

Hybrid Cognition Forum Charter

The Convergence of Human Judgment and Artificial Intelligence

The Business Problem

Organizations do not adopt AI simply to acquire more sophisticated technology. The business case for AI ultimately rests on its ability to create measurable value.

That value is generally expected to come from improving business process performance: doing things faster, better, at lower cost, with higher quality, or with less risk. In many applications, organizations expect improvements across several of these dimensions simultaneously.

At the center of those processes are decisions driving actions. Increasingly, computational systems participate directly in those decisions. As autonomous and agentic systems develop, machines are being given authority—or at least partial authority—to make decisions and initiate actions that previously required human involvement. Typically, the measurable goal is to move the needle on process performance indicators or key performance indicators in a favorable direction.

The Fundamental Challenge for AI Adoption

The technologies organizations are deploying to improve process performance increasingly rely on probabilistic systems commonly referred to as AI. Their results and conclusions are therefore not certainties. Models can be wrong. Confidence can be misplaced. Data and relationships can drift. Causal explanations can be spurious. A system can produce a plausible answer that is incorrect. An autonomous agent can make a locally reasonable decision that produces an undesirable organizational outcome.

AI can therefore simultaneously create business value and decisional risk. Simply keeping a human in the loop does not eliminate that risk. Humans have their own characteristic strengths and failure modes. They bring experience, contextual understanding, judgment, creativity, and an ability to recognize circumstances that may not fit an established model. They are also subject to cognitive bias, limited attention, information overload, selective interpretation, premature closure, organizational pressures, and other constraints.

Computational systems have different strengths. They can process enormous quantities of evidence, identify patterns, maintain consistency, continuously monitor changing conditions, and search alternatives at a scale humans cannot. They also introduce

different vulnerabilities, including model error, probabilistic uncertainty, data and model drift, incorrect causal inference, and failure to recognize when circumstances have moved beyond their operating assumptions.

The question, therefore, is not simply whether a decision should be made by a human or a machine. It is: how should—and how can—human judgment and computational intelligence be best combined across business decision processes to improve process performance, decision quality, organizational performance, and measurable value—and potentially create competitive advantage—while managing the decisional risks inherent in both human and probabilistic machine reasoning? The Hybrid Cognition Forum is a multidisciplinary, practitioner-oriented professional community organized around that question.

Business Decision-Making Exists on a Continuum

We begin with several working premises.

First, business decision-making exists on a continuum. A decision process may be predominantly human, predominantly machine, or some combination of the two. We do not assume that keeping a human involved is inherently better, nor that greater automation is inherently better. Our interest is determining the configuration of human and computational cognition best suited to a particular business environment. That means asking which cognitive and decision functions should be performed by people, which should be performed computationally, which should be shared, and under what circumstances that allocation should change.

The appropriate configuration may evolve. A decision initially requiring substantial human involvement may become suitable for greater machine autonomy as evidence, experience, and confidence accumulate. Conversely, changes in the operating environment may invalidate previous assumptions and require decision authority to be constrained or returned to a human.

The goal is to find configurations that improve process performance and create business value while appropriately managing decisional risk.

Humans and Machines Fail Differently

Both human and machine decision-making have characteristic strengths and failure modes. This is fundamental to Hybrid Cognition. We seek configurations that combine the complementary strengths of human and computational cognition while reducing the failure modes of each when operating independently. This requires understanding when machines outperform humans, when humans outperform machines, and when interaction between the two produces better results than either could achieve alone. It also requires recognizing circumstances in which adding a human to a process creates

unnecessary delay, cost, inconsistency, or error—and circumstances in which removing human judgment creates unacceptable decisional risk.

Governance Is Part of the Decision Architecture

Governance is integral to Hybrid Cognition. As cognition and decision authority move between humans and machines, organizations must determine who or what is authorized to decide and act, over what decisions, within what boundaries, based on what evidence, and with what accountability.

Governance must also establish how performance and uncertainty will be monitored, what controls apply, and what conditions require intervention, escalation, restriction of machine authority, or return of a decision to a human.

Governance therefore should not be treated merely as compliance imposed after a system has been designed. It is part of the decision architecture itself.

KPIs and OODA Are Starting Frameworks

We need useful frameworks both for defining the result to be improved and for examining how decisions produce that result. To get started, the Forum will use widely accepted key performance indicators (KPIs) as outcome anchors and John Boyd's OODA decision-making framework—Observe, Orient, Decide, Act—as a lens for examining the decision process, with particular attention to Orientation.

Well-established KPIs give participants from different industries and disciplines a common, practical language. They are already understood by operating managers, have measurable definitions, and connect directly to business performance. Organizing a case around a KPI keeps attention on the outcome to be improved while allowing competing human-machine configurations, prototypes, methods, and governance approaches to be compared against the same result.

A KPI provides an outcome measure rather than a complete explanation. Forum work will therefore trace the decisions and actions upstream of the KPI: what information was available, how the situation was interpreted, which alternatives were considered, who or what had authority, what action followed, and how the result fed back into the next decision cycle. This makes the KPI the organizing target while Hybrid Cognition remains the subject of inquiry.

On-Time Delivery as a Working Case

On-time delivery (OTD) is a useful initial case because it is widely recognized and easily measured, yet the outcome is produced by a chain of interdependent decisions. Order promising, material availability, capacity, scheduling, production priorities, quality, logistics, customer communication, and exception management can all affect whether

an order arrives on the requested or committed date. The case is understandable across disciplines while still being rich enough to expose the strengths and failure modes of both human and computational reasoning.

For example, a prototype might estimate the probability that a requested delivery date is attainable after considering the relevant operating conditions. An order-entry person might be authorized to accept the requested date when the estimated probability is at least 80 percent. Below that threshold, the customer could be given a probability range, offered alternatives, or the request could be escalated. The Forum can then examine the evidence behind the estimate, whether 80 percent is the right threshold, how uncertainty should be communicated, where human judgment adds value, what authority the system should have, and what conditions should trigger intervention.

OTD also illustrates why a KPI should be considered within a broader performance system. A design that improves delivery performance by excessive expediting, inventory, overtime, or conservative promise dates may simply move cost or risk elsewhere. Related measures such as margin, lead time, expedite cost, inventory, schedule stability, customer satisfaction, and forecast calibration help reveal whether the Hybrid Cognition design creates genuine value rather than a local improvement.

For decision making Boyd emphasized the advantage created by faster and more accurate orientation, enabling more effective decision and action. Modern computational systems offer significant opportunities to improve the speed, breadth, and quality of orientation by integrating diverse evidence, detecting weak signals and patterns, examining competing explanations, assessing causality, challenging assumptions, projecting possible outcomes, and identifying alternative courses of action.

But speed alone is not the objective. Faster orientation has little value if the resulting understanding of the situation is wrong. We are interested in both the speed and quality of orientation, and in how human and computational cognition can be combined to improve both.

OODA is our starting point, not doctrine. We expect the framework to be challenged, extended, and perhaps substantially modified as the Forum develops, different disciplines contribute their perspectives, and evidence accumulates.

Hybrid Cognition Is Cross-Disciplinary

Hybrid Cognition is inherently cross-functional. Effective business decision systems require perspectives that rarely reside within a single profession or discipline. The Forum therefore brings together practitioners, managers, executives, technologists, researchers, academics, product developers, and other professionals. Insights from management, decision science, behavioral science, statistics, operations, finance, engineering, systems thinking, healthcare, project and program management, governance, and other fields can all contribute. A concept developed in one domain may prove unexpectedly valuable in another. The Forum is deliberately not IT-centric.

Technology is an important component of Hybrid Cognition, but so are people, processes, decisions, economics, governance, organizational behavior, and implementation.

From Ideas to Deployable Solutions

The Forum has a strong practitioner orientation. We want to balance technological capability and academic or theoretical insight with the realities of practical implementation at scale. Our progression is:

Ideas → Research → Experimentation → Implementation → Deployable Solutions → Measurable Value

Technology helps establish what is becoming possible. Research and theory help explain why an approach should work, identify limitations, and provide propositions that can be tested. Practitioners bring knowledge of the environments in which solutions must actually operate.

A promising concept still has to confront real data, existing workflows, organizational behavior, incentives, human and machine roles, governance, controls, accountability, escalation, adoption, cost, and integration with existing business processes.

Members are encouraged to share ideas, research findings, tools, methods, frameworks, useful resources, business cases, implementation experience, experiments, successes, failures, and unresolved problems. Where promising concepts emerge, we want to examine what is required to move them beyond discussion or demonstration toward deployable solutions that practitioners can realistically use in operational business environments.

Implementation itself is not the final test. The question is whether the solution produces better process performance: Faster. Better. Lower Cost. Higher Quality. Lower Risk. And whether those improvements ultimately translate into measurable business value.

Economic Value, ROI, and the Real-World Test

Improved process performance does not automatically make an AI or Hybrid Cognition solution economically worthwhile. A proof of concept may demonstrate faster processing, lower operating cost, better predictions, improved quality, or reduced risk and still fail the business test if the cost of implementation and sustained operation outweighs the benefits.

The Forum will therefore consider ROI, cost-benefit analysis, total cost of ownership, scalability, time to value, and ongoing operating and maintenance costs as important dimensions of real-world success. This includes costs that can be overlooked during experimentation: data acquisition and preparation, integration, infrastructure, model and

system maintenance, monitoring, governance, human oversight, training, change management, and adaptation as business conditions change.

The relevant question is not simply “Does it work?” It is also “Does it create sufficient value to justify deploying, operating, governing, and sustaining it?”

Developing the Hybrid Cognition Workforce

Hybrid Cognition also has important implications for education and workforce development. As machines assume greater roles in analysis, orientation, decision support, and autonomous action, the knowledge and skills required of people will change. This applies not only to technical specialists, but also to practitioners, managers, executives, and others who increasingly work with, rely upon, supervise, or govern computational systems.

The Forum will explore what people need to know and be able to do to work effectively in these environments. This may include understanding probabilistic outputs and uncertainty; evaluating evidence; recognizing human and machine failure modes; challenging computational reasoning; exercising judgment when evidence is incomplete or contradictory; understanding the appropriate boundaries of machine authority; governing autonomous activity; and recognizing when intervention or escalation is required.

We are interested not only in what machines can do, but what people need to learn to work effectively alongside them.

Learning From What Actually Happened

Real-world experience is an essential source of knowledge. The Forum will particularly encourage case studies, post-mortems, and lessons learned. Accounts of what worked, what failed, what produced unexpected consequences, what changed during implementation, and why can provide some of the Forum's most valuable learning.

Successful implementations matter, but failures, abandoned projects, disappointing proofs of concept, unexpected costs, governance problems, and solutions that worked technically but failed to create sufficient business value may be equally instructive.

The objective is not simply to showcase success. It is to improve our collective understanding of what works, what does not, under what conditions, and why.

How Forum Working Sessions Can Be Structured

A working session may be organized around a KPI, a business decision problem, or an emerging participant idea. The structure is intended to turn diverse contributions into cumulative learning while leaving room for research presentations, open discussion,

and timely topics. A typical case-oriented session can include the following connected elements.

Frame the outcome and decision process. Define the KPI or business result, the process boundary, the decision points that influence it, and the current allocation of responsibility between people and computational systems.

Bring participant experience into the case. Members share operating experience, data, lessons learned, research, analogous cases, constraints, and ideas. Successes and failures are both useful because they reveal conditions a prototype or theory must survive.

Demonstrate something that can be examined. A prototype may be a decision rule, spreadsheet, simulation, workflow mock-up, dashboard, analytic model, agent, or working product. Early prototypes are welcome; their purpose is to make assumptions and decision logic visible enough to test and improve.

Examine the Hybrid Cognition design. Participants consider the quality of orientation, the division of cognitive work, decision authority, uncertainty, human and machine failure modes, governance, escalation, implementation requirements, economics, and possible unintended consequences.

Define the next test. The session concludes by identifying what was learned, which assumptions remain uncertain, what evidence is needed, and what prototype modification, experiment, field test, or participant contribution should come next.

A case can begin before a finished prototype exists. A participant's operational problem, experience, research finding, or proposed decision method can begin the cycle. Other members may contribute data, challenge assumptions, suggest designs, develop prototype components, or test the approach in another setting. Over multiple sessions, the Forum can build a documented body of cases showing which human-machine configurations work, under what conditions, and with what measurable results.

Prototypes, Product Presentations, and Demonstrations

The Forum welcomes prototypes, product presentations, and demonstrations when relevant to the questions we are exploring. Forum participants, developers, vendors, researchers, and practitioners are encouraged to bring work at different stages of maturity—from an early decision rule or workflow mock-up to an operating product—that illustrates how human and computational cognition can be combined in actual business decision processes.

These sessions should go beyond conventional product promotion. We want to understand the business decision problem being addressed, the process improvement being sought, how cognition and decision authority are allocated between humans and machines, why that configuration was chosen, how uncertainty and failure modes are

handled, how the system is governed, what implementation requires, and what evidence exists regarding performance, cost, ROI, and business value.

Products should be open to the same constructive examination as theories and research. What works? What does not? Where does human judgment add value? Where might it add unnecessary friction? What can appropriately be delegated? What assumptions does the system depend upon? What happens when those assumptions no longer hold? How is decisional risk monitored? When should the system escalate to a human?

A working product or prototype provides an especially valuable bridge between theory and practice because it allows us to determine whether an attractive concept survives contact with real data, workflows, users, governance requirements, organizational constraints, and economic realities.

A Community for Exchange and Collaboration

The Forum should create practical value for the people who participate in it.

Members should encounter ideas they can use, research they might otherwise not encounter, tools and resources they can evaluate, and perspectives from people outside their normal professional disciplines.

The Forum provides a collegial environment where members can bring problems and emerging ideas to one another, compare approaches, challenge assumptions constructively, and learn from successes and failures.

We also want members to establish professional relationships with others who share similar intellectual and practical interests. Those relationships may lead to continuing exchanges of ideas, collaborative experimentation, research, product development, implementation efforts, or simply access to knowledgeable colleagues when a new problem arises.

The Forum is intended to become a community of practice, not simply an audience for presentations or another source of technology news. Contribute what you know. Learn what others know. Challenge ideas together. Build relationships. Put what works into practice.

What We Will Explore

Our mission requires us to push boundaries while keeping the discussion practical. Topics may include machine-assisted perception and orientation; human and machine decision failure modes; probabilistic decision systems; decisional risk; causal analysis and competing hypotheses; falsification and contradiction testing; prediction and decision support; automated managerial decision-making; project and program

management; operations; finance and risk; supply chains; agentic workflows; allocation of human and machine decision authority; bounded autonomy; escalation from machine to human authority; governance of human-machine and autonomous decision systems; education and workforce development; ROI and economic evaluation; and design and implementation of deployable Hybrid Cognition solutions.

Data and evidence are also fundamental to decision quality. The Forum will explore data quality, provenance, source reliability, relevance, data generation and preparation, methods for evaluating evidence, and the effects of changing data and relationships on both human and computational decision-making.

We are particularly interested in business processes where decisions once made predominantly by humans are increasingly shared with or delegated to machines and autonomous agents. These environments raise important questions: What should be automated? What should remain human? What should be shared? What constraints should govern machine authority? How should decisional risk be measured and managed? What evidence should trigger human intervention? What skills do people need as these roles change? And how do we know whether the resulting system actually creates value?

The Forum is not limited to AI technology itself. Our subject is the business decision system: how information is sensed and interpreted, how situations are understood, how competing explanations are considered, how alternatives are developed, where judgment and decision authority should reside, how decisions become actions, how those decisions are governed, and how results feed back into subsequent orientation.

Staying Focused on the Business Proposition

The business case begins with process performance:

AI Adoption → Faster · Better · Lower Cost · Higher Quality · Lower Risk → Measurable Value

Better decision-making is central to that improvement, while probabilistic AI adds decisional risk. Hybrid Cognition addresses that intersection:

Human + Machine Strengths → Better Orientation → Better-Governed Decisions and Actions → Better Process Performance → Measurable Value → Potential Competitive Advantage

And measurable value must ultimately survive the economic test:

Performance Improvement → Lifecycle Benefits, Costs, and Risk → ROI and Economic Value

The goal is the right, governed configuration of human and computational cognition. Evidence must show that it reduces decisional risk, improves process performance, builds the required human capability, and creates sufficient economic value.

The Hybrid Cognition Forum meets for one hour every other Wednesday at 12:00 PM Central via Zoom. Current information is available at www.ratio-weekly.com.