Anticipatory Intelligence Using Polyagents

A perennial concern of analysts and the decision-makers they support is anticipating the future. Conventional forecasting technologies are too slow, too brittle, and too labor-intensive. For the past two decades, Axon's scientists have been reverse-engineering the mechanisms used by natural swarms (both animal and human) to develop a revolutionary forecasting technology, the polyagent, that has repeatedly outperformed both humans and competing algorithms. Axon has embedded these mechanisms in its AXON:OS platform, which now makes industrial-strength polyagent forecasting available to a wide range of applications, including cyber-security, insider threat, geopolitical dynamics, combat prediction, and tracking persons of interest.

This white paper summarizes five challenges of anticipatory intelligence, introduces four technical concepts that allow AXON:OS to address them, and summarizes the performance of these mechanisms in our previous, research applications.

Five Challenges of Anticipatory Intelligence

"Prediction is hard, especially about the future." At least five challenges cloud our crystal ball.

Complexity.—The future is driven by many factors that interact in complex ways. Individual people and institutions ("actors") are constrained by external influences (such as terrain, the weather, and the state of the economy), as well as by one another. They are constrained not only by the laws of physics, but also by cognitive forces such as emotions and goals, which are often conflicting, and may be physically unattainable. Predictive mechanisms such as differential equations or Bayesian networks are effective in physical systems, but crafting such formalisms for human behavior is extremely labor intensive, when it is possible at all. A common approach is to neglect the complexities in order to force the problem into a closed mathematical form, leading to a brittle solution.

Incomplete information.—We often do not have all the information we would like about the beliefs, desires, and plans of the people and organizations whose behavior we wish to anticipate. This challenge increases the labor demands on analysts to make realistic estimates of missing knowledge, and renders our models brittle if our guesses are wrong.

Data.—How the future unfolds depends not only on the constraints that drive the actors, but also on the specific state of the world that they encounter. But identifying the relevant data to take into account is slow and labor-intensive, because of the volume of data we already have, the velocity with which new data arrives, and the variety of data sources that we could query.

Emergence.—Imagine that we could solve the previous challenges: that we had complete knowledge of each actor's cognitive state, that we could enumerate all of the constraints on each actor, and that we could identify all of the relevant data. We would still be frustrated by the phenomenon of emergence, which is the tendency of a collection of interacting actors to exhibit behavior that is not explicitly represented in the behavior of the individual actors. A model that simply aggregates the predicted behaviors of individual agents using statistical tools such as means and standard deviations will be extremely brittle, yet many tools (such as the differential equations underlying System Dynamics models) are based on population averages.

Probabilities.—Given the uncertainty in forecasting, a single forecast will be dangerously brittle. We should view any forecast, not as the answer to the problem, but as a possible answer among many others. What we really need is a collection of forecasts so that we can estimate the probability of any given outcome. The challenge here is that generating multiple forecasts using conventional tools is slow, and may not yield results in the time frame that decision makers require.

The Technical Road to Anticipatory Intelligence

Decades of research into these challenges have led Axon to identify four main mechanisms, which are now available for robust deployment in the AXON:OS platform.

Agent Orientation.—Early commercial software was organized around a functional decomposition of the problem. Such an architecture requires closed-form representations of the complex constraints that drive behavior. A major breakthrough is taking an agent-oriented view of the problem, decomposing it not by overall functions, but by the individual actors, each represented by a software agent. Each agent runs a program that seeks to imitate how the corresponding actor would respond to various environmental conditions (including other agents). This decomposition of the problem allows the system to represent individual differences among actors, rather than relying on population averages that hide the impact of emergence.

Mark Slonecker, President
mark.slonecker@axonai.com

Dr. Sven A. Brueckner, Chief Scientist
sven.brueckner@axonai.com

Agents are usually localized in some structure (such as a geospatial map, a social network, a narrative-causal graph, or a model of goals and sub-goals). Conventional forecasting methods pulls these structures inside a computer algorithm, limiting the size and complexity of the constraints that can be taken into account. An agent approach inverts the problem, distributing the algorithm (in the form of the agents) over the structure. As a result, the processing can scale to handle problems of arbitrary size and complexity, drawing on distributed hardware if needed.

Synthetic Evolution.—Asking analysts to estimate the things we don't know about actors is slow and brittle. In nature, populations of agents learn to respond to their environment by a process of trial and error, reinforced by the survival of those who succeed and the disappearance of those who do not. This process can be extremely fast, as the rapid change in drug-resistant microorganisms shows. We can build software agents that capture basic behavioral notions (i.e., "I should respond to the presence of other agents") but not the details ("Should I be attracted or repelled by agents of a given type? How strongly?"). Then we can evolve them in real time against the observed behavior of the actors they represent, learning the information we don't know and keeping it up to date even if the actors' behaviors change.

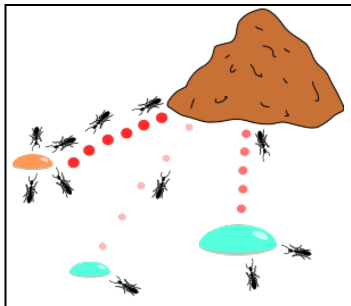

Figure 1: Pheromone-based ant path formation

Self-Organization.—Emergence is a challenge for conventional forecasting mechanisms. Agents in the natural world (including bacteria, social insects, birds, predators, and even humans) exploit it as an important tool. For example, an individual ant might have a hard time surveying a large area for food, but if we allow it to communicate its individual experience to other ants through chemical markers ("pheromones"), and if it responds appropriately to the markers deposited by others, the colony as a whole can efficiently exploit food sources that are far from the next (Figure 1). An individual human farmer might have a hard time setting the right price for his crop, but if we provide him with a marketplace where many buyers and sellers come together, along with a set of rules for bidding in this market, the market as a whole can emergently find a price that matches available supply and overall demand (Figure 2). In both cases, the system harnesses emergence to achieve overall self-organization.

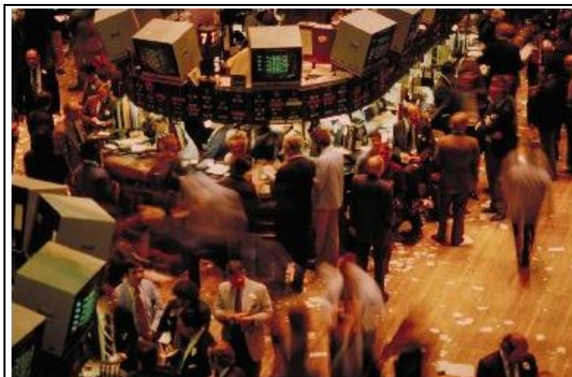

Figure 2: Trading floor of the NYSE

The common feature to these and other examples of self-organization is that agents communicate with one another, not directly, but through a shared environment. In 1959, the French biologist Pierre-Paul Grassé coined the term "stigmergy" the Greek words "stigma," meaning "mark," and "ergon," meaning "work," to describe this form of indirect communication [1]. It is widely observed not only among lower animals, but also in humans [2]. Stigmergy provides the feedback between the individual agent and the overall community that enables self-organization to occur. As a result, instead of being an obstacle to forecasting, emergence actually improves the process, yielding "swarm intelligence" that can guide not only natural populations, but also swarms of computational agents.

The Polyagent.—Traditional applications of Swarm Intelligence deploy a single swarm (collection of autonomous agents) from which the desired functionalities emerge. For instance, in swarm robotics, a single agent controls each robot and complex feats (e.g., distributed persistent surveillance, material logistics) are achieved through the execution of the simple behavioral rules of each such agent. Large-scale self-organization of peer-to-peer network overlays, like the BitTorrent file-sharing system (Figure 3), also emerges from the collective of identical client (agent) processes. But in forecasting applications, where we need to anticipate many possible futures, we can assign a swarm of agents to each actor. Now the actor is represented, not just by an agent, but by a polyagent, and the distribution of individual agents in the swarm gives a distribution over the possible futures that the actor may experience. Techniques invented by Axon's engineers exploit stigmergy to execute multiple futures extremely rapidly. One system running on a Windows desktop in the year 2000 maintained a distribution over 107 possible futures while running an order of magnitude faster than real time.

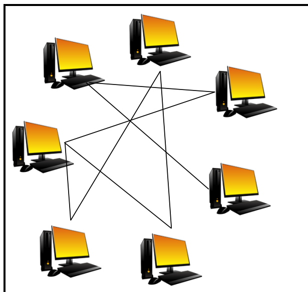

Figure 3: Peer-to-Peer Network (e.g., BitTorrent)

An additional benefit of the use of multiple agents to represent a single actor in a polyagent is that they can be initialized to begin their simulation with the recent past, and then evolved

Mark Slonecker, President
mark.slonecker@axonai.com

Dr. Sven A. Brueckner, Chief Scientist
sven.brueckner@axonai.com

against the actor's observed behavior to maintain a realistic model of the actor's behavioral rules.

AXON:OS.—To deploy these mechanisms rapidly available in real-world settings, AXON AI has developed and is continuously refining AXON:OS, a general-purpose software framework for the rapid development of swarming applications. We have already realized more than 10 different applications and demonstrations in AXON:OS, and each new project contributes to the further expansion, refinement, and hardening of the framework.

With a long history of designing, implementing, and deploying swarming and polyagent models and with the growing sophistication of the proprietary AXON:OS application development framework, the team at AXON AI is in a unique position to pro-vide robust, scalable, and adaptive solutions to complex challenges in many domains.

Examples

The mechanisms embodied in AXON:OS have been developed over three decades of research and applied to a wide range of demonstration projects, in which they have demonstrated clear performance benefits over both people and competing technologies. For example:

- In human-staffed war-games of company-level urban combat, polyagents were able to predict the locations of enemy forces twenty minutes into the future more accurately than game-theoretic or Bayesian reasoners, and also more accurately than human military officers with combat experience [3].
- We trained polyagents on the behavior of terrorists planting IEDs based on past attacks, and then used them to predict future attacks. The resulting predictions were more accurate than those made by humans with experience in IED-prone theaters.
- In a crowd-prediction application, a polyagent model gave police a useful overview of likely crowd movements up to an hour into the future, allowing anticipatory deployment of officers to maintain order [4].

Now AXON:OS allows rapid extension of these promising results to other domains, enabling true anticipatory intelligence for complex national security problems.

References

[1] P.-P. Grassé. La Reconstruction du nid et les Coordinations Inter-Individuelles chez Bellicositermes Natalensis et Cubitermes sp. La théorie de la Stigmergie: Essai d'interprétation du Comportement des Termites Constructeurs. Insectes Sociaux, 6:41-84, 1959.

[2] H. V. D. Parunak. A Survey of Environments and Mechanisms for Human-Human Stigmergy. In D. Weyns, F. Michel, and H. V. D. Parunak, Editors, Proceedings of E4MAS 2005, vol. LNAI 3830, Lecture Notes on AI, pages 163-186. Springer, 2006. http://abcresearch.org/papers/E4MAS05HHS.pdf.

[3] H. V. D. Parunak. Real-Time Agent Characterization and Prediction. International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'07), Industrial Track, pages 1421-1428, ACM, Honolulu, Hawaii, 2007. http://abcresearch.org/papers/AAMAS07Fitting.pdf.

[4] H. V. D. Parunak, H. S. Brooks, S. Brueckner, and R. Gupta. Dynamically Tracking the Real World in an Agent-Based Model. Fourteenth International Workshop on Multi-Agent-Based Simulation (MABS 2013) at AAMAS 2013, IFAAMAS, Minneapolis, MN, 2013. http://www.abcresearch.org/papers/MABS2013CAVE.pdf.

Mark Slonecker, President
mark.slonecker@axonai.com

Dr. Sven A. Brueckner, Chief Scientist
sven.brueckner@axonai.com