthcare, education, and society in general.

1 Introduction

Much of human activity in society can be seen as decision-making: choosing a course of action given the current problem and general knowledge about the domain, in order to optimize a set of goals. Decision-making is hard because there are often multiple stakeholders and multiple conflicting objectives, each problem may be different from those seen before, and information about the problem and the domain may be incomplete or uncertain. For instance in the medical domain, there are patient, caregiver, hospital, insurance, legal, economic, and societal considerations; the tests and diagnoses may be unreliable, the patient history only partially known, medical knowledge incomplete, and the treatment's effectiveness in conflict with its cost and side effects. Many domains have grown so complex that it is difficult for human decision makers to understand the consequences of their actions, let alone arrive at optimal decisions reliably.

On the other hand, AI technologies have recently improved substantially, and are now in a position to improve complex decision making in the real world. The main new perspective is multi-agent decision making. Orchestration of multiple agents makes it possible build teams of AI agents that collaborate to solve complex tasks. The agents can be diverse, employing multiple technologies, and as those technologies get better, they can be integrated into the larger system.

Several frameworks have emerged for multi-agent orchestration, enabling autonomous or collaborative AI agents to plan, reason, and act in concert [2, 3, 9]. These frameworks aim to overcome the limitations of relying on a single large language model (LLM) or agent for complex tasks. Instead of expecting one AI to be an all-knowing expert, multi-agent orchestrators break problems into smaller pieces and assign them to specialized agents that can collaborate.

While there is significant value in orchestrating the agents well, ultimately such systems depend on the quality of the individual agents. The agents in such a system are typically LLMs prompted or fine-tuned to act as experts of various kinds. However, it is also possible to include more specialized agents, including those that are trained with machine learning methods on specialized data to predict

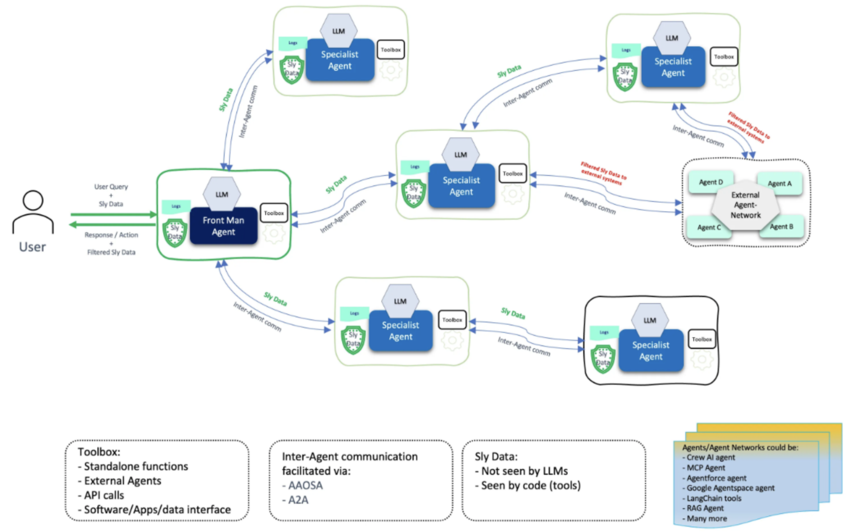

Figure 1: **Data-driven definition of a NeuroSAN multiagent system**. The agents interact with each other in the defined network using the AAOSA protocol, where an agent breaks down and delegates tasks or actions to its down-chain agents, and the coordinating agent (the "FrontAgent") consolidates the response and sends it back to the user. An agent network can be hierarchical, sequential, or DAG-oriented; each node can be a single agent, another Network of Agents, or a tool constructed through traditional programming. In this manner, complex tasks can be broken down to manageable challenges, resulting in AI systems that are more powerful, robust, easier to deploy and maintain, and safer.

the outcomes, and to recommend actions that lead to best possible such outcomes. As part of such a system it is important to estimate how reliable the predictions are, as well as how reliable the LLM outputs are.

In NeuroSAN+NeuroAI, these previously separate technologies are brought together to support decision-making in a synergetic manner. This paper reviews these elements and gives examples and suggestions on how real-world applications can be built based on them.

2 Orchestrating Multiple Agents with NeuroSAN

NeuroSAN [Figure 1; 13] is a recently proposed framework for multiagent orchestration, different from most other such frameworks in four ways: First, it is a data-driven approach to multi-agent orchestration, allowing entire agent networks to be defined declaratively rather than through hard-coded logic. It allows users — from machine learning engineers to domain experts — to quickly build multi-agent applications without extensive coding, using declarative configuration files. Second, agent interactions in NeuroSAN are managed through the AAOSA [Adaptive Agent-Oriented Software Architecture; 6, 13] communication protocol, making it possible to solve complex tasks by dynamically delegating subtasks to LLM-powered agents. Third, NeuroSAN treats non-LLM tools as agents, expanding the capabilities of the AI system in a uniform manner. Fourth, NeuroSAN is an open-source Python library that can run either embedded in the application or as a standalone service.

In NeuroSAN, multi-agent networks are specified entirely in a *data configuration* file rather than in code. NeuroSAN uses an editable HOCON format to define agents, their relationships, and behaviors. In practice, this approach allows subject matter experts or architects — not just programmers — to design complex agent networks by describing what each agent does and how they connect. The NeuroSAN engine then interprets this config at runtime to spin up the agents accordingly. This data-driven approach lowers the barrier to creating and managing multi-agent systems.

At the top level, NeuroSAN provides a flexible orchestration loop that manages inter-agent communication. Agents converse in natural language (prompts) but follow a structured protocol to decide how to route queries. NeuroSAN agents adhere to an *interaction pattern called AAOSA*, developed for multi-agent systems in the 1990s and adapted to LLM-based AI in NeuroSAN [6, 13]). AAOSA essentially gives each agent a guideline on how to respond to an inquiry. An agent will check if it is the appropriate one to answer; if not, it knows how to call on other agents for help. This inter-agent communication protocol ensures that queries get dynamically passed to the agent best suited to handling them, without a human having to manually route the questions. It is a decentralized decision-making mechanism baked into the agents' behavior. One agent (a coordinator or a "front agent") acts as the main interface to the user, and can delegate subtasks to downstream specialist agents. This approach allows complex tasks to be broken into a hierarchy of smaller tasks. NeuroSAN natively supports such hierarchies and even nested agent networks. In fact, one agent network can call another sub-network as part of its process, enabling the reuse and modular design of agent teams.

Real-world tasks often involve actions that pure LLMs cannot perform, such as querying a database, doing heavy calculations, or calling an external API. NeuroSAN addresses this need by allowing agents to *invoke coded tools* — essentially custom Python functions or classes that do deterministic operations outside the LLM's domain. These tools are integrated seamlessly into the agent networks. An important distinction is that NeuroSAN treats these tools as agents: an agent can ask a tool to do something and get back results as part of its reasoning loop. This approach opens up a hybrid approach to problem solving — part natural language (LLM reasoning) and part traditional programming.

Multiagent systems are a natural next step from agents in the evolution of AI architectures. Such systems can *have better abilities* than single-agent systems: large problems can be broken down to manageable size through divide-and-conquer, and multiple perspectives can be brought together in complex tasks where a single line of thought may get stuck. General agents can be combined with specific agents, and even proprietary ones. The system can be improved and maintained in a modular manner where new agents are included over time. Multiagent systems also offer opportunities to *improve AI safety*, by providing redundancy, and also by including agents that are experts at evaluating performance and catching and repairing errors. Human experts can also be included in such networks.

3 Discovering Decision Strategies with NeuroAI

Complementing the multiagent orchestration, machine learning of individual agents can play a large role in improving decision-making. Multiple technologies are brought together in NeuroAI to make this approach effective (Figure 2).

First, many organizations now routinely collect data on their operations, making it possible to take advantage of these technologies. The most immediate role for them is in *predictive analytics* [7]. Data science methods can be used to model datasets in order to predict what will happen in various scenarios in the future.

Second, extending their obvious role as chatbots, LLMs can be used in this role as well: they can be prompted with a situation and asked to generate possible future scenarios. The difference from data science is that LLMs predictions are largely based on general knowledge. Thus, they are broadly applicable, but it is difficult to tell how reliable they are. In either case, the predictions are consumed by human decision-makers, who also bring their own expertise, as well as biases, into the decision-making process.

Third, predictive intelligence is only part of an effective agent. Other technologies allow taking advantage of *prescriptive intelligence* as well [5]. This means using AI to generate recommendations for decisions to reach optimal outcomes. In essence, the human decision-maker can explore potential tradeoffs of conflicting objectives, and receive the best implementations for each tradeoff. For instance, AI can suggest the best treatment for a given tolerance for side effects and a budget. The decision-maker can explore alternatives and see their effect on the goals, arriving at the best solution supported by the data and his/her expertise and preferences.

Prescriptive technologies discover good decision policies. They are distinctly different from predictive techniques and LLMs. Optimal decisions are not known, so gradient-based techniques cannot be used; instead, policies are generated through trial-and-error search such as evolutionary computation or reinforcement learning. The *search is intelligent*, guided by fitness (or reinforcement) information. It takes advantage of existing solutions through recombination, and the generation of new ideas through mutation. Evolutionary search is particularly powerful in this role because it is population-based, making it possible to explore solutions more broadly than techniques based on the refinement of

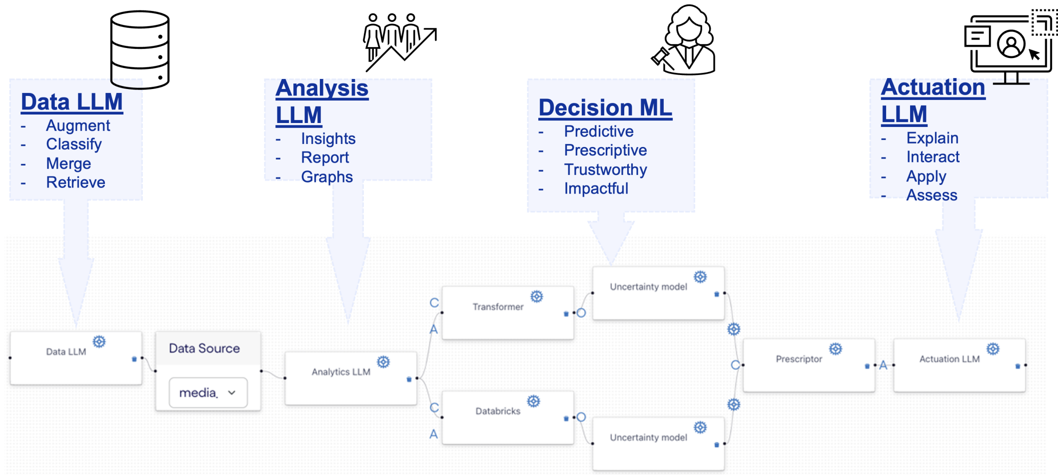

Figure 2: **The NeuroAI integration of multiple technologies into a data-based decision-making agent.** The core of the system is a prescriptive model that discovers optimal decision strategies, evaluated with a predictive model during search, and with confidence estimated through uncertainty models for each objective, such as performance and cost. LLMs are used to augment and integrate datasets, provide insights into the data, and to explain and actuate the decisions. As part of a system of agents in NeuroSAN, such agents can be built to support human decision-making in business, engineering, healthcare, education, and society.

single solutions. It can therefore discover *creative solutions* that would be difficult for other methods and even humans to find.

Although prescriptive technologies have existed for a long time, LLMs can now be used to take advantage of them on a larger scale, broader scope, and improved transparency and trustworthiness in decision-making. The idea is to use LLMs to *empower the data and the decisions*. That is, LLM agents can be used to make the datasets more useful by cleaning up errors, filling in missing data, integrating multiple datasets, and extending the data with insights and common sense. They may be used to transform the recommendations into human-readable form, and to implement them through an actuator interface.

To build a powerful decision-making agent, NeuroAI integrates LLMs and prescriptive modeling with a predictive model. A data science or GenAI model (a neural network, random forest, or transformer) serves as a *surrogate* for the world, making it fast and safe to evaluate the decision-strategy candidates. For instance, a recommended treatment is evaluated through the predictive model to see how good it is, instead of having to apply it to an actual patient.

Fourth, uncertainty estimation, rule-set evolution, and LLM technologies are integrated into NeuroAI to make the results trustworthy. The predictive model is paired with a Gaussian Process model in order to estimate the *uncertainty* in the predictions. The decisions are then made explainable: The prescriptive model is implemented as a set of rules that are *transparent and interpretable*. These rules are further interpreted with LLMs in natural language, making it easier for the decision-maker to trust the recommendations.

As a concrete example, a NeuroAI agent was created for recommending medication for a diabetes patient in a hospital, in order to release them rapidly with a lower risk of readmission [12]. LLMs were used to augment the data and glean insights from it; a random forest was trained to predict the two outcomes given a patient description and prescribed medication; a Gaussian Process based uncertainty model [11] was trained to estimate confidence in these predictions; a rule set was evolved to recommend medication towards fast release and low risk of readmission; LLMs were used to express these rules in natural language, and to make the prescription order once the physician approves it.

4 Examples, Challenges, and Opportunities

Several real-world applications have already been built with NeuroSAN and NeuroAI. They include land-use optimization to minimize carbon emissions with minimal change (from ITU/UN's Project Resilience; 16, 4), optimizing non-pharmaceutical interventions in the COVID-19 pandemic (from the XPRIZE competition; 10, 15), and optimizing the design of an artificial pancreas [8]. Practical business applications abound as well, for instance optimizing mixed marketing strategies, pricing optimization for maximal revenue and margin, personalizing training exercises for education and rehabilitation, and discovering software testing strategies to maximize accuracy and minimize cost and time.

AI technologies are still new, and there are known challenges: they may exhibit unwanted biases, it is difficult to separate fact from fiction in their output, their training may not include the latest data, and their responses may be generic rather than specific. However, as part of the orchestration, it is possible to build safeguards to identify potential issues, mitigate them, and alert human operators when needed. As such, it is already possible to take advantage of these systems, and even more so in the future as they improve.

NeuroSAN is intended as a common framework for the community, making it easy to deploy. It is a Python library that can be integrated into an application programmatically; it can also be run as a centralized agent orchestration service so that it can be shared by many clients. It is an open-source project [1], allowing the community to extend it and adapt it to new uses. Similarly, NeuroAI is available to the community through Project Resilience [4]. In that framework, developers can bring in their datasets and use the technology to build AI-For-Good applications that address e.g. the 17 UN Sustainable Development Goals [14].

5 Conclusion

Thus, effective AI-driven decision making is possible today. By bringing together multiple technologies, systems such as NeuroSAN+NeuroAI can be built that combine data collection and analysis, prediction, prescription, uncertainty estimation, explainability, and practical user interfaces. Multiagent coordination then makes it possible to bring together multiple areas of expertise, stakeholders, and conflicting objectives to address complex problems, including those in business, engineering, healthcare, education, and society.

References

[1] Cognizant AI Lab. NeuroSAN data-driven agents. https://github.com/cognizant-ai-lab/neuro-san, 2025. Retrieved 8/31/25.

[2] CrewAI. The leading multi-agent platform. https://www.crewai.com, 2025. Retrieved 7/29/2025.

[3] Abul Ehtesham, Aditi Singh, Gaurav Kumar Gupta, and Saket Kumar. A survey of agent interoperability protocols: Model context protocol (MCP), agent communication protocol (ACP), agent-to-agent protocol (A2A), and agent network protocol (ANP). *arXiv:2505.02279*, 2025.

[4] Olivier Francon. Project resilience platform. https://github.com/Project-Resilience/platform, 2025. Retrieved 8/31/25.

[5] Olivier Francon, Santiago Gonzalez, Babak Hodjat, Elliot Meyerson, Risto Miikkulainen, Xin Qiu, and Hormoz Shahrzad. Effective reinforcement learning through evolutionary surrogate-assisted prescription. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2020)*, pages 814–822, 2020.

[6] Babak Hodjat and Makoto Amamiya. Adaptive interaction using the adaptive agent oriented software architecture (AAOSA). *arXiv:cs/9812015*, 1998.

[7] John D. Kelleher, Brian MacNamee, and Aoife D'Arcy. *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies*. MIT Press, Cambridge, MA, 2015.

[8] Ashok Khanna, Olivier Francon, and Risto Miikkulainen. Optimizing the design of an artificial pancreas to improve diabetes management. In *Proceedings of the IJCAI workshop on Advanced Neural Systems for Next-Generation Biomedical Intelligence*, 2025.

[9] LangChain. Langgraph. https://langchain-ai.github.io/langgraph/, 2025. Retrieved 7/29/2025.

[10] Risto Miikkulainen, Olivier Francon, Elliot Meyerson, Xin Qiu, Darren Sargent, Elisa Canzani, and Babak Hodjat. From prediction to prescription: Evolutionary optimization of non-pharmaceutical interventions in the COVID-19 pandemic. *IEEE Transactions on Evolutionary Computation*, 25:386–401, 2021.

[11] Xin Qiu, Elliot Meyerson, and Risto Miikkulainen. Quantifying point-prediction uncertainty in neural networks via residual estimation with an I/O kernel. In *Proceedings of the International Conference on Learning Representations*, pages 2146–2180, 2020.

[12] Hormoz Shahrzad, Babak Hodjat, and Risto Miikkulainen. EVOTER: Evolution of transparent explainable rule-sets. *ACM Transactions on Evolutionary Learning and Optimization*, pages 1–30, 2024. Vol 5, Issue 2, Article 11.

[13] Deepak Singh. NeuroSAN is all you need — A data-driven multi-agent orchestration framework. https://medium.com/@evolutionmlmail/neuro-san-is-all-you-need-a-data-driven-multi-agent-orchestration-framework-563fbd31a735, 2025. Medium; Retrieved 7/29/2025.

[14] United Nations. The 17 goals. https://sdgs.un.org/goals, 2025. Retrieved 8/31/25.

[15] XPRIZE. Pandemic response challenge. https://www.xprize.org/challenge/pandemicresponse, 2023. Retrieved 8/31/2025.

[16] Daniel Young, Olivier Francon, Elliot Meyerson, Clemens Schwingshackl, Jacob Bieker, Hugo Cunha, Babak Hodjat, and Risto Miikkulainen. Discovering effective policies for land-use planning with neuroevolution. *Environmental Data Science*, 4:e30, 2025.