WHITE PAPER

COGILITY

HCEP: Expert AI for High-Consequence Decision Making

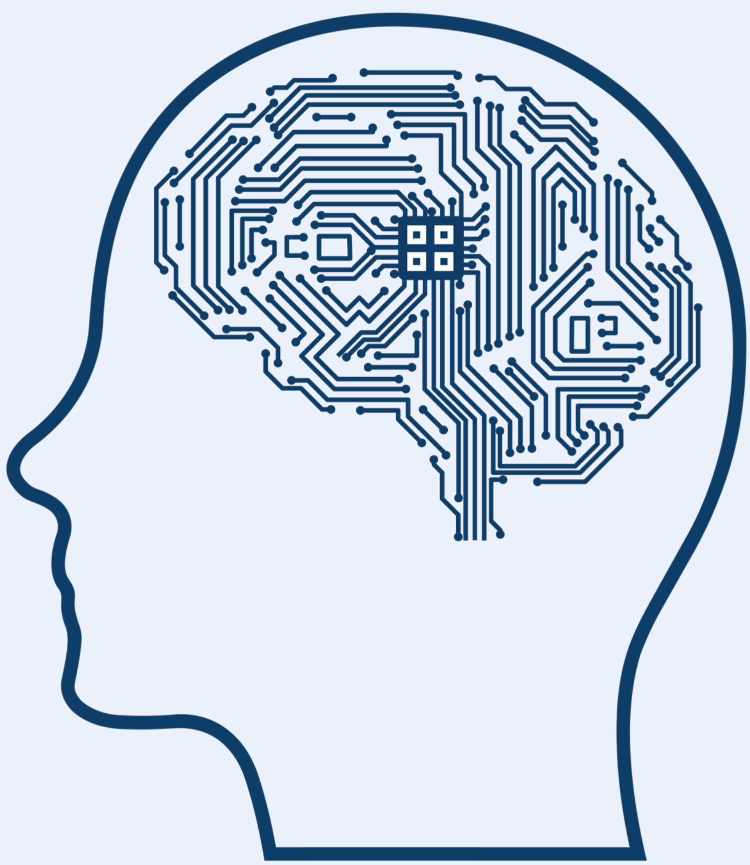

Executive Summary

Complex Event Processing has evolved beyond simple data tracking into a sophisticated form of Expert AI. By encoding human expertise into hierarchical models, the Cogynt.ai Continuous Decision Intelligence Platform allows organizations to establish situation awareness and identify meaningful patterns—such as insider threats or cyber risks—in real time. ***Cogynt.ai's Hierarchical Complex Event Processing (HCEP)*** *model implements an expert system that mimics human reasoning to distill massive data streams into actionable intelligence, delivering trusted decision-making in high-stakes environments.*

Organizations today face increasingly complex, high consequence decision environments—ranging from national security to financial fraud to healthcare risk—where the cost of a wrong decision can be catastrophic. In these domains, AI systems must do more than analyze data; they must support **transparent, explainable, and trusted** decision making that mirrors human expert reasoning. Cogility's **HCEP** model, implemented in the **Cogynt.ai Continuous Decision Intelligence Platform**, delivers exactly this capability: an Expert AI system that can detect meaningful patterns in massive data streams, sustain contextual awareness over time, and proactively assess evolving risks.

HCEP integrates the foundational principles of historical AI approaches—expert systems, hierarchical pattern recognition, and complex event processing—with modern, scalable technologies such as Apache Kafka, Apache Flink, Kubernetes, and real time streaming analytics. By encoding human expertise into hierarchical, multi layer event models, **HCEP** transforms raw, noisy data into high level "complex events" that reflect the knowledge structures and inference processes used by trained analysts. This results in AI driven insights that are both operationally powerful and inherently explainable.

A core strength of HCEP is its **structured sensemaking**, achieved through abstraction hierarchies that decompose expert hypotheses into observable evidence patterns. Combined with **stateful processing**, which preserves contextual memory indefinitely, the system can reason about partially matched patterns, evolving situations, and inferred causality—capabilities essential for real time situation awareness. **HCEP's Continuous Risk Assessment** further enhances decision support by applying probabilistic reasoning that updates risk dynamically as new information arrives.

Executive Summary (Continued)

Cogynt.ai also incorporates practical tools that accelerate analytic insight, including customizable **lexicons** for extracting meaning from unstructured text, manual actions that allow expert intervention and oversight, and the **EPCL no code authoring environment** that democratizes model creation for non technical users.

In contrast to opaque machine learning "black box'' systems, **HCEP** offers **explainability by design**. Every assessment is grounded in a traceable causal chain of events, enabling analysts and decision makers to understand why a risk was flagged and to trust the system's conclusions. This alignment with human cognitive processes—hypothesis → evidence → conclusion—makes **HCEP** a uniquely powerful platform for mission critical intelligence.

HCEP's operational value has been demonstrated in real world, high stakes missions:

- During military operations in Iraq and Afghanistan, **HCEP** based analytics enabled rapid identification of Improvised Explosive Devices (IEDs), reducing analytic cycles from days to hours while scaling expert intuition across the force.
- **Cogynt.ai's HCEP** expert AI modeling delivers a proven, scalable Insider Risk Management (IRM) solution for two major U.S. government organizations by replicating expert IRM reasoning to transform raw data into transparent, proactive, whole person analytics and explainable, actionable intelligence.

Myriad additional potential applications of Cogynt.ai's HCEP expert AI behavioral analytic modeling are envisioned that span high-risk/high complexity domains as maritime security, financial crime, and healthcare monitoring—virtually any domain requiring transparent, defensible inference from complex, multi modal data.

In sum, **HCEP** marks a paradigm shift from passive monitoring to active Expert AI: a system that continuously senses, interprets, and forecasts complex situations with transparency, precision, and expert level reasoning. In high consequence environments, **HCEP** elevates organizations from reactively monitoring overwhelming data to truly understanding it—delivering expert, explainable, and proactive intelligence to gain a decisive operational advantage.

HCEP Origins: A Synthesis of Historical AI Approaches

Cogynt.ai's **Hierarchical Complex Event Processing (HCEP)** model builds on early foundations of artificial intelligence (AI) applications in decision support. As described in a recent whitepaper, Tracing the Evolution of Artificial Intelligence in Decision Support Systems (Greitzer, 2025), there are traces of Selfridge's (1959) Pandemonium model and of knowledge-based expert systems (e.g., Feigenbaum, 1977) in HCEP's hierarchical pattern processing, in addition to the direct influence of Luckham's (2002) Complex Event Processing (CEP) model enabling real-time detection and response to event patterns through hierarchical rule sets. Cogility's expert AI solution, as implemented in its **Cogynt.ai** platform, serves as an archetype for integrating historical AI concepts and advanced decision intelligence tools into practical, modern solutions.

Expert systems use codified knowledge and a rule-based inference engine for transparent, explainable decision-making. Machine Learning (ML) systems—particularly deep neural networks and support vector machines—have introduced powerful but less explainable forms of decision support. ML systems excel in handling high-dimensional data and uncovering complex relationships but often lack transparent mechanisms for explaining their decisions. Classification-based decision systems include tree-based ML models that have varying degrees of transparency. For example, decision trees use hierarchical structures resembling flow charts that visualize potential outcomes, with branches that are straightforward to interpret representing individual decisions. Random Forests aggregate hundreds or thousands of decision trees by independently averaging their predictions. This yields higher predictive accuracy while reducing transparency (explainability).

Cogynt.ai's HCEP model builds on early foundations of AI in decision support

David Luckham's CEP model extended the principles of expert systems by providing architectures for detecting and responding to patterns of events in real time. CEP systems utilize hierarchical rule sets to process streams of data, identify significant event patterns, and trigger actions based on predefined logic. This approach reflects the influence of both expert-based decision-making to emphasize interpretability and control over the reasoning process and hierarchical pattern processing to enhance real-time processing. This unique implementation of pattern-based, event-driven decision support produces a hierarchy of events, consistent with the experts' knowledge, that provides more understandable and explainable outputs.

The incorporation of hierarchical knowledge structures in AI expert system based models supporting information analysis/risk assessment is consistent with research in cognitive psychology (e.g., Bower, 1970; Rasmussen, 1985; Purcell & Kiani, 2016; Eckstein & Collins, 2020) indicating that human knowledge and decision-making are structured hierarchically, allowing the brain to manage complex tasks by organizing information in nested, abstract levels—facilitating memory retrieval and enabling efficient multi-level decision making. Building on this research, Greitzer and colleagues proposed a conceptual model for insider risk assessment based on hierarchical, pattern-based inference processes (see "Pattern-Based Inference for Threat Assessment" text box, below)—a conceptual framework that Cogynt.ai's continuous, streaming **HCEP** implementation operationalizes. With this foundation established, the next section describes how **HCEP** is implemented within the **Cogynt.ai** platform.

Pattern-Based Inference for Threat Assessment

Greitzer & Frincke (2010) proposed a conceptual model for predictive insider threat classification, using a "shredded puzzle" metaphor to emphasize the challenge of inferring behavioral patterns (puzzle pieces) from large, fragmented data streams. Their framework relies on a knowledge base of indicators and heuristic models describing insider behavior, where indicators represent interpretations of intentions and actions derived from observed data. This work led to the development of the comprehensive SOFIT (Sociotechnical and Organizational Factors for Insider Threat) insider risk indicator knowledge base (Greitzer et al., 2018). The model defines a multilayer inference process at increasing levels of abstraction, from **Data → Observations → Indicators → Behaviors**, in which raw data are transformed into observations, observations into indicators, and indicators into behavioral assessments. Some risk indicators are directly observable, while others must be inferred. The approach stresses the need to integrate both human (psychosocial) and cyber indicators, recognizing that combining these domains is essential yet methodologically challenging for predictive insider threat analysis.

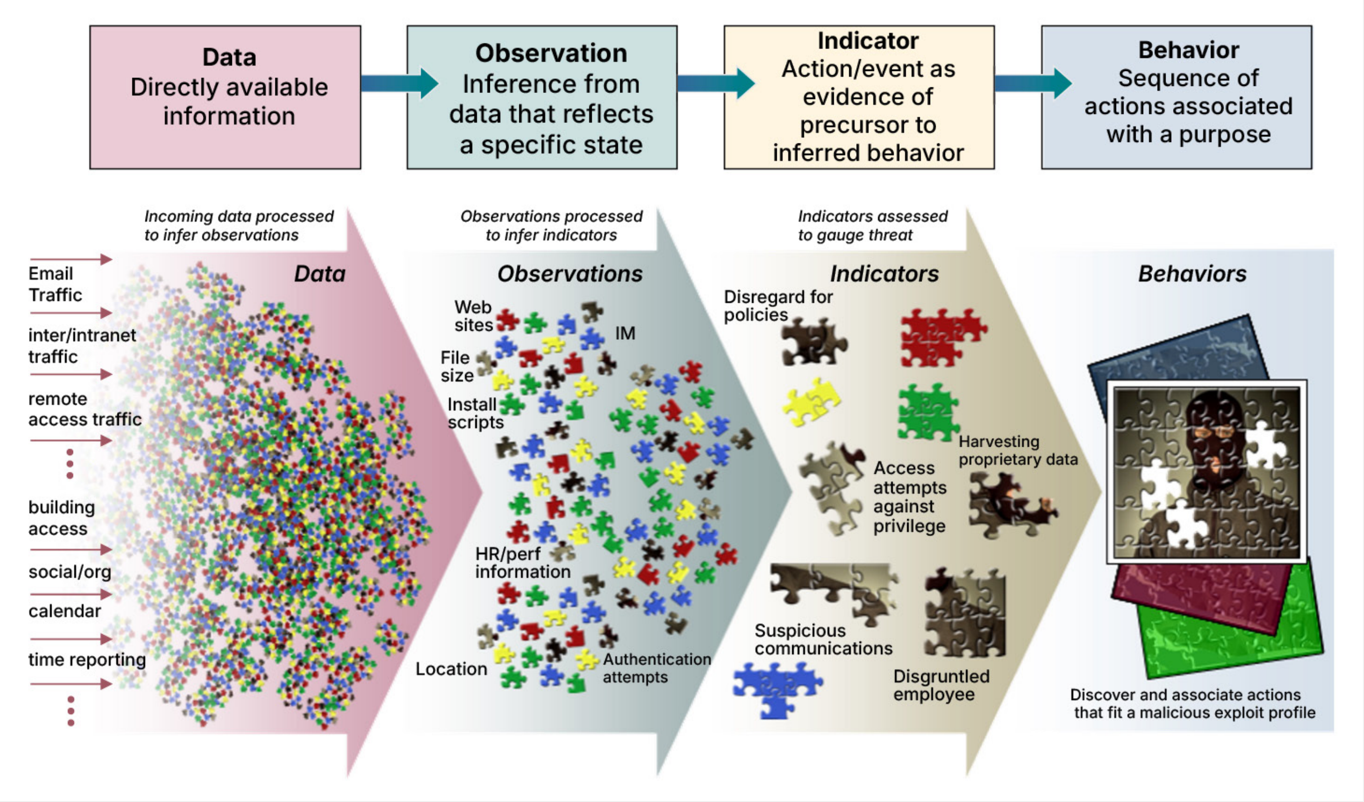

Shredded Puzzle Metaphor Illustrating Risk Analysis Process

HCEP Foundational Concepts

HCEP was developed by applying practical problem-solving methods derived from computer science, software engineering, systems thinking, and AI/expert systems approaches aimed at discovering patterns in data. The root of this approach is the ability to apply the powers of abstraction and decomposition, which provide the basis for HCEP's multi-layer event processing, concept matching, and creation of new "abstract events." Cogynt.ai's HCEP capabilities, described next, are readily employed by non-technical users to detect patterns at scale.

Example of Complex Event Processing: A Wedding

Suppose someone wishes to know about all the weddings taking place at any given time in Los Angeles on any given day. For this example, we define a western wedding requiring observation of four events: “bells ringing,” “man in a tuxedo,” “woman in wedding dress,” and “gifts,” and by various means these events must be observed inroughly the same location and within one hour. If all these events are observed and consistent with the defined constraints, this would comprise a fully matched “pattern” in HCEP, allowing us to make the assertion that a wedding is occurring.

The HCEP analysis is illustrated at right with elements indicating events, event patterns (combinations of events), and constraints (such as consistent timing and geographic location shown as lines with arrows). If we were to observe just two of these events - “man in tuxedo” and “woman in white dress” - this partially matched pattern may still add to our assessment that this could be wedding, but if the observations were made in separate locations and/or at different times (e.g., more than a mile apart, more than an hour apart), these “constraint violations” would lead us to reject this inference.

The more completely the events match the pattern and meet the constraints, the greater the confidence (probability output) in the HCEP assertion. Statistical inferences and other more sophisticated considerations (such as requirements about the order of events) may be applied to enhance the accuracy of the decision.

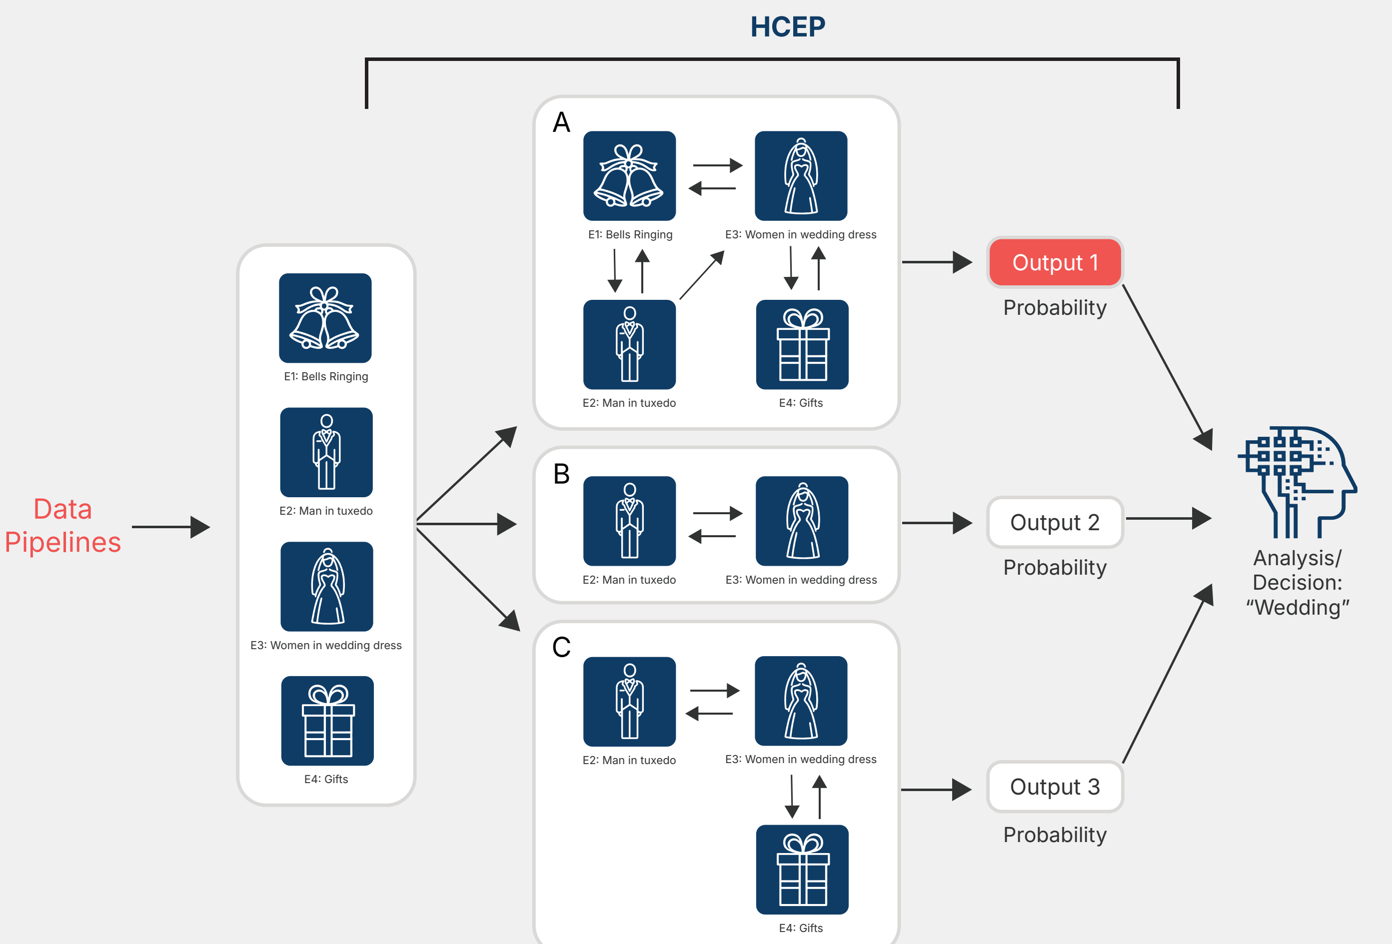

Computation Hierarchy (Structured Sensemaking)

HCEP uses abstraction hierarchies to mirror how experts decompose problems. Abstraction hierarchies are event patterns organized in a hierarchical form. The hierarchy is typically conceived and developed from the top-down (see Figure 1) starting with a hypothesis to discover specific patterns in data that will answer the hypothesis (e.g., the hypothesis, "Is an insider threat developing?"). The HCEP structure is decomposed to a level that can match observed event data, which is referred to as the leaf or lowest level of the hierarchy. Bottom-up processing of data is where the pattern-matching process takes place. The system processes millions of low-level events and matches them against these structures. Matched patterns create new "complex events" that trigger event patterns at the next level, and so on. The employment of computation hierarchical concepts helps avoid combinatorial explosion, which often occurs in rule-based systems that attempt to address every permutation of a rule. In **HCEP**, we define multiple variants of a pattern at one level that each produce the same complex event (and therefore support the same inference). This abstraction hierarchy structure avoids information overload by abstracting raw data into higher-level "complex events," allowing decision-makers to focus on the narrative rather than the noise. The ability to manage complexity is key to HCEP and its ability to scale and handle extraordinarily complex analytic problems.

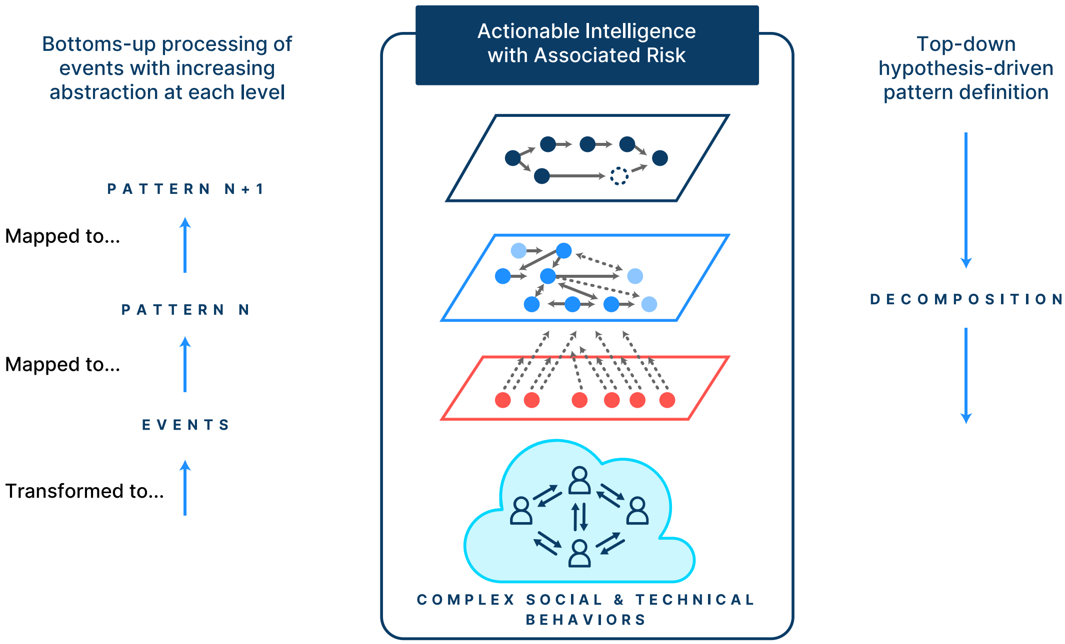

Figure 1. HCEP Computation Hierarchy

Stateful CEP (Contextual Memory)

True awareness requires context. Experts do not analyze facts in isolation; they remember what happened yesterday. **HCEP** is **stateful**, meaning it retains knowledge of matched or partially matched patterns indefinitely. Thus, the system tracks the "state" of a behavior over time (e.g., remembering a badge-in event from 48 hours ago to contextually analyze a file download today). This allows the system to mirror human reasoning and intuition, such as recognizing patterns even when they are incomplete and asserting a *potential* match even though only partial evidence is observed. Stateful processing represents our knowledge about patterns we care about, enabling us to infer potential outcomes (thus supporting more proactive analyses). This also supports *inferred* causality (Luckham, 2002) between event relationships, which is crucial when having to explain how a particular assessment was determined.

Continuous Risk Assessment

In addition to being stateful, **HCEP** supports quantitative risk assessment to determine the likelihood of events occurring in the future. In high-consequence scenarios, risk is not static; it evolves. **HCEP** employs **probabilistic models** to perform **Continuous Risk Assessment**. Unlike rigid rule-based systems, HCEP uses probabilistic reasoning, assigning conditional probabilities to events. These probabilities are subjective but are established through careful expert knowledge elicitation studies (e.g., Greitzer, in press) and reflect what the analyst or subject matter expert believes is important. As new data arrives, the system mathematically updates the likelihood of a threat or opportunity. This transforms the system from a passive monitor into a proactive forecasting tool, giving analysts insights into developing situations so they can mitigate risks before they materialize. This continuous risk assessment gives us insights about developing situations based on risk or opportunity, which allows us to proactively mitigate risks or better leverage opportunities.

Lexicons

A lexicon is vocabulary of a language or branch of knowledge. For **Cogynt.ai** and **HCEP**, we commonly analyze unstructured text such as reports, and the narratives in these reports define specific terms that analysts look for in extracting knowledge about what they are analyzing. Cogynt.ai makes this easy by allowing the user to define a lexicon, which can be organized in the form of a tree and decomposed as a hierarchy of terms. These terms, in some cases, serve as the matching relationship between the detected lexicon term and the event type, where we treat the matched word as an event. **Cogynt.ai** also supports word variability such as tense, stemming, and wild card replacements. **Cogynt.ai**'s use of lexicons facilitates rapid extraction of complex narratives that relate to patterns and behaviors that are most relevant to the analyst.

Manual Actions

Because they are derived from human knowledge, **HCEP** models are expert systems. This means that, like any analytic, **HCEP** can get the wrong answer depending on the context. **Cogynt.ai** supports manual actions enabling the human user to interact with the hierarchical processing and modify a matched event pattern. Thus, when needed, an analyst can manually generate an event for a specific instance of a pattern. This manual intervention can reverse a pattern, revise its risk level, or change the nature of the hypothesis matched. This is propagated throughout the entire hierarchy to update the risk scores. If this situation occurs frequently, it suggests that the model requires modification. Human/expert oversite is a critical advantage that is incorporated into the **Cogynt.ai** platform.

Cogynt.ai's HCEP model builds on early foundations of AI in decision support

Event Pattern Constraint Language

The Event Pattern Constraint Language (EPCL) is a domain-specific, declarative, graphical/no-code programming language developed by Cogility. The Cogynt.ai Authoring tool allows users to create models, discover schemas of input data sources (i.e., Kafka topics), and create event pattern hierarchies. Another key aspect of EPCL and the Cogynt.ai Authoring tool is support for computations that allow for the mapping of data from input to output, using mathematical calculations and/or transformations of the incoming data for matched event patterns. The result of this process can be passed on to higher level event patterns, thus supplying context for the sociotechnical phenomenon under study.

Cogynt.ai Authoring tool allows users to create models, discover schemas of input data sources and create event pattern hierarchies

Cogynt.ai Technology Enablers

Key technology enablers have allowed Cogility to develop HCEP in its current form, as shown in Figure 2. HCEP models are developed by analysts using the Cogynt.ai Authoring Tool. Data are streamed or batched to the platform from an external source to Apache Kafka, a stream processing engine with stream storage. Apache Kafka enables an open core architecture. Apache Pinot21 and Apache Superset support real-time streaming Business Intelligence tools.

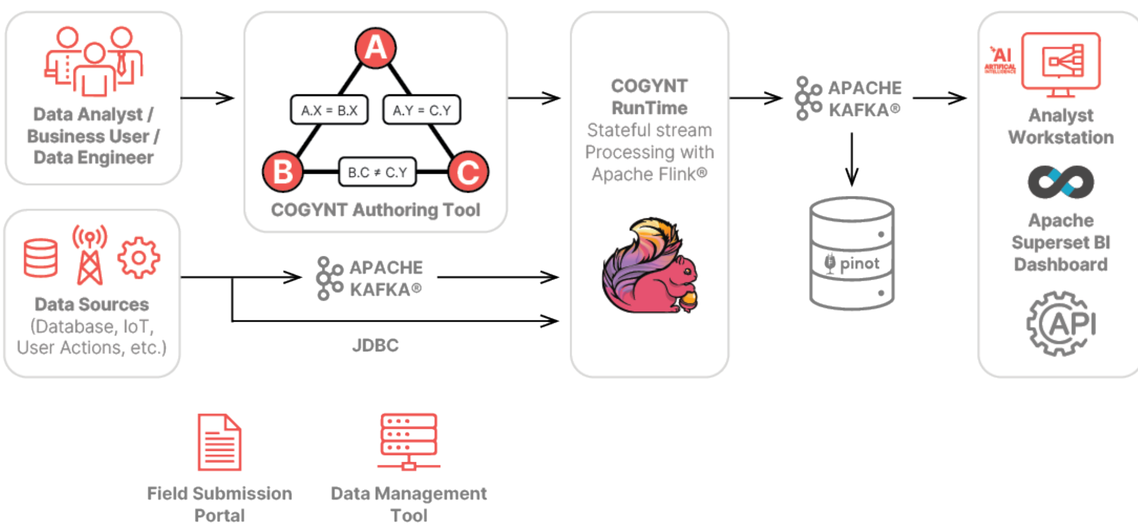

Figure 2. Cogynt.ai Platform High Level Architecture

Expert AI Applications of Cogynt.ai

Countering Improvised Explosive Devices During Operations Enduring Freedom and Iraqi Freedom

The "Expert AI" capability of **HCEP** was proven during Operations Enduring Freedom and Iraqi Freedom to counter Improvised Explosive Devices (IEDs)—an environment where analysts' lives literally depended on the quality of their intelligence.

- **The Problem:** Soldiers faced lethal threats from IEDs constructed from common materials, requiring rapid identification of bomb "signatures" amidst noisy data.
- **The Expert Solution:** Cogility developed an analytic that matched chemical and component signatures (e.g., "Wood Board," "Carbon Rod") against known IED profiles defined by subject matter experts.
- **The Result:** The system, JIST, allowed junior analysts to produce high-relevance intelligence in an hour—a task that previously took PhD-level researchers days—effectively scaling expert intuition across the force.

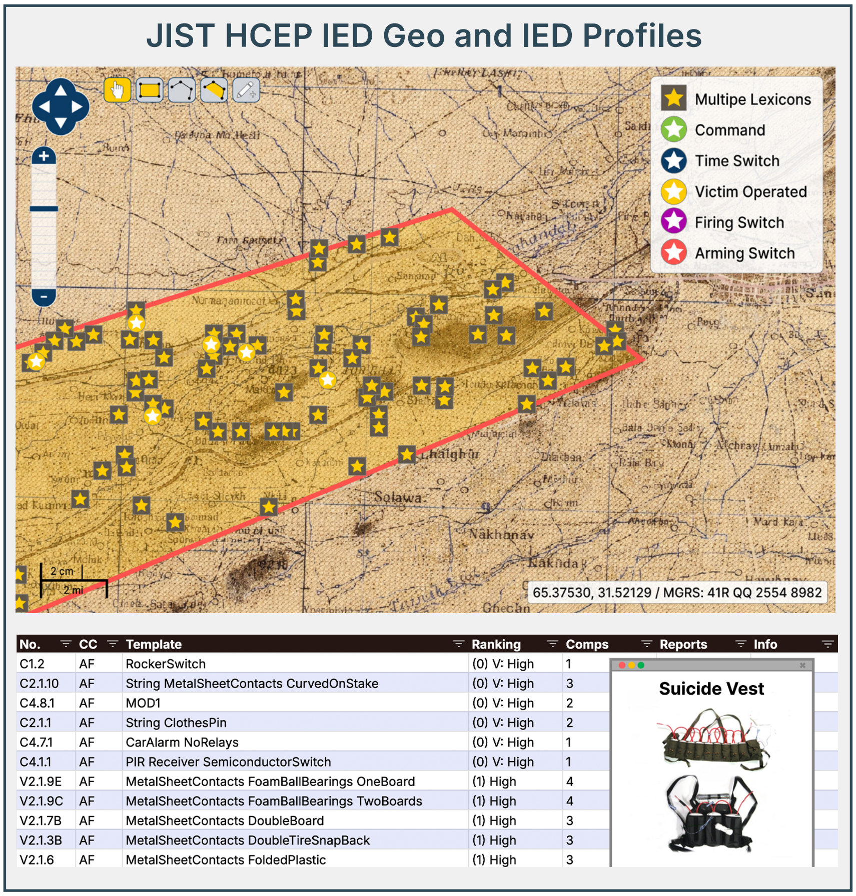

No.	CC	Template	Ranking	Comps	Reports	Info
C1.2	AF	RockerSwitch	(0) V: High	1		
C2.1.10	AF	String MetalSheetContacts CurvedOnStake	(0) V: High	3		
C4.8.1	AF	MOD1	(0) V: High	2		
C2.1.1	AF	String ClothesPin	(0) V: High	2		
C4.7.1	AF	CarAlarm NoRelays	(0) V: High	1		
C4.1.1	AF	PIR Receiver SemiconductorSwitch	(0) V: High	1		
V2.1.9E	AF	MetalSheetContacts FoamBallBearings OneBoard	(1) High	4		
V2.1.9C	AF	MetalSheetContacts FoamBallBearings TwoBoards	(1) High	4		
V2.1.7B	AF	MetalSheetContacts DoubleBoard	(1) High	3		
V2.1.3B	AF	MetalSheetContacts DoubleTireSnapBack	(1) High	3		
V2.1.6	AF	MetalSheetContacts FoldedPlastic	(1) High	3		

Countering Insider Threats

Cogility has developed an expert AI decision intelligence solution for Insider Risk Management (IRM) that applies **HCEP** to deliver a transparent, trusted risk assessment capability. The system integrates the comprehensive SOFIT insider threat indicator knowledge base and models human expertise to support a leading edge, whole person, proactive IRM program.

Cogynt.ai is purpose built to address the demanding information processing, analytic, and workflow challenges faced by insider risk teams across government and industry. Its patented **HCEP** hierarchical pattern processing engine analyzes high volume, heterogeneous technical and behavioral data streams in real time. Structured and unstructured data are processed from the bottom up—transforming raw data into explainable, actionable intelligence identifying relevant potential risk indicators and higher level behavioral assessments.

- **The Problem:** IRM analysts must monitor and interpret massive volumes of real time data, yet traditional tools force them into reactive workflows with high rates of false positives and false negatives.
- **The Expert Solution:** Cogility's behavioral analytic approach uses HCEP to replicate how IRM experts classify and interpret complex behavioral patterns, providing transparent, proactive, whole person analytics.
- **The Result:** The Cogynt.ai decision intelligence platform delivers a proven, scalable IRM solution capable of supporting very large enterprises.

Cogynt.ai IRM Model

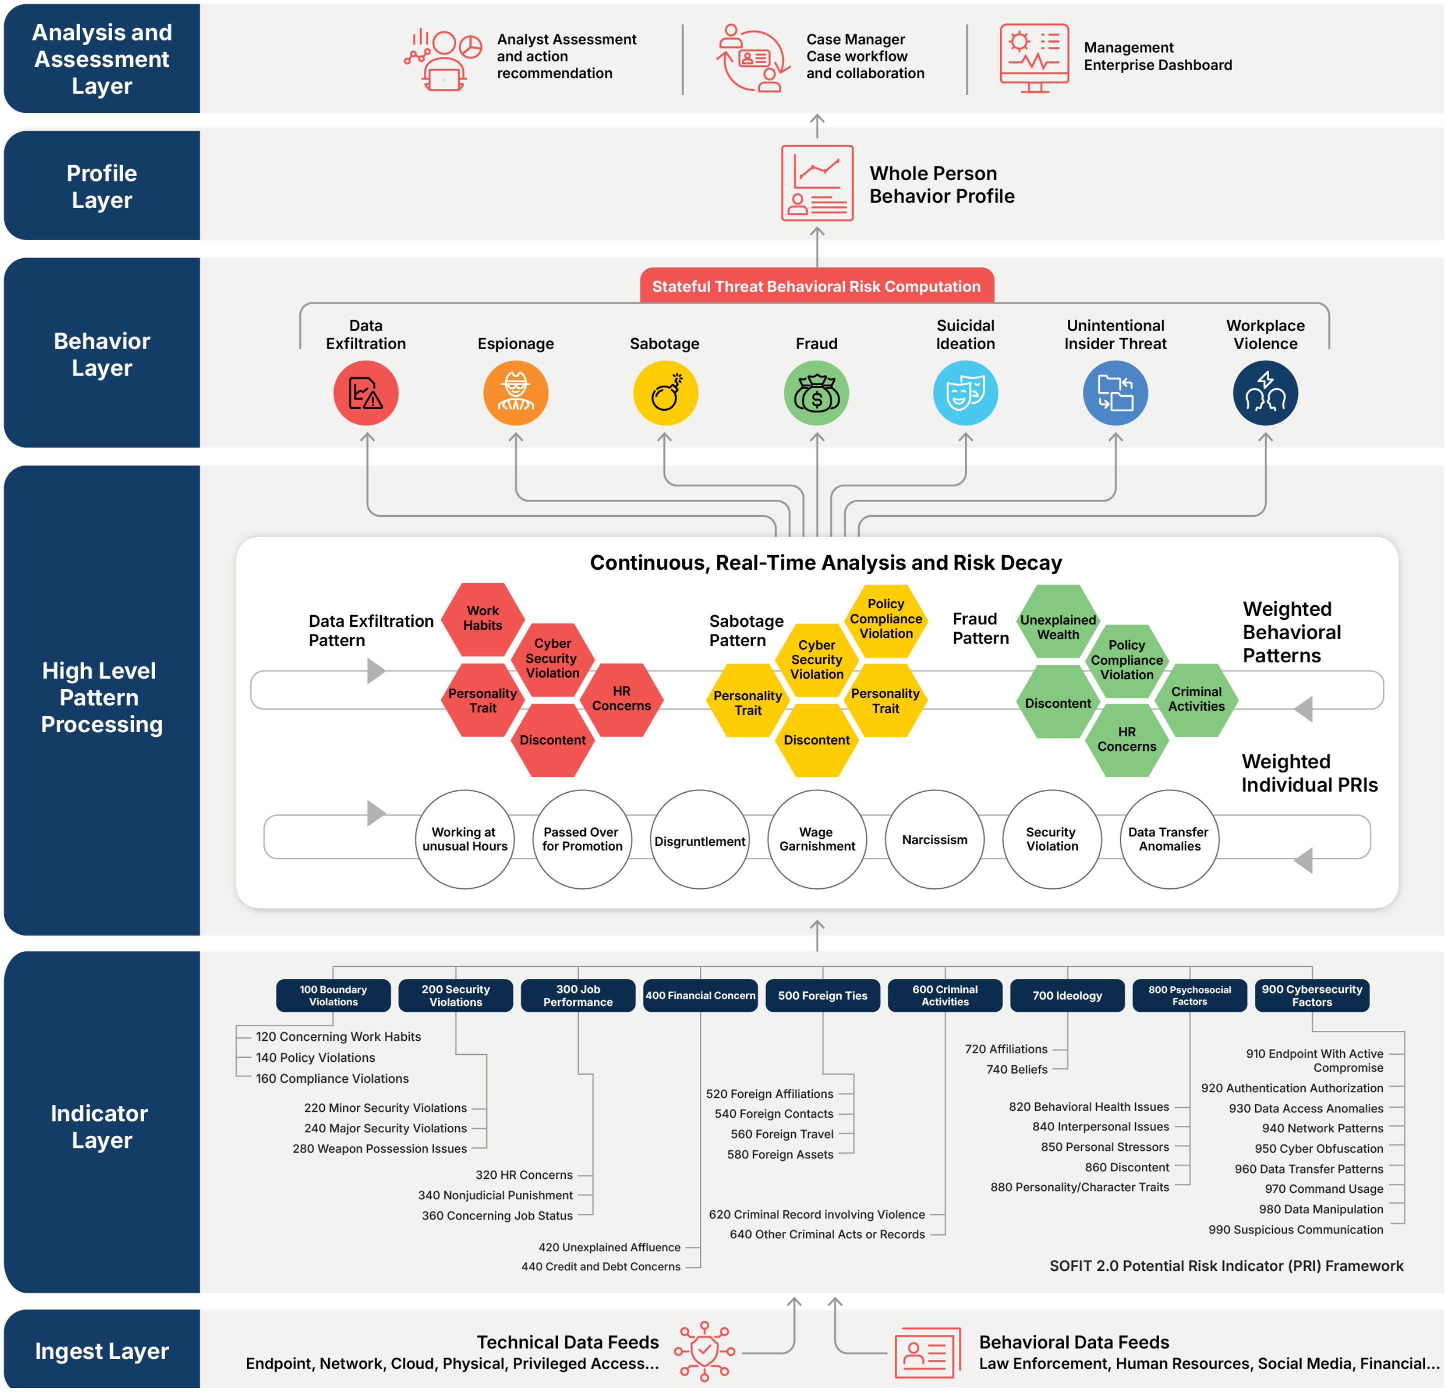

Other Applications

Consider several other applications and potential Cogynt.ai HCEP solutions:

Maritime Smuggling Detection

The application of **HCEP** to maritime smuggling detection follows a clear, expert modeled workflow that identifies suspicious maritime behavior patterns using hierarchical event constructs, continually ingests and matches vessel telemetry data against these patterns, tracks evolving behaviors by retaining long-term context, performs dynamic probabilistic risk updates, and delivers explainable alerts. This allows **HCEP** to detect smuggling risk early, explain its reasoning, and support confident, time critical decisions in a domain where subtle anomalies matter.

Cogynt.ai's expert AI HCEP model encodes analyst tradecraft and domain knowledge in diverse high complexity/ high consequence decision environments

- **Encoding of Expert Knowledge.** Cogility's **HCEP** model can support maritime smuggling detection by encoding expert knowledge about suspicious vessel behavior into hierarchical event pattern models. Analysts define key indicators—such as **AIS tracker deactivation (AIS_Gap)** and **vessel to vessel rendezvous**—which **HCEP** monitors continuously in real time. As streaming data arrives, low level events are matched bottom up to generate higher level "complex events" representing smuggling related behaviors. D:\Cogility Software\Marketing Material\New-AI-HCEP-paper\L49
- **Stateful Retention of Context.** Because **HCEP** is stateful, it retains partial matches and long term contextual information. This allows it to recognize multi step smuggling patterns that unfold over hours. For example, when Vessel X goes "dark," the system raises risk from *Low* to *Elevated*. Hours later, if the vessel reappears and meets another vessel at sea, the additional evidence drives the risk to *Critical*.
- **Continuous Risk Forecasting.** **Cogynt.ai** can update smuggling risk data dynamically as each new piece of evidence arrives, shifting the system from reactive alerting to proactive forecasting.

When risk becomes significant, HCEP produces an explainable alert, including the causal chain—AIS disabled, rendezvous, high risk zone—so operational commanders receive a transparent justification for interdiction decisions. This explainability is essential for trusted action in high consequence maritime operations.

Financial Enterprise Operational and Regulatory Security

Cogility's HCEP model provides a powerful, transparent, and scalable expert-AI framework ideally suited for large financial institutions facing high-volume data, sophisticated fraud patterns, and strict regulatory expectations.

- **Expert-Level Reasoning at Enterprise Scale.** HCEP encodes human expertise into hierarchical event-pattern models that reflect how analysts reason about fraud, insider risk, or anomalous financial behavior. These abstraction hierarchies decompose complex hypotheses into observable evidence, allowing the system to transform massive data streams into explainable, high-level intelligence. This supports financial-sector needs for interpretable, defensible risk assessments.
- **Continuous, Context-Aware Detection of Risk.** Financial threats often unfold over days or weeks, not seconds. Because **HCEP** is stateful, it retains memory of prior behaviors—such as anomalous logins, unusual account activities, or previous alerts—enabling pattern recognition across long temporal windows. This mirrors the way human experts track behavior over time and supports proactive detection of developing risks like account takeover, money laundering, or insider fraud.
- **Proactive, Probabilistic Risk Forecasting. HCEP's** continuous risk assessment assigns and updates conditional probabilities as new evidence arrives. Rather than static rules, the system mathematically evaluates how evolving behaviors influence overall risk—ideal for financial environments where early detection of emerging fraud or compliance violations is critical.
- **Integration of Structured and Unstructured Data.** Financial enterprises generate massive amounts of both structured (transactions, access logs) and unstructured (case notes, communications) data. **HCEP's** lexicon-based text analytics enable rapid extraction of meaningful signals from unstructured sources, enriching risk models with behavioral and contextual insights.
- **Transparent, Auditable, and Regulatory-Ready Intelligence.** Financial institutions must justify decisions to regulators, auditors, and internal oversight teams. **HCEP** offers explainability by design, producing clear causal chains that show why a risk was flagged—crucial in regulated domains such as anti-money laundering (AML), Know your Customer (KYC), and insider-risk programs.
- **Human-in-the-Loop Control for High-Consequence Decisions.** Because risk decisions in finance can have severe consequences, manual actions allow analysts to override or refine system-generated events, with changes propagating through the entire model. This ensures expert oversight and continuous model improvement.
- **Modern, Scalable, Cloud-Native Infrastructure. HCEP** leverages technologies such as Apache Kafka, Apache Flink, and Kubernetes, enabling real-time processing of high-volume financial data with elastic scalability. This architecture is essential for large enterprises with massive transaction throughput and global operations.

Together, these capabilities deliver trusted, explainable, proactive AI-driven security perfectly aligned with the operational and regulatory demands of modern finance.

Healthcare Enterprise Information Security

Cogility's HCEP model delivers a powerful Expert AI capability that aligns precisely with the needs of large healthcare enterprises, which must manage high volume, heterogeneous data while protecting sensitive patient information and meeting strict regulatory requirements.

- **Expert Modeled, Transparent Reasoning. HCEP** encodes human expertise into hierarchical event pattern models that mirror how analysts reason about complex behaviors and risks. This structure transforms massive data streams—clinical data, access logs, communications—into **explainable, high level intelligence** essential for regulated environments like healthcare.
- **Continuous, Context Aware Risk Detection.** Healthcare security threats often unfold gradually (e.g., credential misuse, insider access to Electronic Medical Records/EMRs, anomalous device behavior). Because **HCEP** is **stateful**, it retains memory of prior events and partially-matched patterns over long time periods, enabling accurate detection of evolving risks such as insider misuse or unauthorized access.
- **Proactive Risk Forecasting. HCEP** performs **Continuous Risk Assessment** using probabilistic reasoning. As new data arrives—such as unusual access sequences, repeated login anomalies, or changes in patient data handling—the system updates risk dynamically, enabling proactive mitigation rather than reactive response.
- **Integration of Clinical, Operational, and Behavioral Data.** Healthcare systems generate extensive *structured* data (Electronic Health Record/EHR logs, device telemetry, scheduling systems) and *unstructured* data (clinical notes, incident reports). HCEP's **lexicon-based analysis** extracts meaningful signals from unstructured text, enriching risk models with contextual and behavioral indicators.
- **Human in the Loop Oversight.** Security and compliance decisions in healthcare are high consequence. **HCEP** supports **manual actions** that allow analysts to revise inferences, correct patterns, and adjust risk assessments—ensuring expert oversight and improving model quality over time.
- **Scalable, Real Time Architecture.** Healthcare enterprises require high throughput, real time monitoring across distributed systems (hospitals, clinics, cloud services, devices). **HCEP's** architecture—built on **Kafka, Flink, and Kubernetes**—provides scalable, continuous analysis

Together, these features make HCEP exceptionally well suited to large healthcare organizations requiring **trusted, transparent, and proactive AI driven security.**

Summary & Conclusions

In environments such as national security, global banking, or healthcare, the cost of a wrong decision is catastrophic: Trust in – and the ability to explain a decision – are as critical as the decision itself.

A fundamental limitation of many modern AI approaches (like deep learning neural networks) is that they often function as "Black Boxes" — opaque systems where even the developers cannot explain why a specific risk was flagged. HCEP delivers transparent, trusted AI.

HCEP represents a paradigm shift from passive data monitoring to active Expert AI. By structuring data analysis through the lens of human expertise—using hierarchies for situation awareness, stateful memory for context, and computational modeling to support inference and continuous risk assessment—Cogynt.ai allows organizations to replicate the decision-making expertise of their best analysts. In high-consequence domains where trust is paramount, HCEP turns raw data into a continuous stream of explainable, actionable intelligence.

If you'd like to learn more about Cogynt.ai DIP for continuous decision support, please visit www.cogility.com or contact Cogility sales at sales@cogility.com.

HCEP Enables Transparent, Trusted AI

- **Transparency by Design:** Because models are built on explicit hierarchies defined by human experts, every output has a traceable lineage.
- **Auditability & Provenance:** When the system flags a risk, an analyst can "drill down" to see the specific causal chain. You do not just get a risk score; you get the evidence (e.g., "Risk is High because Pattern A and Pattern B converged").
- **Cognitive Alignment:** HCEP models mirror the "Hypothesis a Evidence a Conclusion" workflow of human analysts. This alignment fosters trust, drastically reducing the time required to validate alerts in time-sensitive situations.

COGILITY

Visit www.cogility.com to obtain more information and request an expert demo.

Cogility
15495 Sand Canyon Ave. #150
Irvine, CA. 92618

sales@cogility.com
+1 949.398.0015

02/26

References

Bower, G. H. (1970). Organizational factors in memory. Cognitive Psychology, 1, 18-46. https://www.sciencedirect.com/science/article/abs/pii/0010028570900034?via%3Dihub

Eckstein, M. & Collins, A. G. E. (2020). Computational evidence for hierarchically structured reinforcement learning in humans. Proceedings of the National Academy of Sciences (PNAS), 117(47), 29381-29389. https://www.pnas.org/doi/epdf/10.1073/pnas.1912330117

Greitzer, F. L., & Frincke, D. (2010). Combining Traditional Cyber Security Audit Data with Psychosocial Data: Towards Predictive Modeling for Insider Threat Mitigation. In Insider Threats in Cyber Security, ed. CW Probst, J Hunter, D Gollmann & M Bishop, pp. 85-113. Springer, New York. https://ieeexplore.ieee.org/document/8424651

Greitzer, F. L. (2025). Tracing the Evolution of Artificial Intelligence in Decision Support Systems. Cogility whitepaper. https://cogility.com/resources/ai-evolution-decision-intelligence-expert-systems/

Greitzer, F. L. (in press). From patterns to predictions: Insider risk modeling with a pattern-based behavioral analytic model. Paper submitted for publication.

Feigenbaum, E. A. (1977). The art of artificial intelligence: Themes and case studies of knowledge engineering. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, 1014–1029.

Luckham, D. (2002). The power of events: An introduction to complex event processing in distributed enterprise systems. Addison-Wesley.

Purcell B. A., & Kiani, R. (2016). Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy. Proceedings of the National Academy of Sciences,113(31), E4531-E4540. doi: 10.1073/pnas.1524685113

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decision making and system management. IEEE Transactions on Systems, Man, and Cybernetics.

Selfridge, O. G. (1959). Pandemonium: A paradigm for learning. In Symposium on Mechanization of Thought Processes (pp. 511–529). London: HMSO.