Using Data Science Tools: A Case Study to Identify Weak Signals

Volume 11, Number 4, 2026
ISSN 2957-7160 (Online)
ISSN 2957-7799 (Print)

DISCLAIMER:

OPEN publications are produced by Allied Command Transformation/Strategic Plans and Policy; however OPEN publications are not formal NATO documents and do not represent the official opinions or positions of NATO or individual nations. OPEN is an information and knowledge management network, focused on improving the understanding of complex issues, facilitating information sharing and enhancing situational awareness. OPEN products are based upon and link to open-source information from a wide variety of organizations, research centers and media sources. However, OPEN does not endorse and cannot guarantee the accuracy or objectivity of these sources. The intellectual property rights reside with NATO and absent specific permission OPEN publications cannot be sold or reproduced for commercial purposes. Neither NATO or any NATO command, organization, or agency, nor any person acting on their behalf may be held responsible for the use made of the information contained therein. The views expressed in this article are solely those of the authors and may not necessarily represent the views of NATO, Allied Command Operations, or Allied Command Transformation, or of their affiliated organizations.

Let us know your thoughts on "Using Data Science Tools: A Case Study to Identify Weak Signals " by emailing us at:

editor@openpublications.org

www.openpublications.org

CREDITS

CONTRIBUTING AUTHORS

Dr. Enis Karaarslan,
Assoc. Prof. Dr. Muğla Sıtkı Koçman University

OPEN CAPABILITY LEADER

COL Stefan Lindelauf

OPEN LEAD EDITOR

Mr. Jeffrey Reynolds

OPEN OPERATIONS MANAGER

COL Alexios Antonopoulos

ACTION OFFICER

Dr. Mehmet KINACI, LTC Jason STEWART

OPEN EDITORIAL REVIEW BOARD

LTC Tor-Erik Hanssen
CDR Alan Cummings
LTC Claus Slembeck
LTC Anders Wedin
LTC Gabor Farkas
LTC Emmanuel Senoussi
MAJ Ante Milas
LT CDR Cem Caglayan
Dr. Bohdan Kaluzny
Ms. Klodiana Thartori
Ms. Jedidja Van Keulen
Mr. Helmar Storm
Mr. Theodore Rubsamen
Mr. Christopher Hall
Mr. Benjamin Davenport
Mr. Eric Thompson

TECHNICAL EDITOR

Dr. Maureen Archer

ART DESIGNER

PO2 Allayah Klein

CONTENTS

EXECUTIVE SUMMARY

This paper tackles the challenge of identifying weak signals that are early, subtle indicators of emerging threats and opportunities critical for strategic foresight within NATO. It begins by reviewing traditional methods, such as horizon scanning and expert panels, noting their value while also acknowledging their limitations in terms of speed, scale, and objectivity.

The study presents a transformative approach using data science tools. A layered model demonstrates how Large Language Models (LLMs), Artificial Intelligence (AI), AI agents, and Digital Twin technologies can automatically scan vast, continuously updated open-source intelligence (OSINT) and web data streams, detecting patterns and assessing credibility beyond human capacity.

Comparisons with commercial platforms highlight relative strengths and gaps.

Practical examples across Social, Technological, Economic, Environmental, and Political (STEEP) domains illustrate how these technologies can enhance NATO's anticipatory capabilities. The report emphasizes the essential balance between automated detection and human judgment, positioning technology as a force multiplier rather than a replacement for human judgement.

Ultimately, this paper provides military and executive leaders with a clear understanding of how advanced tools can strengthen strategic warning, inform policy, and improve organizational resilience.

Keywords: Weak Signals, Strategic Foresight, Horizon Scanning, Early Warning, Artificial Intelligence (AI), Large Language Models (LLMs), AI Agents, Agentic AI, Digital Twin, STEEP Analysis (Social, Technological, Economic, Environmental, Political), Anticipatory Governance, Human-in-the-Loop (HITL), Decision Support, NATO Strategy, Risk Assessment, Data Mining, Natural Language Processing (NLP)

ABSTRACT

This study examines the challenge of identifying weak signals as early and subtle indicators of emerging risks and opportunities within NATO's strategic foresight activities. It first outlines the strengths and limitations of traditional practices such as horizon scanning and expert-based assessments, particularly in terms of scale, timeliness, and analytic consistency. The paper introduces a data-driven model that uses Large Language Models (LLM), artificial intelligence (AI) systems, autonomous agents, and Digital Twin technologies to process continuously evolving open-source intelligence and web-based information flows. By detecting patterns and evaluating credibility across big data environments, these tools extend the reach of human analysis. Examples drawn from the Social, Technological, Economic, Environmental, and Political (STEEP) domains illustrate how these capabilities can reinforce NATO's anticipatory posture. This study emphasizes the need for a balanced relationship between automated detection and expert judgement, positioning advanced technologies as enablers that support rather than replace human decision-making. The findings offer military and executive leaders a concise perspective on how emerging analytical capabilities can strengthen strategic warning, inform policy, and enhance institutional resilience.

INTRODUCTION

The rapid pace of technological, social, and environmental change presents considerable challenges for strategic foresight and risk management. In this complex landscape, the ability to identify and interpret weak signals has become critical for maintaining organizational resilience (Ansoff, 1975) and operational readiness. Weak signals are subtle, early indicators of emerging trends or events that are often ambiguous, fragmented, and difficult to distinguish from background noise. Their core value lies in enabling the anticipation of changes in complex environments through careful observation and contextual interpretation, turning otherwise hidden patterns into actionable insights (Van Veen and Ortt, 2021).

In military cyber threat intelligence, weak signals (minor anomalies, unusual patterns, low-level probing, etc) can reveal threats before they escalate into visible attacks. Effective detection depends on interpretive processes that assess subtle cues to anticipate risks, allowing intelligence teams to manage uncertainty proactively (Ćwik and Świerszcz, 2019). Cyber operations are particularly challenging because the same properties that make weak signals observable also permit deception; attackers may conceal tracks, mimic benign behaviour, or disguise probing as harmless activity (Gartzke and Lindsay, 2015). Kello's concept of the virtual weapon illustrates how cyber operations exist in a grey zone between peace and war, creating continuous, low-level tension that amplifies the strategic importance of weak-signal monitoring (Kello, 2019). Organizations can anticipate adversary intentions and transform latent risks into informed, proactive decision-making by integrating careful detection of weak signals with awareness of potential deception.

The concept of weak signals has been studied for decades. The challenge is developing practical methods for their consistent detection and interpretation (Hiltunen, 2010). The idea that early, subtle indicators can help predict emerging trends is well-established. However, identifying and acting on these signals in a systematic way continues to be a key challenge in studies and strategic planning.

Horizon scanning methodologies were introduced as a solution first, in which individuals and groups analyse vast amounts of information, identifying changes, anomalies or patterns that might suggest future developments. These methods rely heavily on human judgment. The involvement of subject matter experts (SMEs) with their deep knowledge, often organized into advisory groups or panels was the crucial element of these methodologies. Experts can recognize significant information that might be overlooked by automated systems or non-specialists (Tetlock, 2005).

It is also important to recognize the limitations of human-centric methodologies. Studies on group dynamics show that expert panels are open to cognitive biases, such as groupthink or overconfidence. These can lead to dismissing unconventional signals that are valuable (Janis, 1972). The composition of advisory groups can introduce selection biases, potentially overlooking signals that fall outside traditional frameworks (Schoemaker et al, 2013). These biases are particularly important in a military context, where new threats and emerging risks do not always

fit within established paradigms. The effective construction of expert-led groups is critical as it can improve signal detection by pooling diverse perspectives and knowledge biases (Morgan, 2014).

Traditional methods often miss weak signals due to their rarity, novelty, or diffuse character (Mühlroth & Grottke, 2018). Moreover, these methodologies alone cannot meet today's needs in speed especially with the amount of information to be analysed in scale. There is a challenge of detecting weak signals via data science and developing methods that are sensitive, robust, and validated. Emerging technologies can be used to improve current processes and better prepare for emerging challenges. The volume and variety of available data in the big data landscape provides new opportunities to be discovered.

Recent studies demonstrate promising directions. Text mining and NLP have been applied to heterogeneous sources to identify signals that later proved significant (Griol-Barres et al., 2020). Patent analysis shows that weak signals may be more predictive of breakthrough innovations than already strong signals (Bzhalava et al., 2022). Literature-based clustering and graph convolutional networks have been used to detect signals and forecast their growth over time (Ha et al., 2023).

Digital twin technologies are emerging as complementary tools. Studies demonstrate their value for anomaly detection in physical and industrial systems, with applications from safety monitoring to forecasting operational bottlenecks using explainable AI (Tancredi et al., 2022; Calvo-Bascones et al., 2023; Iyer et al., 2025).

Large Language Models (LLMs), AI agents and agentic AI (incorporating multi-agent systems and collaborative agent teams) demonstrate potential in zero-shot and weakly supervised environments. In this context, zero-shot capabilities enable models to perform tasks without specific prior training or task-specific examples (Brown et al., 2020). While current performance in anomaly detection remains inconsistent(Alnegheimish et al., 2024), these systems offer advantages where labelled data is scarce. Recent applications include LLM-based frameworks designed to detect early credibility signals in online content, thereby improving the identification of emerging misinformation (Leite et al., 2025). Such multi-agent architectures facilitate the decomposition of complex tasks by allowing specialized agents to function as a unified team.

These academic studies and emerging products on the market show that combining data science, digital twins, and emerging AI methods can be used to strengthen weak-signal detection, providing decision makers with earlier and more reliable insights for strategic warning and policy planning.

WEAK SIGNAL DETECTION MODEL OF ADVANCED DATA SCIENCE TOOLS

We propose a weak signal detection model that is structured as a five-layer system built upon a secure data foundation (see Figure 1). The architecture is designed to support NATO by enabling earlier warning, improved policy planning, and more resilient strategic foresight.

Each layer performs a distinct transformation of information, ensuring that weak, ambiguous, or fragmentary signals are progressively strengthened into decision-quality outputs. This progression can be conceptualized as a signal maturing from its initial state (S0) through successive stages of validation and enrichment (S1, S2, S3). The following subsections introduce the fundamentals of a given layer. Appendix A compares well-known analytical tools that can be used for each layer, and Appendix B lists key metrics of each layer.

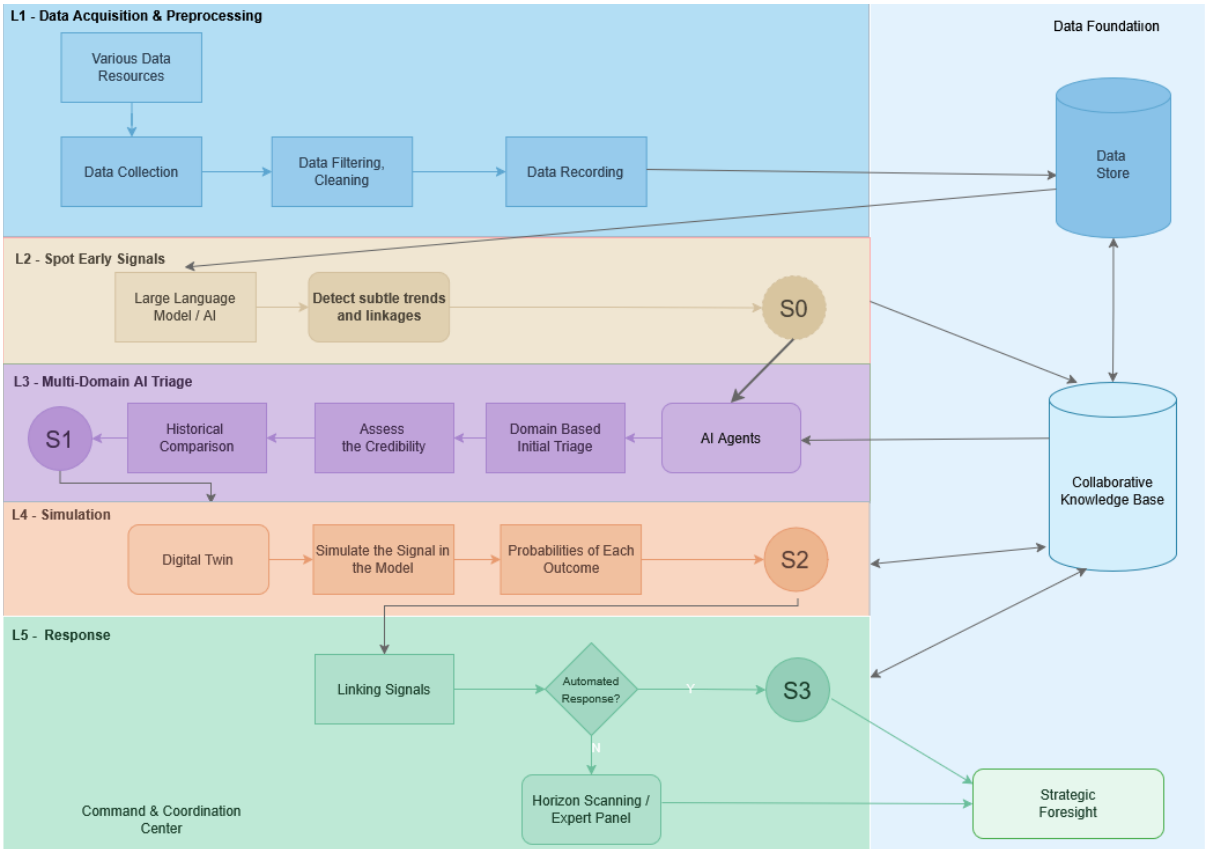

Fig 1. Layered Model for Weak Signal Detection

The model aims to provide a systematic, reproducible, and scalable framework that orchestrates the signal detection process to form a strategic foresight in a coordinated and strategic manner. It is built upon specialized analytical units and is designed to evolve through continuous refinement, thereby enhancing its overall effectiveness. Institutional learning and collective memory are interlinked, with organizational awareness achieved through the systematic connection of dispersed insights. A collaborative knowledge base serves as the repository for collected experience and shared understanding.

The core functions are envisaged to operate within a central command and control structure; however, the framework also supports distributed deployment across participating states and institutions. Each organization may establish its own system, while federated mechanisms enable the integration of national and institutional nodes into a global framework. This will strengthen both interoperability and resilience, while preserving the confidentiality and integrity of sensitive information.

2.1. Data Foundation

The proposed system is built on a dual infrastructure: a secure data store for structured information and a collaborative knowledge base to support strategic foresight. Sensor feeds, intelligence reports, and open-source information are ingested, time-stamped, and stored immutably in the data store. Each entry records its origin (first source), timestamp, classification, confidence, and provenance (processing and custodial history) to ensure full traceability.

The infrastructure draws on these resources across its layers, providing auditable and verifiable inputs for all operations. Three options are considered for organizing information:

• Data Lakes for raw and unstructured data,

• Relational Databases for structured metadata and operational records,

• Vector or Graph Databases to represent complex relationships, linkages, and historical comparisons.

Together, these components form the foundation of the collaborative knowledge base, enabling the system to capture, contextualize, and connect weak signals across domains, thereby enhancing situational awareness and strategic foresight.

2.2. Data Acquisition & Preprocessing Layer (L1)

This layer orchestrates the collection, filtering, cleaning, and recording of data. Data pipelines ensure a structured flow of information. They are implemented through either ETL (Extract, Transform, Load) or ELT (Extract, Load, Transform) approaches. The choice between these approaches depends on whether data is processed before or after storage. Such an infrastructure will have to deal with big data and involve large-scale data processing. Large-scale data environments rely on distributed processing frameworks, while smaller-scale or pilot systems may utilize lightweight solutions. Raw data may consist of fragmented or misleading contents. The aim is to transform it into a reliable and consistent form. Steps will include the following:

• Multiple data formats are harmonized into a unified schema,

• Time and language normalization, entity recognition, and integrity checks are applied according to established standards,

• Data cleaning reduces noise and safeguards quality through processes such as normalization and deduplication.

2.3. Spot Early Signals Layer (L2)

This layer is designed to enable the early detection of emerging signals within the data repository. Techniques such as statistical anomaly detection, clustering, and similarity searches help uncover linkages, trends, and clusters that may signal emerging developments.

Both proprietary and open-source Large Language Model (LLM) solutions exist for this purpose, ranging from scalable enterprise analytics suites to specialized libraries developed for research communities. LLMs provide application programming interfaces (API) that allow integration and querying. Open source LLMs can be installed on each organization's own infrastructure. Algorithms for anomaly detection, pattern recognition, and hybrid fusion methods

can be applied.
The key requirement is the ability to process large volumes of heterogeneous data and reveal low-visibility signals that would otherwise remain undetected. Outputs of this layer are candidate signals that will be checked in the following layers.

2.4. Multi-Domain AI Triage Layer (L3)

This layer focuses on domain-specific assessment and credibility scoring, thereby investigating the quality of signals detected. This layer will use autonomous codes that are called AI agents, agentic AI (multi-agent) systems. The process includes:

• Initial triage by domain,

• Credibility assessment through historical comparison, and

• Cross-checking against existing knowledge bases.

By structuring the workflow into specialized analytical streams (e.g., technical, financial, geopolitical), the model ensures that weak signals are assessed in context and understanding if it is a considerable signal that should be investigated in detail.

2.5. Simulation Layer (L4)

Simulation provides the means to test the potential impact of signals through digital representations of relevant systems. Digital twins and scenario models allow organizations to assess possible outcomes, assign probabilities, and understand causal linkages. Graph-based models can further illustrate interconnections between signals, events, and entities, helping decision-makers grasp the possible effects before they happen.

2.6 Response Layer (L5)

The final layer delivers strategic foresight by transforming candidate signals into decision-quality outputs. Automated responses can be executed without delay where pre-approved rules exist for time-critical and low-risk scenarios. In all other cases, the system will operate as human-in-the-loop (HITL). Human analysts will examine candidate signals, test hypotheses, validate or discard findings, and provide structured feedback that strengthens institutional memory.

This human oversight is enabled through analyst workbenches, approval and release workflows, and feedback mechanisms. These tools not only ensure accountability and governance but also allow domain expertise to be systematically incorporated into AI model retraining and refinement. In this way, the system continuously improves while maintaining human control.

Response mechanisms may include forecasting modules, decision-support engines, or pre-approved response templates. Outputs are delivered in appropriate formats such as secure dashboards, concise alert briefs, or structured policy notes. These will highlight what has occurred, why it matters, the assessed confidence level, and recommended courses of action. Every product is linked to its full provenance trail, ensuring defensibility in strategic deliberations and enabling trusted coordination with Allies.

COMMERCIAL PRODUCTS VERSUS WEAK SIGNAL MODEL

Certain commercial products are applicable to weak signal detection, while not built for this purpose, and most are focused on cyber security. Table 1 provides a comparison based on online white papers and public LLMs. The evaluation of weak-signal detection solutions can be framed along several key dimensions such as data breadth, early warning capability, capturing effectiveness in detecting subtle or emerging signals, automation and AI support, actionability for leaders, scalability, and integration.

Solution	Main Purpose	Best For	Key Limitation
Palantir Gotham / Foundry	Connecting and analysing many different types of data to see the "big picture".	Strategic decision-making when you need to combine data from many separate sources.	- High cost and complexity - Requires significant setup.
Recorded Future	Providing real-time cyber threat alerts from the open and dark web.	Early warning of cyberattacks and online threats.	Focus only on cyber threats, not broader intelligence.
IBM i2 Analyst's Notebook	Visually mapping hidden connections between people, organizations, and events.	Criminal investigations and link analysis (e.g., counter-terrorism).	Relies heavily on manual analysis.
Microsoft Sentinel	Cloud-based security monitoring to detect unusual user and system behaviour.	Protecting large corporate IT networks from insider threats and attacks.	Primarily for IT security, not strategic intelligence.
Darktrace / Exabeam	Using AI to automatically detect anomalies and potential threats on the network.	Finding advanced, stealthy cyber attacks that traditional tools miss.	- Can produce false alarms - Focused on network data.

Table 1. Commercial platforms

While a full performance assessment is beyond the scope of this study, Table 2 presents a comparison of commercial platforms against the proposed five-layer model. Commercial tools are typically focused on cyber or operational tasks, provide only limited cross-domain fusion, and often provide partial or vendor-dependent provenance. In contrast, the proposed model delivers comprehensive cross-domain integration, full traceable provenance, mission-tuned sensitivity, analyst-in-the-loop oversight, and explainable, decision-quality outputs explicitly designed to support strategic warning and policy planning.

Dimension	Commercial Platforms	Five-Layer Weak-Signal Model
Decision focus	Optimized for cyber/operational tasks or analyst queries	Explicitly tuned to provide strategic warning and decision-quality outputs
Cross-domain fusion	Some tools need heavy customization; others are domain-specific	Built-in layered design integrates multiple domains consistently
Provenance & audit	Partial lineage; varies by vendor	Full, immutable chain-of-custody for every signal and transformation
Human-in-the-loop	Optional or add-on analyst consoles	Core component: analyst workbenches, approval workflows, feedback loops
Sensitivity vs. false alarms	Often optimized for low false positives in Security Operation Center (SOC) context	Mission-weighted: high-recall signals supported by Human-in-the-loop triage
Explainability	Some Machine Learning decisions opaque	Every output includes rationale and traceable provenance for accountability
Strengths	- Fast deployment for specific domains, - Mature vendor support, - Rich visual analytics	- Tailored to strategic foresight, - Flexible for cross-domain integration, - Fully auditable, - Supports institutional memory
Limitations	- Limited to pre-defined workflows, - Partial auditability, - Often opaque ML	Requires initial setup of layered architecture and analyst resources
General Evaluation	Effective for tactical and operational intelligence; may miss low-visibility signals	Optimized for early warning and policy planning; explicitly designed for strategic decision-making

Table 2. Commercial Platforms vs Five-Layer Model

RECOMMENDATIONS FOR NATO

This section presents the STEEP Analysis of AI-Enhanced Capability, the importance of human involvement in the process, and strategies for selecting the best software approach.

4.1. STEEP Analysis of AI-Enhanced Capability

STEEP Analysis of AI-Enhanced Capability is shown in Table 3, in which the key contribution and the impact on NATO's anticipation ability are summarized. Below the table are definitions of each dimension.

Social dimension deals with the rise of hyper-connected societies and information warfare. Examples are deepfake campaigns targeting elections or coordinated botnets amplifying divisive narratives to weaken social cohesion ahead of a hybrid attack. AI-enhanced capability can ensure enhanced OSINT (Open-Source Intelligence) (Bennett & Livingston, 2018) to detect orchestrated influence operations. It can correlate sentiment analysis with economic data or event triggers to forecast periods of heightened vulnerability of a region.

Technological dimension deals with proliferation of dual-use technologies and Internet of Things(IoT). Examples include the widespread

use of drones (that have backdoors) in critical infrastructure or vulnerabilities in a new satellite communication protocol. AI is also transforming international security competition, so there is a need for technological horizon scanning (Horowitz, 2018). An AI-enhanced capability can scan for the weaponization potential of emerging technologies, predicting which critical infrastructures are most likely to be targeted and the probable methods of attack.

Dimension	Key Contribution	Impact on NATO's Anticipation Ability
Social	Real-time analysis of social media to detect disinformation flows	Early warning with - Predicting Social Unrest - Enhanced OSINT (Open-Source Intelligence)
Technological	Scanning for the weaponization potential of emerging technologies	Early detection with - Threat Horizon Scanning - Cyber Vulnerability Forecasting
Economic	Detection of anomalies in global supply chains and financial flows	Identification of economic coercion and energy weaponization; assurance of defence-industrial continuity. - Financial Intelligence (FININT) Fusion - Supply Chain Risk Mapping
Environmental	Modelling the destabilizing effects of environmental change (drought, migration).	Forecasting climate-induced humanitarian crises and new operational domains with - Environmental Early Warning - Resource Conflict Prediction
Political	Monitoring subtle shifts in adversary diplomatic activity.	Detection of geopolitical intent and Alliance realignments early with - Diplomatic Signal Detection - Alliances & Adversary Modelling

Table 3: STEEP Analysis of AI-Enhanced Capability: Sharpening NATO's Strategic Foresight

Economic dimension deals with weaponization of economic interdependence (Farrell & Newman, 2019). Examples include the strategic acquisitions of European port facilities by state-owned enterprises, or subtle shifts in energy supply patterns and pricing that could be used as coercive leverage. AI-enhanced capability will help in detecting patterns of economic coercion or pre-positioning for sanctions evasion. It can identify single points of failure in the NATO-wide defence supply chain and help in predicting disruptions from geopolitical events, greatly enhancing resilience planning.

Environmental dimension deals with climate change as a threat (Mach, 2019). Examples include the rapid Arctic ice melt or increased frequency of droughts leading to mass migration and instability. AI-enhanced capability will help in forecasting climate-induced humanitarian crises and new operational domains (e.g., the Arctic). It can anticipate future conflicts over resources, allowing NATO to engage in preventive diplomacy and readiness planning.

Political dimension deals with the erosion of the rules-based international order and the rise of authoritarian alignment. Examples include the increased diplomatic coordination and joint military exercises between adversarial states. AI-enhanced capability will help in detection of geopolitical intent (King & Lowe, 2003) and Alliance realignments. It can detect subtle shifts in rhetoric, revealing strategic intent and potential pivots in policy before overt actions are taken. It can predict the formation of adversarial blocks or proxies.

4.2. Recommendations on Human Involvement

Human-in-the-loop integrates human judgment into the AI process, whereas human-out-of-the-loop denotes full automation aimed at rapid automatic responses. This approach still involves disadvantages and challenges, as summarized in Table 4. An optimal solution is a hybrid model in

Aspect	Human-Out-of-the-Loop (Full Automation)	Human-In-The-Loop (Human Judgment)
Advantages	• **High Speed:** Responds in milliseconds. • **Scalability:** Can monitor and respond to millions of events simultaneously. • **24/7 Consistency:** Operates continuously without fatigue. • **Perfect for Defined Basic Rules:** Excellent for executing pre-authorized, simple playbooks.	• **Contextual Understanding:** Interprets nuance, irony, and political intent. • **Moral & Ethical Reasoning:** Can reason with the laws of armed conflict and ethical principles. • **Alliance Cohesion:** Can form crucial consensus-building among member states. • **Accountability:** Establishes a clear chain of command and responsibility.
Disadvantages	• **Lacks Nuance:** Cannot understand strategic deception or political context. • **Accountability Void:** Difficulty attributing political and legal responsibility for erroneous actions. • **Exploitable:** Can be predictable; adversaries may learn to exploit its automated rules.	• **Decision Lag:** Slower response time due to necessary human deliberation. • **Inconsistency:** Human decisions can vary due to bias, fatigue, or experience. • **Limited Coverage:** Humans cannot process the volume and speed of modern data feeds.
Key Challenges	• **Accidental Escalation:** A misidentification could trigger an unintended and disproportionate conflict. • **Setting Boundaries:** Defining the strict, politically-approved Rules of Engagement (ROE) for each situation is extremely difficult. • **Trust:** Requires immense confidence in the system's accuracy, which is hard to achieve and verify.	• **Signal-to-noise problem:** Critical alerts may be overlooked amidst thousands of false positives. • **Information Overload:** Synthesizing big data into a timely decision is a major cognitive burden. • **Training:** Requires continuous awareness training for operators to understand and trust alerts without over-relying on them.

Table 4. Comparison of human-out-the-loop and human-in-the-loop systems

which human-in-the-loop applies to defined tactical domains, with humans setting the rules while automation executes them under continuous monitoring. Samples of its usage are automated cyber defence such as blocking malicious Internet Protocol (IP) addresses and electronic warfare countermeasures. For strategic decisions, human-in-the-loop is preferable, with AI recommending courses of action and humans retaining final verification and command.

4.3. Recommendations on Selecting Software Approach

Sample software approaches for each required function are presented in Table 5. Implementing each function demands varying levels of expertise in areas such as distributed systems, machine learning, natural language processing (NLP), data engineering, and visualization. Organizations often require smaller internal teams and can delegate more operational responsibility to the vendor when they choose proprietary solutions.

Function	Free Sw. Approach	Proprietary Sw. Approach	Expertise Required
Data Processing & Storage	Apache Spark, PostgreSQL, Elasticsearch	Splunk, Databricks, Snowflake	High (Distributed systems)
Machine Learning Framework	Scikit-learn, TensorFlow, PyTorch	DataRobot, H2O.ai, SageMaker	Medium-High (ML engineering)
Anomaly Detection	PyOD, Isolation Forest	Azure Anomaly Detector, AWS Lookout	Medium (Statistics/ML)
Natural Language Processing	spaCy, NLTK, Transformers	OpenAI API, Google Cloud NLP	Medium (NLP specialization)
Orchestration & Workflow	Airflow, Dagster, Prefect, Argo, Kubeflow, Flyte, Luigi	AWS Step Functions, Azure Logic Apps, GCP Workflows, Databricks Workflows, ADF, Cloud Composer, MWAA, Dataiku	High (DevOps/Data engineering)
Visualization & Dashboarding	Grafana, Kibana, Superset	Tableau, Power BI, Grafana Cloud	Medium (Data visualization)

Table 5: Data Science Tools Comparison for Weak Signal Detection

A critical strategic decision involves selecting a licensing model, either free software or proprietary, for each necessary function. The specific challenges associated with each model are detailed in Table 6. Free software is fundamentally characterized by the principles of freedom and sovereignty. This freedom encompasses not only the right to use, distribute, and modify the code but also provides freedom from vendor lock-in. This freedom encompasses not only the right to use, distribute, and modify the code but also provides freedom from "vendor lock-in." Vendor lock-in creates a strategic risk by creating a dependency on a product where transitioning away is incredibly difficult, expensive, or disruptive. However, achieving these benefits with free software typically involves a longer implementation time and a higher requirement for in-house expertise. In return, organizations gain maximum customization flexibility, as they retain the right to modify the code to meet their precise needs.

Challenge	Free Software Approach	Proprietary Solution
Licensing Freedom	High (Use, study, modify, share without restrictions)	Restricted (Subject to EULA, licensing fees, audits)
Vendor Lock-in	None (You own and control the stack)	High (Dependent on vendor's roadmap, pricing, and existence)
Implementation Time	Long (6-12 months for full integration)	Short (2-4 months with professional services)
Initial Cost	Low (no licensing fees)	High (licensing fees)
Long-term TCO (Total cost of ownership)	Predictable & Controllable (Costs are primarily expertise-based)	Predictable but Recurring (mandatory support and update fees)
Required Expertise	High (in-house data engineers, ML experts)	Medium (vendor training and support available)
Customization Flexibility	Complete (Freedom to modify any part of the software)	Limited to vendor capabilities and roadmap
Security & Compliance	Transparency (You can verify the code yourself)	Delegated Responsibility (Reliant on vendor's claims and certifications)
Time to Value	Slow (build from scratch)	Fast (pre-built solutions and templates)

Table 6: Implementation Challenges: Free Software vs. Proprietary Solutions

Additionally, the free software model adopts an approach of "security through transparency", meaning the software code can be audited by anyone. This transparency is a powerful feature for verification, but also constitutes a liability, as it necessitates internal capability to manage and respond to the findings. Key decision factors for leadership are summarized in Table 7. The final choice for an organization depends on its overarching strategy, the expertise of its team, project timelines, budget, and specific needs for customization and compliance.

Factor	Choose Free Software When	Choose Proprietary Software When
Strategy	You want complete control and ownership	You want to focus on your core functions
Team	You have strong in-house technical expertise	You have limited technical resources
Time	You have a long-term timeline (more than a year)	You need a solution quickly (less than 6 months)
Budget	You have limited initial capital but can fund ongoing costs	You have capital budget and prefer predictable operating expenses
Customization	You have unique requirements not met by commercial tools	Your needs align well with available commercial solutions
Compliance	You have specialized security/compliance needs	You can leverage vendor certifications and audits

Table 7: Key Decision Factors for Leadership

KEY RESULTS

Large Language Models, artificial intelligence systems, autonomous agents, and Digital Twin technologies can be used together to detect weak signals. The key results are given in the following subsections.

Gaps, Limitations, and Opportunities

Digital twin (DT) environments currently provide promising tools for simulating and monitoring physical or operational systems. However, most DT applications focus on detecting clear anomalies rather than subtle, early weak signals. Sensitivity, false-positive rates, and lead time are often not optimized for faint or gradually emerging indicators, and reliance on traditional machine learning limits adaptability to novel signal types. Large Language Models (LLMs), AI agents and agentic AI demonstrate potential in zero-shot or weakly supervised settings. This capability allows systems to generalize across diverse domains by utilizing natural language instructions or limited high-level guidance. However, they face challenges, including high computational cost, limited context windows, potential hallucinations, and suboptimal detection performance compared to specialized supervised models. Their ability to detect very early weak signals in high-noise environments in the presence of minimal data remains underexplored.

Future Research Directions

Future research should explore the integration of digital twin frameworks with LLMs and agentic AI, which offers a promising path forward. As in the proposed hybrid model, digital twins can provide realistic simulations of system behaviour, LLMs can interpret unstructured or semantic cues, and AI agents can explore trajectories and amplify weak signals. Key research priorities include reducing lead time for signal detection, improving the balance between false positives and negatives, and enhancing domain transferability and robustness in minimally supervised pipelines.

Key Implications of the Five Layer Model:

Key implications of the proposed model include:

• Timely foresight, enabling earlier detection of adversary intent or systemic risks,

• Evidence-based policy, with decisions supported by traceable, multi-layered analyses,

• Resilience under contestation, ensured through layered processing maintained database integrity even in degraded conditions, and

• Governance and trust, achieved via human oversight and full provenance to establish credibility in sensitive contexts.

Practical Recommendations:

Practical recommendations include leveraging free software solutions and avoiding vendor lock-in. NATO members, in collaboration with industry and academic partners, can implement the proposed five-layer model, which ingests data from multiple sources into a canonical data foundation, performs ensemble detection, and enforces analyst workbench and approval workflows. While full implementation may be time-consuming, module-by-module activation is feasible and expected to yield significant benefits. Success should be evaluated based on mission value rather than detection metrics alone, with decision contribution and provenance completeness tracked as primary key performance indicators.

CONCLUSION – SO WHAT FOR NATO

Weak signals can precede significant strategic, operational, or security developments, yet they are often difficult to detect using traditional methods due to their ambiguity and low visibility. Artificial intelligence offers NATO the ability to enhance the detection of these early indicators of emerging threats. AI tools can process large volumes of complex data, identify unusual behaviours, recognize emerging patterns, and forecast potential developments before they fully materialize, enabling decision-makers to act proactively rather than reactively.

In the era of big data, AI adoption is inevitable, but speed must be balanced with wisdom to ensure that NATO's actions remain effective, legitimate, and ethically sound. Human judgment should remain central in decision-making, especially for strategic actions that require contextual understanding, moral reasoning, and collective oversight by the North Atlantic Council. A hybrid approach is therefore optimal; automation can support defensive, time-critical tasks within strict boundaries, while human experts should guide high-level, strategic decisions.

AI solutions, particularly those based on free software and open-source platforms, can be implemented without requiring prohibitively expensive infrastructure, making them practical and widely accessible. Leveraging open-source and preferably free software licenced tools also enables NATO to collaborate with academic institutions and the broader Linux community, benefiting from innovation, transparency, and ongoing community-driven improvements.

By integrating AI-driven weak signal detection into intelligence workflows, NATO can enhance situational awareness, anticipate adversary intentions, strengthen operational readiness, and maintain strategic stability, all while fostering a flexible, sustainable, and community-supported technological ecosystem.

REFERENCES

Alnegheimish, S., Nguyen, L., Berti-Equille, L., & Veeramachaneni, K. (2024). Large language models can be zero-shot anomaly detectors for time series?. arXiv preprint arXiv:2405.14755.

Ansoff, H. I. (1975). Managing strategic surprise by response to weak signals. California management review, 18(2), 21-33. https://doi.org/10.2307/41164635

Bennett, W. L., & Livingston, S. (2018). The disinformation order: Disruptive communication and the decline of democratic institutions. European journal of communication, 33(2), 122-139.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. Advances in neural information processing systems, 33, 1877-1901.

Bzhalava, L., Kaivo-Oja, J., Hassan, S. S., & Gerstlberger, W. D. (2022). Identifying entrepreneurial discovery processes with weak and strong technology signals: a text mining approach. Open Research Europe, 2, 26.

Calvo-Bascones, P., Voisin, A., Do, P., & Sanz-Bobi, M. A. (2023). A collaborative network of digital twins for anomaly detection applications of complex systems. Snitch Digital Twin concept. Computers in Industry, 144, 103767.

Ćwik, B., & Świerszcz, K. (2019). Efficiency of Weak Signals' Detection: Interpretive Aspects of Threat Signal Perceiving.

Erik Gartzke & Jon R. Lindsay (2015) Weaving Tangled Webs: Offense, Defense, and Deception in Cyberspace, Security Studies, 24:2, 316-348, DOI:10.1080/09636412.2015.1038188.

Farrell, H., & Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. International security, 44(1), 42-79.

Griol-Barres, I., Milla, S., Cebrián, A., Fan, H., & Millet, J. (2020). Detecting weak signals of the future: A system implementation based on text mining and natural language processing. Sustainability, 12(19), 7848.

Ha, T., Yang, H., & Hong, S. (2023). Automated weak signal detection and prediction using keyword network clustering and graph convolutional network. Futures,152, 103202.

Hiltunen, E. (2010). Weak signals in organizational futures learning. Helsinki School of Economics.

Horowitz, M. C. (2018). Artificial intelligence, international competition, and the balance of power (May 2018).

Iyer, S. V., Sangwan, K. S., & Dhiraj. (2025). A cognitive digital twin for process chain anomaly detection and bottleneck analysis. Journal of Industrial and Production Engineering, 42(1), 65-87.

Janis, I. L. (1972). Victims of groupthink: A psychological study of foreign-policy decisions and fiascoes.

Kello, L. (2019). The Virtual Weapon and International Order. In The Virtual Weapon and International Order. Yale University Press.

King, G., & Lowe, W. (2003). An automated information extraction tool for international conflict data with performance as good as human coders: A rare events evaluation design. International Organization, 57(3), 617-642.

Leite, J. A., Razuvayevskaya, O., Bontcheva, K., & Scarton, C. (2025). Weakly supervised veracity classification with LLM-predicted credibility signals. EPJ Data Science, 14(1), 16.

Mach, K. J., Kraan, C. M., Adger, W. N., Buhaug, H., Burke, M., Fearon, J. D., ... & von Uexkull, N. (2019). Climate as a risk factor for armed conflict. Nature, 571(7764), 193-197.

Morgan, M. G. (2014). Use (and abuse) of expert elicitation in support of decision making for public policy. Proceedings of the National academy of Sciences, 111(20), 7176-7184.

Mühlroth, C., & Grottke, M. (2018). A systematic literature review of mining weak signals and trends for corporate foresight. Journal of Business Economics, 88(5), 643-687.

Schoemaker, P. J., Day, G. S., & Snyder, S. A. (2013). Integrating organizational networks, weak signals, strategic radars and scenario planning. Technological Forecasting and Social Change, 80(4), 815-824. https://doi.org/10.1016/j.techfore.2012.10.020

Tancredi, G. P., Vignali, G., & Bottani, E. (2022). Integration of digital twin, machine-learning and industry 4.0 tools for anomaly detection: An application to a food plant. Sensors, 22(11), 4143.

van Veen, B. L., & Ortt, J. R. (2021). Unifying weak signals definitions to improve construct understanding. Futures, 134, 102837.

APPENDIXES

A. Free Software vs. Proprietary Software for Layered Model

Layer / Function	Primary Tool Category	Free Software	Proprietary Software	Key Metric for Comparison
Data Foundation	Data Lakehouse Format	Apache Iceberg, Apache Hudi	Delta Lake on Databricks, Snowflake	Query Performance (QPS), Scalability Limits, ACID Compliance, Backup/Recovery Time
	Vector Database	Weaviate, Chroma, Qdrant	Pinecone, AWS Aurora PG Vector	
	Graph Database	Neo4j Community, JanusGraph	Neo4j Enterprise, Amazon Neptune	
	Relational Database	PostgreSQL	Oracle	
L1: Data Acquisition & Preprocessing	Data Pipeline (DP) Orchestration	Apache Airflow, Prefect	Databricks Workflows, Azure Data Factory	Data Pipeline Reliability, Processing Speed (GB/minute), Data Quality Score, Storage Cost per TB/month
	Data Processing	Apache Spark	Databricks, Google Dataflow	
	Data Quality	Great Expectations	Monte Carlo, Soda Cloud	
L2: Spot Signals	LLMs & Embeddings	Llama 3, BGE-M3	OpenAI GPT	Signal Detection Recall, Precision, Embedding Quality (MTEB), Inference Latency, Cost per 1M Tokens
	Frameworks	Transformers, SentenceTransformers	OpenAI text-embedding-3-large	
L3: AI Triage	AI Agent Framework	LangChain, LlamaIndex	LangChain Commercial Offerings, Microsoft Autogen	Triage Accuracy, False Positive/Negative Rates, Processing Throughput, Domain Adaptation Time
	Evaluation	PyTorch, Hugging Face	Galileo, Arthur	
L4: Simulation	Simulation Library	SimPy, Mesa	AnyLogic, Simul8	Simulation Speed, Model Accuracy, Scenario Coverage, Computational Cost per Simulation
	Causal Inference	DoWhy	Commercial SAS/SPSS modules	
	Visualization	Grafana (for dashboards), Plotly Dash	Tableau, Microsoft Power BI	
L5: Response	Decision Engine	Custom rules + Drools	SAS Decision Manager, Pega	Decision Latency, Expert Consensus Time, Strategy Implementation Rate, Audit Trail Completeness
	Collaboration	JupyterHub + Apache Superset	Databricks Notebooks, Power BI	
	Forecasting	Prophet, statsmodels	Azure Anomaly Detector, Amazon Forecast	
	Workflow	Apache Airflow	ServiceNow, Microsoft Power Automate	

B. Key Metrics of Each Layer

DB + Layers	Key Metric	Definition	Priority Level
Data Acquisition & Preprocessing	Query Performance	Queries per second at peak load	High
	Data Freshness	Time from data arrival to query availability	Critical
	Scalability Limits	Maximum data volume before performance degradation	High
	ACID Compliance	Level of transactional reliability	Critical
	Backup/Recovery Time	Recovery time after failure	Medium
	Vector Search Accuracy	Recall@K for similarity searches	High
	Graph Traversal Speed	Time for 5-hop relationship queries	High
L1: Data Acquisition & Preprocessing	Pipeline Reliability	Uptime percentage of data ingestion pipelines	Critical
	Processing Speed	Amount of data processed per minute (GB)	High
	Data Quality Score	Composite metric for accuracy/completeness	Critical
L2: Spot Signals	Signal Detection Recall	Percentage of true signals identified	Critical
	Precision	Percentage of identified signals that are meaningful	High
	Embedding Quality	MTEB benchmark score	High
	Inference Latency	Embedding generation time (ms)	Medium
L3: AI Triage	Triage Accuracy	Percentage of signals correctly prioritized	Critical
	False Positive Rate	Percentage of low-importance signals flagged as high-priority	High
	Processing Throughput	Number of signals processed per hour	High
	Domain Adaptation Time	Time to adapt to new threat domains	Medium
L4: Simulation	Simulation Speed	Time to complete 1000 scenario iterations	High
	Model Accuracy	Percentage deviation from historical outcomes	Critical
	Scenario Coverage	Number of distinct scenarios model can handle	Medium
	Computational Cost	Cost per simulation hour	Medium
L5: Response	Decision Latency	Time from signal detection to finalized response	Critical
	Expert Consensus Time	Average time for panel to reach agreement	High
	Strategy Implementation Rate	Percentage of decisions successfully executed	High
	Audit Trail Completeness	Percentage of decision steps with full documentation	Medium

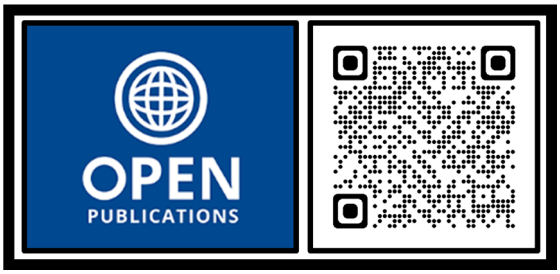

Using Data Science Tools: A Case Study to Identify Weak Signals

www.openpublications.org