vehere

Vehere Vision AI

See the unseen. Stop the unknown.

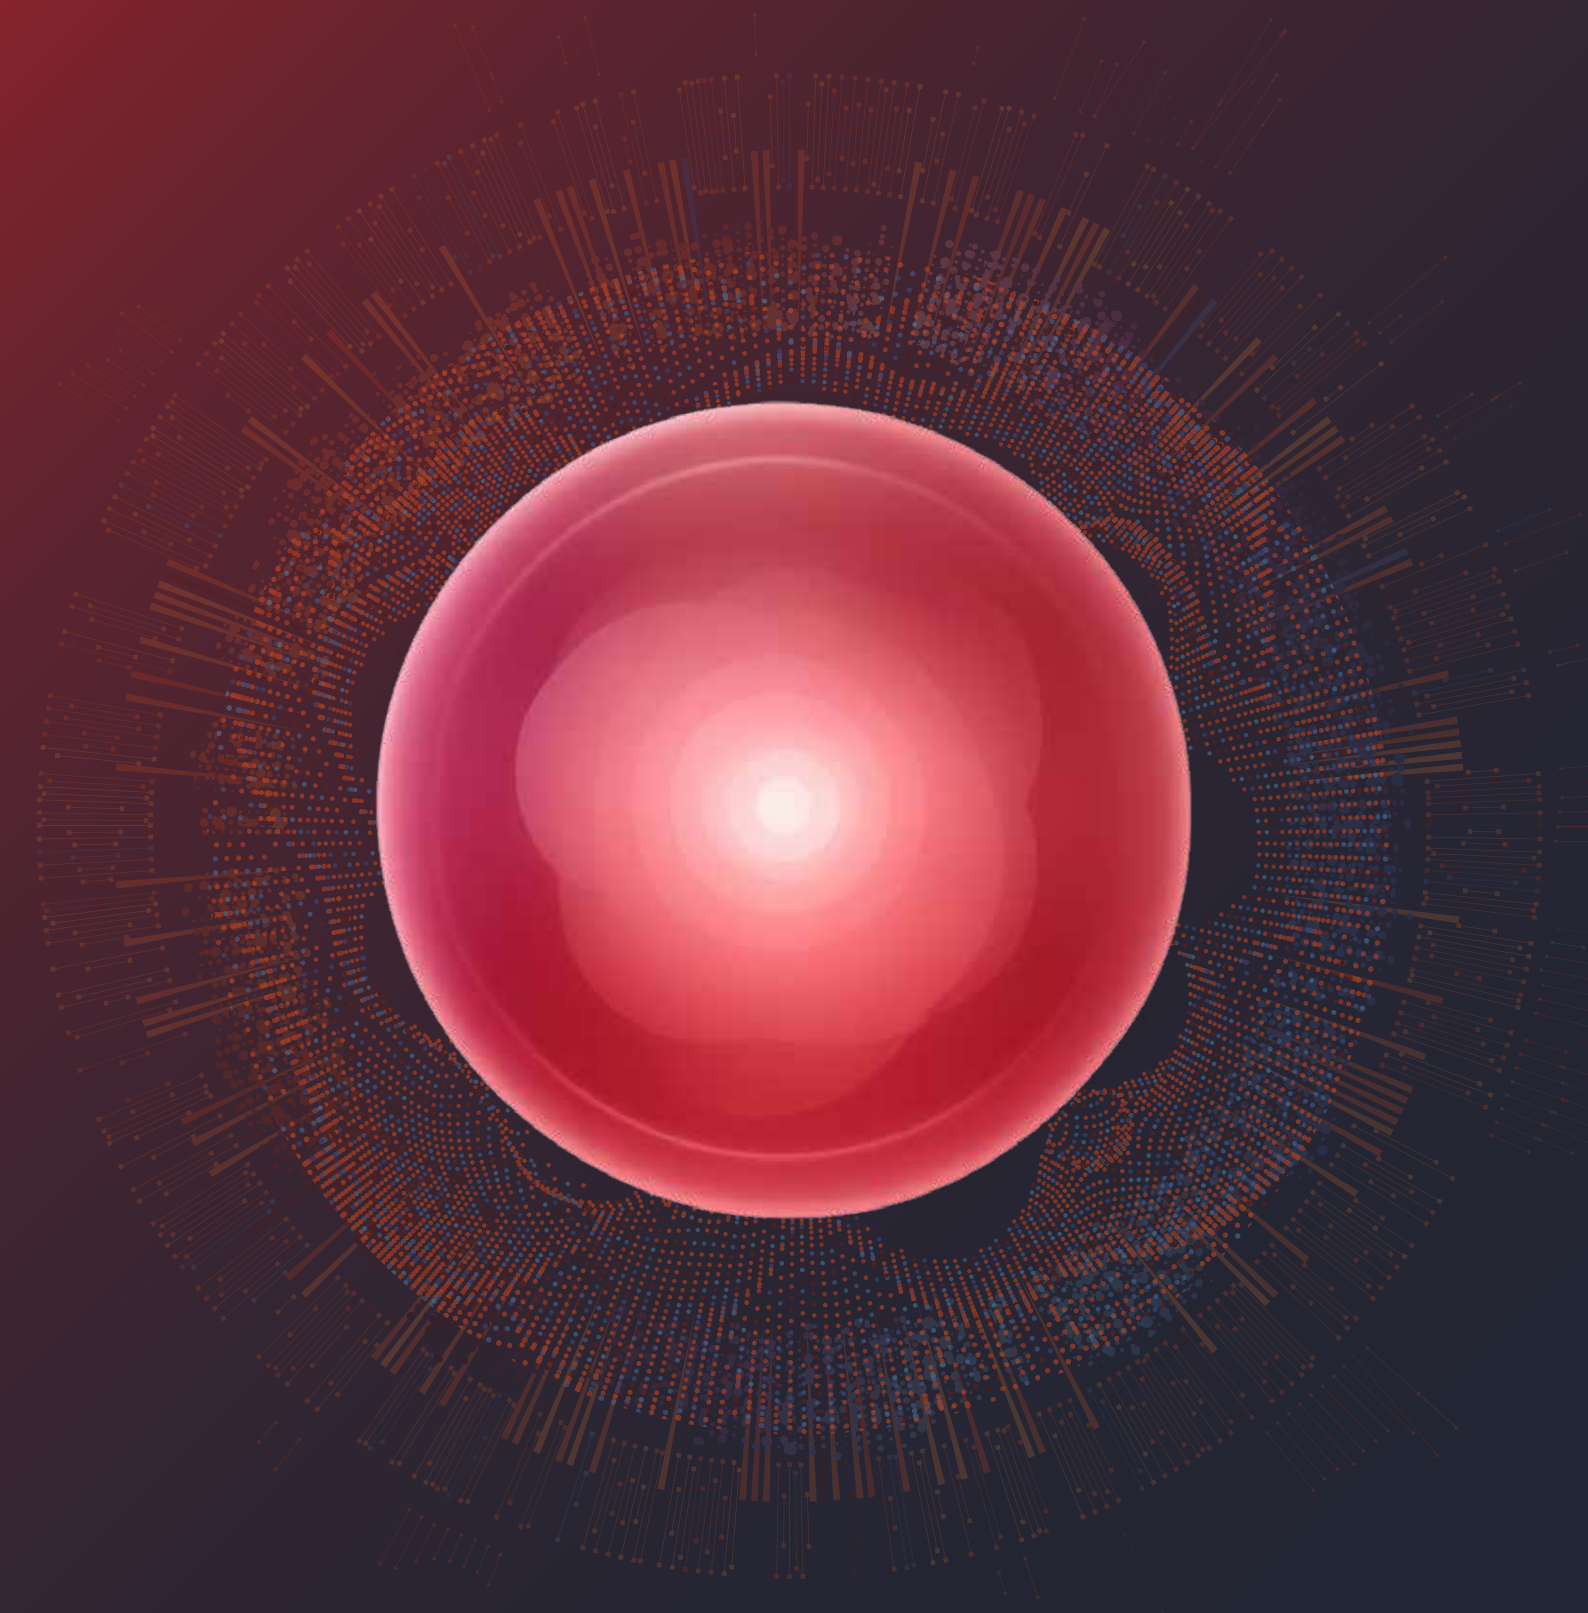

Vehere Vision AI

See the unseen. Stop the unknown.

Detects everything from evasive attacks to subtle insider anomalies defended by a hybrid, multi-layered AI powered defense

Introduction

At Vehere, we are guided by a singular principle- 'Visibility is the foundation of security'. With true visibility comes the power to equip defenders with intelligence that can transform this visibility into decisive, meaningful action. Vehere delivers this intelligence through its powerful AI powered hybrid defense framework, anchored by Vehere Vision AI - the network-aware intelligence layer that unifies large language models, supervised and unsupervised learning, and deep behavioral analytics.

Vehere Vision AI ensures that visibility is not just collected, but understood, contextualized, and translated into clear, actionable insights. Vehere Vision AI powers the full spectrum of Vehere solutions from SIGNT for cyber defense to Lawful interception, Network Detection and Response and Firewall serving as the common intelligence layer across them all.

In this paper, we will explore the core concepts of Artificial Intelligence and Machine Learning, tracing the progression from foundational AI principles to advanced methods such as deep learning and representation learning. We will detail Vehere's hybrid defense framework, outlining the role of each algorithm, both supervised and unsupervised and how they work in concert. We will also discuss Vehere Vision AI, an intelligence layer that underpins every component of the system. Working together, these AI powered components, amplified by Vehere Vision AI, create a cohesive, multi-layered defense. It proactively identifies threats in real time- ranging from highly evasive external attacks to subtle insider anomalies while also taking decisive actions on behalf of the analyst.

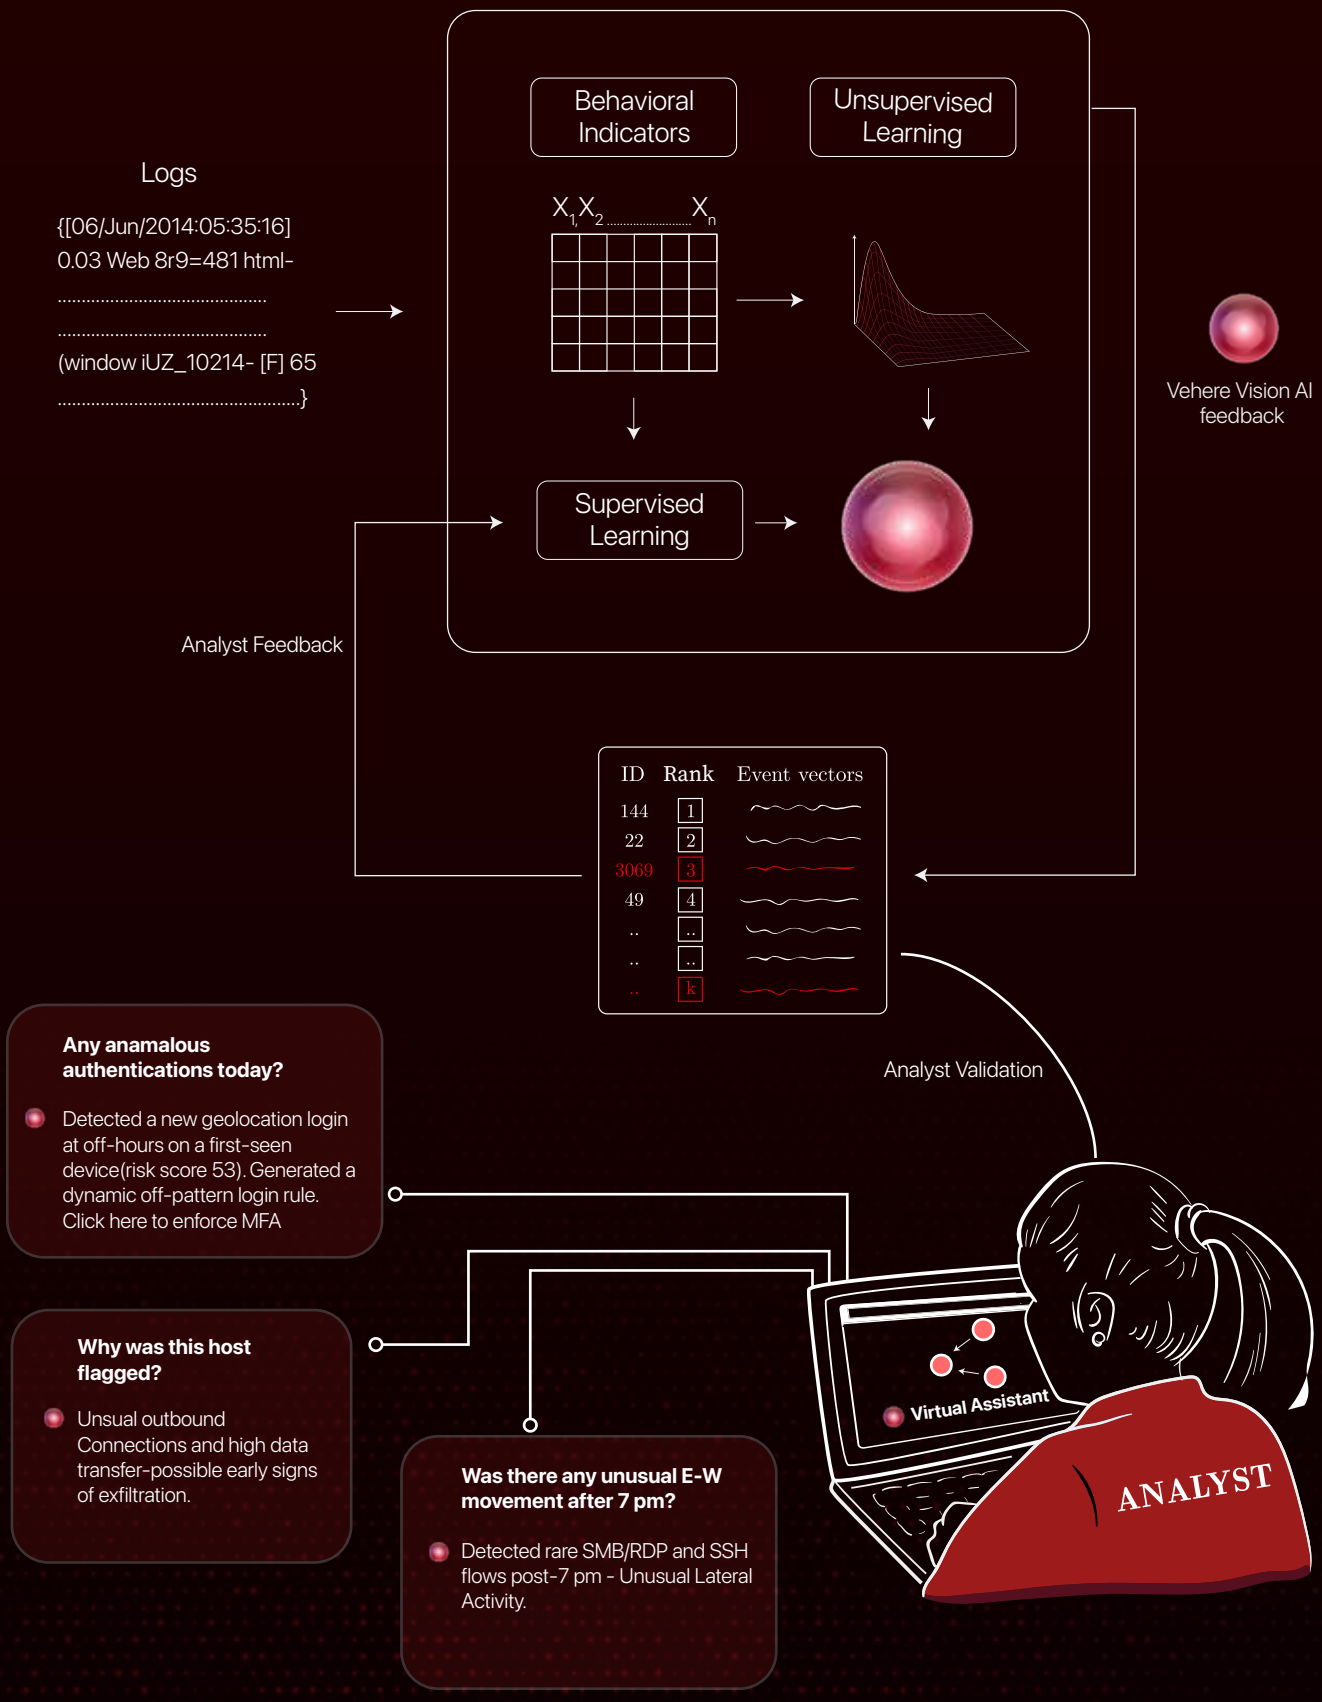

Figure 1.1: The flow illustrates how logs are transformed into features, evaluated by learning models, and surfaced through Vehere Vision AI, which provides context, explanations, and recommended actions.

Sommer and Kalyanaraman, 2016. AI²: Training a big data machine to defend.

About Vehere

From our origins in SIGINT in cyber defense and counter-terrorism, where every signal could mean the difference between prevention and consequence, Vehere was built on one defining principle: Deliver complete network visibility, amplified with purpose-built Artificial Intelligence, and give defenders the strategic advantage in the cyber battlefield. Vehere Vision AI carries this forward as the embedded intelligence layer, turning visibility into foresight and autonomous action.

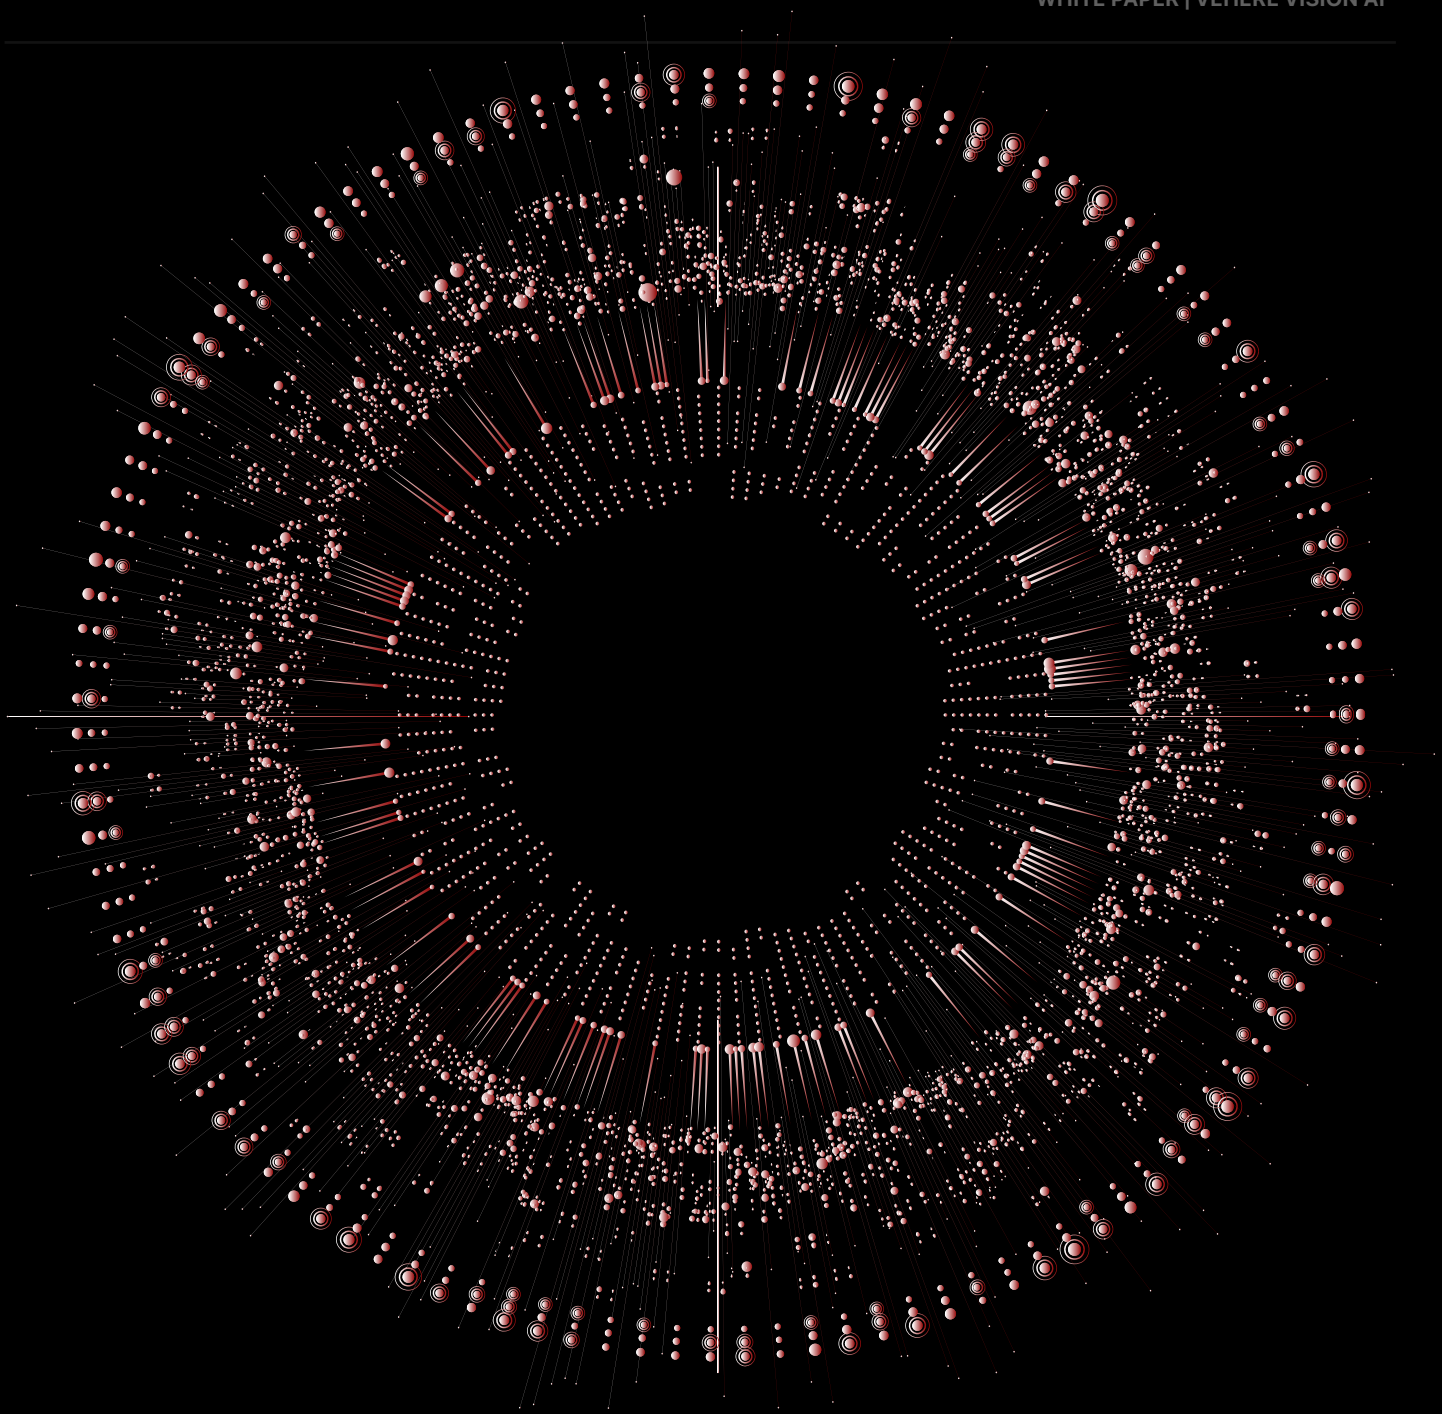

What is AI?

Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) are often used interchangeably with the mistaken assumption that they refer to the same discipline or represent a hierarchy of capability. In reality, while these fields are closely related, each has its own distinct definition, scope and application. Understanding their nuanced distinctions is essential to demystifying the trio and accurately determining their relevance in both technical discourse and practical application.

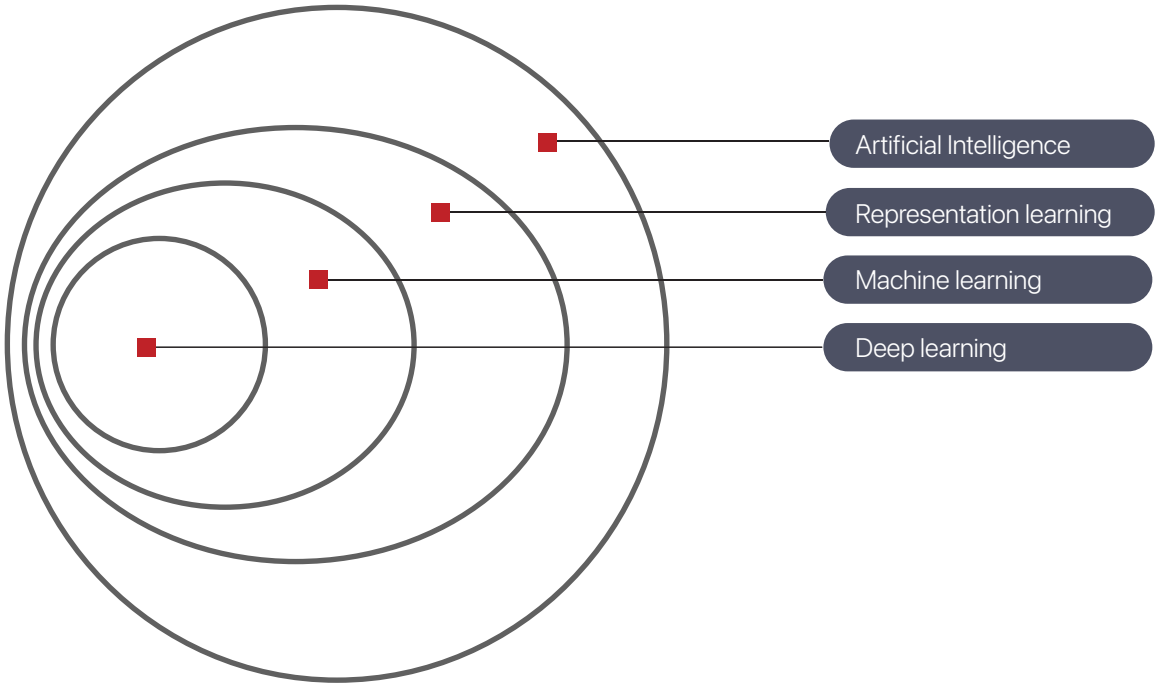

Figure 1.2 Relationship between different AI

Artificial Intelligence:

Artificial Intelligence is a broad discipline of computer science dedicated to creating systems that can perform tasks requiring human-like cognitive abilities, such as reasoning, learning, perception and decision making. Humans perform many of these tasks intuitively like recognizing speech, identifying faces, or interpreting context in conversation because these rely on perception and pattern recognition that are difficult to express in formal computational rules. AI addresses this by learning from examples, adapting to new inputs and continually refining its performance while progressively closing this gap between human intuition and machine capability

Machine Learning:

Machine Learning is a branch of artificial intelligence that enables computers to learn from data and improve their performance without being explicitly programmed for every step. Instead of relying on fixed, human-written rules, ML systems find patterns and relationships in information, then use those insights to make predictions, classify data, or detect unusual events.

In practical terms, a machine learning task starts with an example, a set of measurable characteristics, known as features, describing an object or event we want the system to analyze. By processing many such examples, the system learns how these features relate to outcomes, allowing it to handle complex problems that are too dynamic or nuanced for traditional, hard-coded solutions.

Goodfellow, Bengio and Courville, 2016

Representation Learning:

Representation Learning is a type of machine learning in which the system learns not only how to map inputs to outputs, but also how to automatically discover the best way to represent the input data. Instead of relying on hand-crafted features designed by humans, representation learning algorithms identify and encode the most relevant patterns or "features" in the data, often achieving much better performance and adaptability.

For example, imagine a security system that monitors network traffic. At first, it only sees raw data like IP addresses, timestamps, and packet sizes. Through representation learning, it can transform this raw information into more meaningful patterns, such as "typical employee browsing," "video streaming," or "possible data exfiltration." This new way of representing the data makes it much easier for the system to detect unusual or malicious behavior.

Deep Learning:

Deep Learning is a type of machine learning that teaches computers to understand the world by building concepts in layers. Simpler ideas form the foundation for more complex ones. Instead of relying only on human-designed features, deep learning automatically discovers high-level patterns hidden in raw data, such as recognizing shapes in an image or distinguishing speech accents.
The "deep" in deep learning refers to the multiple layers of processing that transform the input step-by-step into more abstract and useful representations.

For example, in image analysis, the first layer might detect edges, the next might identify corners, the following layers could detect shapes like wheels, and the final layers could recognize cars.

This layered approach allows deep learning systems to handle the complexity and variability of real-world data such as different lighting, angles, or background noise while discarding irrelevant details. As a result, deep learning has become a powerful method for tackling problems, where identifying and disentangling important factors from raw data is extremely challenging.

Refer image 1.3 contrasting machine learning and deep learning. Where Machine Learning separates feature extraction and classification, giving limited outputs (e.g., "car" or "not a car"), Deep Learning integrates both, automatically learning complex patterns to deliver richer insights (e.g., car make, model, and color).

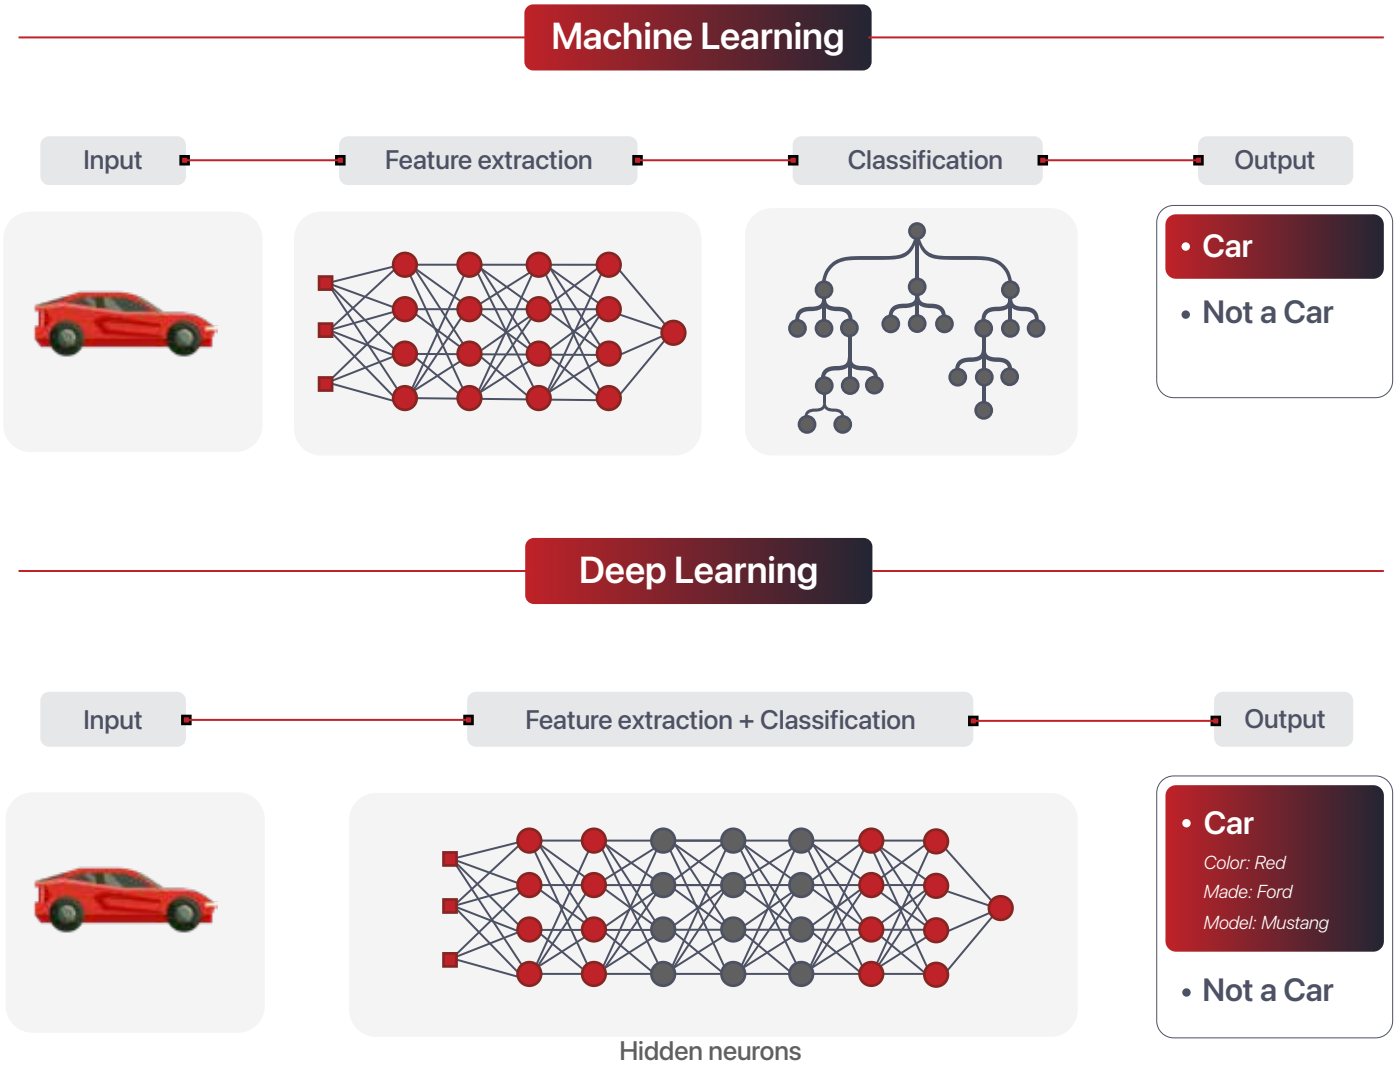

1.3 Contrasting machine learning and deep learning with an example of a car

Types of **Algorithm** Learning Techniques:

Unsupervised

Machine learning algorithms can be broadly categorized as unsupervised or supervised by what kind of experience they are allowed to have during the learning process.

One of the oldest datasets studied by statisticians and machine learning researchers is the Iris dataset-Fisher, 1936. It is a collection of measurements of different parts of 150 iris plants. Each individual plant corresponds to one example. The features within each example are the measurements of each of the parts of the plant: the sepal length, sepal width, petal length and petal width. The dataset also records which species each plant belonged to. Three different species are represented in the dataset.

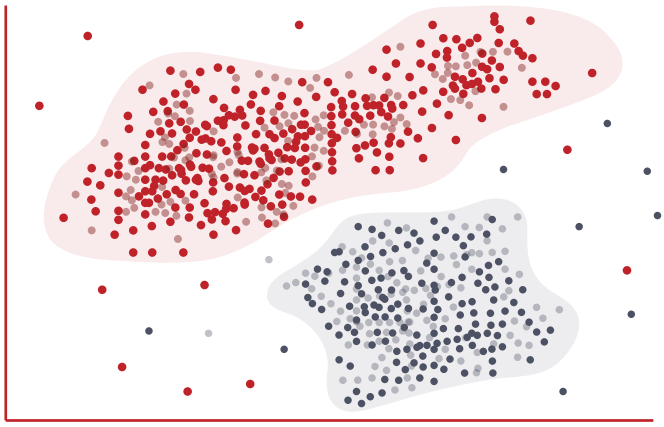

1.4: Depiction of the unsupervised clustering method where labelled data trains a model to predict outcomes

Unsupervised learning algorithms experience a dataset containing many features, then learn useful properties of the structure of this dataset. Some other unsupervised learning algorithms perform other roles, like clustering, which consists of dividing the dataset into clusters of similar properties.

For example, in the Iris dataset, an unsupervised algorithm would use only the iris measurements to group irises into clusters based on similarities in features, without any prior knowledge of their species. (refer image 1.4)

Colab, n.d. "The Use of multiple measurements in Taxonomic Problems", Ronald Fisher (1936)

Supervised

Supervised learning algorithms are trained on datasets that include both features and corresponding labels. The algorithm learns the relationship between inputs and outputs to make predictions on new data.

For example, the Iris dataset includes measurements of each flower along with its species label. A supervised learning algorithm can use these labelled examples to learn how to classify irises into different species based on their measurements.
(refer image 1.5)

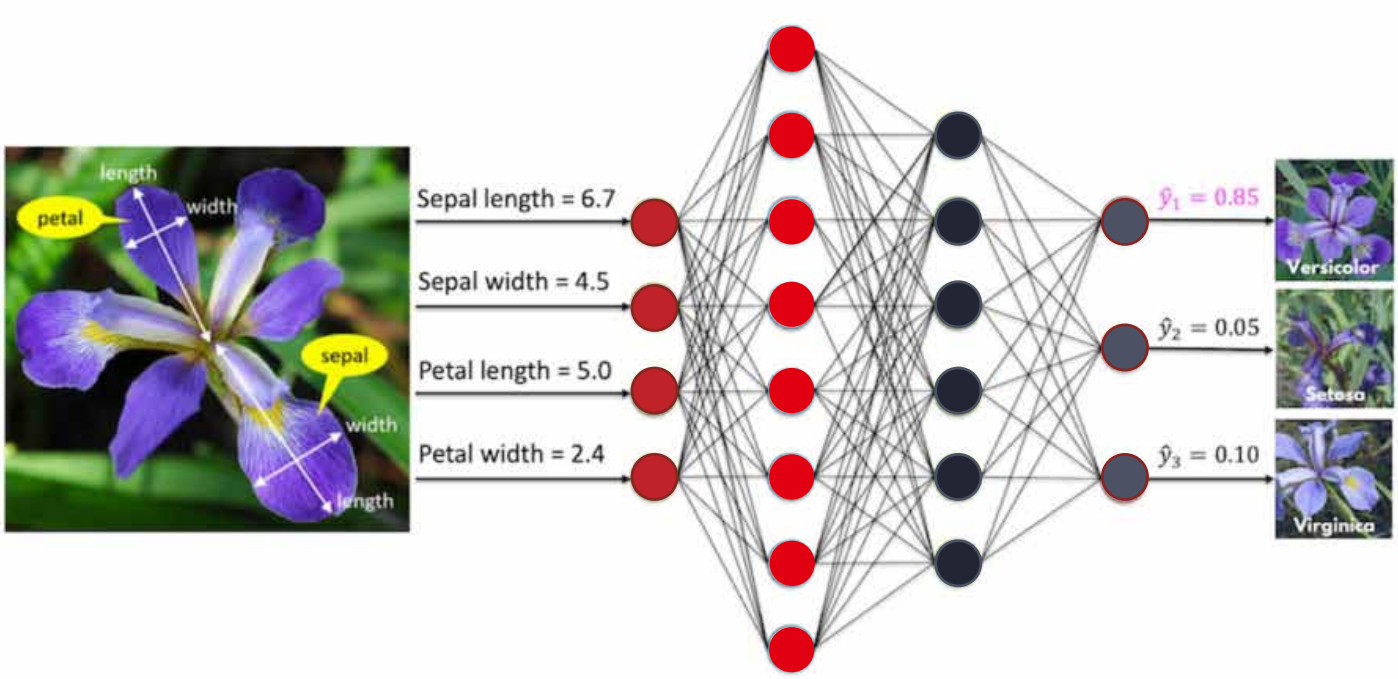

1.5: Depiction of a supervised algorithm, where labelled data
(Iris measurements paired with species) trains a model to predict outcomes

Within these broad approaches, are a range of different learning algorithms, with researchers continuing to create new ones. To complicate things further, algorithms can be combined to form even more complex systems. The question then arises, how does a data scientist choose the right algorithm or algorithms to solve a particular problem? Or can one algorithm be superior to all others no matter what the problem is?

Colab, n.d. "The Use of multiple measurements in Taxonomic Problems", Ronald Fisher (1936)

Unsupervised Learning

- K-Means Clustering
- Apriori Algorithm

Deep Learning Models

- Long Short-Term Memory Net- works (LSTM)

Semi-Supervised Learning

- Artificial Neural Networks (ANN)

Supervised Learning Algorithms

- Naïve Bayes Classifier Algorithm
- Support Vector Machine (SVM) Algorithm
- Linear Regression Algorithm
- Decision Trees Algorithm
- Random Forests Algorithm
- K-Nearest Neighbours Algorithm

1.6: Illustrative overview of widely used Machine Learning algorithms that form the foundation of modern AI applications

No Silver Bullet Principle

The No Silver Bullet Principle states that there is no single algorithm, methodology, or technology that can address every problem with equal effectiveness.

> Not only are there no silver bullets now in view, the very nature of software makes it unlikely that there will be any"
>
> — Frederick P. Brooks, Jr.

In cybersecurity, threats vary widely in nature, from insider anomalies and malware detection to phishing campaigns and ransomwares. The characteristics of these problems such as data type, volume, velocity, and labeling availability mean that a single model or approach cannot perform optimally across the board. Below are a few variations which show that "AI" is not a single approach, but a spectrum of methods, each suited to different environments and operational demands.

- **Domain specificity:** An algorithm fine-tuned for network anomaly detection may fail when applied to spear-phishing detection, because the underlying data patterns and features are fundamentally different.

- **Data variance:** Some problems have rich labeled datasets, ideal for supervised learning; others lack labels and require unsupervised or semi-supervised techniques.

- **Operational constraints:** In high-throughput environments, even the most accurate algorithm may be unusable if it cannot process data in real time.

How does the data scientist choose the most optimum algorithm to use?

Every solution comes with trade-offs, and its success depends heavily on the specific nature of the problem, the available data, and the operational constraints.

Occam's Razor Principle:

Occam's Razor is a problem-solving principle that states: when faced with multiple explanations or solutions, choose the simplest one that adequately fits the data. Simplicity here does not mean "oversimplification". Rather, it means avoiding unnecessary assumptions, complexity, or overengineering when a more straightforward approach is equally effective.

For example, if both a 50-layer deep neural network and a well-tuned decision tree achieve the same detection accuracy for phishing emails, Occam's Razor favors the decision tree for its simplicity, interpretability, and lower maintenance overhead.

Finding the best fit: How data scientists select optimum algorithms:

Data scientists typically follow a structured, iterative process to choose the right model for a given problem. The choice is rarely about "picking the most advanced algorithm". It's about finding the best fit for the problem, data, and operational requirements.

Brooks, F.P. Jr. (1987) 'No Silver Bullet: Essence and accidents of software engineering'

In light of the above, data scientists follow several best practices when choosing and applying ML algorithms for NDR

- **Data driven algorithm selection**: The nature of network telemetry dictates the choice of algorithm. Structured data such as flows or aggregated logs often align with tree-based models or clustering techniques. Sequential data like logs benefits from LSTM or other time-series models, while unstructured and high-dimensional inputs (e.g., DNS queries, packet payloads) call for deep learning to automatically derive features.

- **Aligning models with detection goals:** Algorithms are mapped to the threat behavior being addressed. Well-defined threats with abundant examples are handled using supervised learning for precision. Emerging or unknown threats lean on unsupervised anomaly detection. Sequential threats such as lateral movement are best modeled with sequence-aware algorithms while zero-days or insider threats are better detected through anomaly detectors and autoencoders.

- **Balancing Accuracy and False Positives:** Data scientists optimize for both high detection rates and low false alarms. Each algorithm's inherent bias (sensitivity vs. specificity) is carefully weighed. For example, supervised classifiers are tuned for precision against known threats but complemented with anomaly detectors to catch novel attacks.

- **Interpretability for Analyst confidence:** In regulated or mission-critical environments, clarity is as important as accuracy. Interpretable models such as decision trees or rule-based approaches are leveraged to provide transparent reasoning behind detections. Deep learning models need to be enhanced with explainability frameworks to ensure that the SOC analysts understand the "why" behind alerts.

- **Scalability and Operational Fit:** Algorithms are selected with real world infrastructure in mind. The data pipeline is optimized to pre-process and reduce data volumes, ensuring models run at scale. Many organizations use a modular design to update individual models as threats evolve without disrupting operations.

- **Continuous Adaptation:** Threat landscapes evolve rapidly. Performance is continuously benchmarked against new attack techniques. Threat intelligence (new IoCs, tactics) and real-world incident feedback feed into retraining cycles.

The most effective approach is to adopt a hybrid strategy combining broad anomaly detectors that capture unforeseen threats with targeted detectors that identify known malicious behaviors with high precision

Inspired by Goodfellow et al., 2016

AI in threat detection:

Without AI, threat detection largely depends on static, rule-based mechanisms that identify threats only when they match known, predefined patterns. With AI, detection moves beyond fixed rules leveraging models that can learn normal network behavior, identify deviations, and flag potential zero-day threats.

AI's ability to model "normal" makes it possible to detect unknown or emerging threats that traditional systems would miss. If the goal is to uncover threats outside the boundaries of prior knowledge, AI is essential.

Modern Network Detection and Response (NDR) solutions leverage multiple AI paradigms to address the diversity and sophistication of today's cyber threats. Hence, most effective NDR solutions combine several methodologies into a unified, hybrid defense framework.

How Vehere Vision AI thinks, explains and defends?

Much like a prism focuses scattered light into clarity, Vehere Vision integrates multiple AI subsystems into one concise, operationally effective result. Beyond clarity, it enriches visibility with context, transforms detections into dynamic rules, and acts decisively on behalf of the analyst. This layered approach enables precise identification of known threats while simultaneously uncovering novel threats and zero-day attacks.

Inspired by Goodfellow et al., 2016, Deep learning. Cambridge, MA: MIT Press Rong, J. (2024) 'Applications of machine learning in cybersecurity'

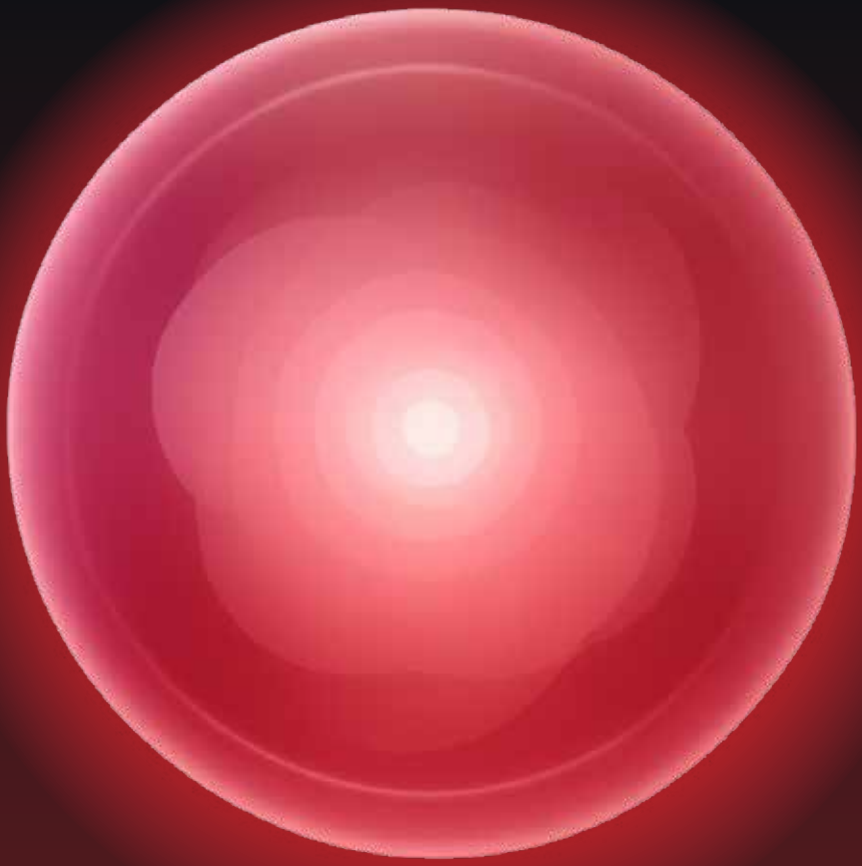

V E H E R E

VISION AI

SEE THE UNSEEN. STOP THE UNKNOWN.

Vehere Vision AI underpins all Vehere solutions, including SIGINT for cyber defense, Lawful interception, Network Detection and Response and Firewall

Multiple Models. One Mission : Eliminate blindspots

Vehere Vision AI is the orchestrator of multiple AI algorithms, rather than just a single predictive model

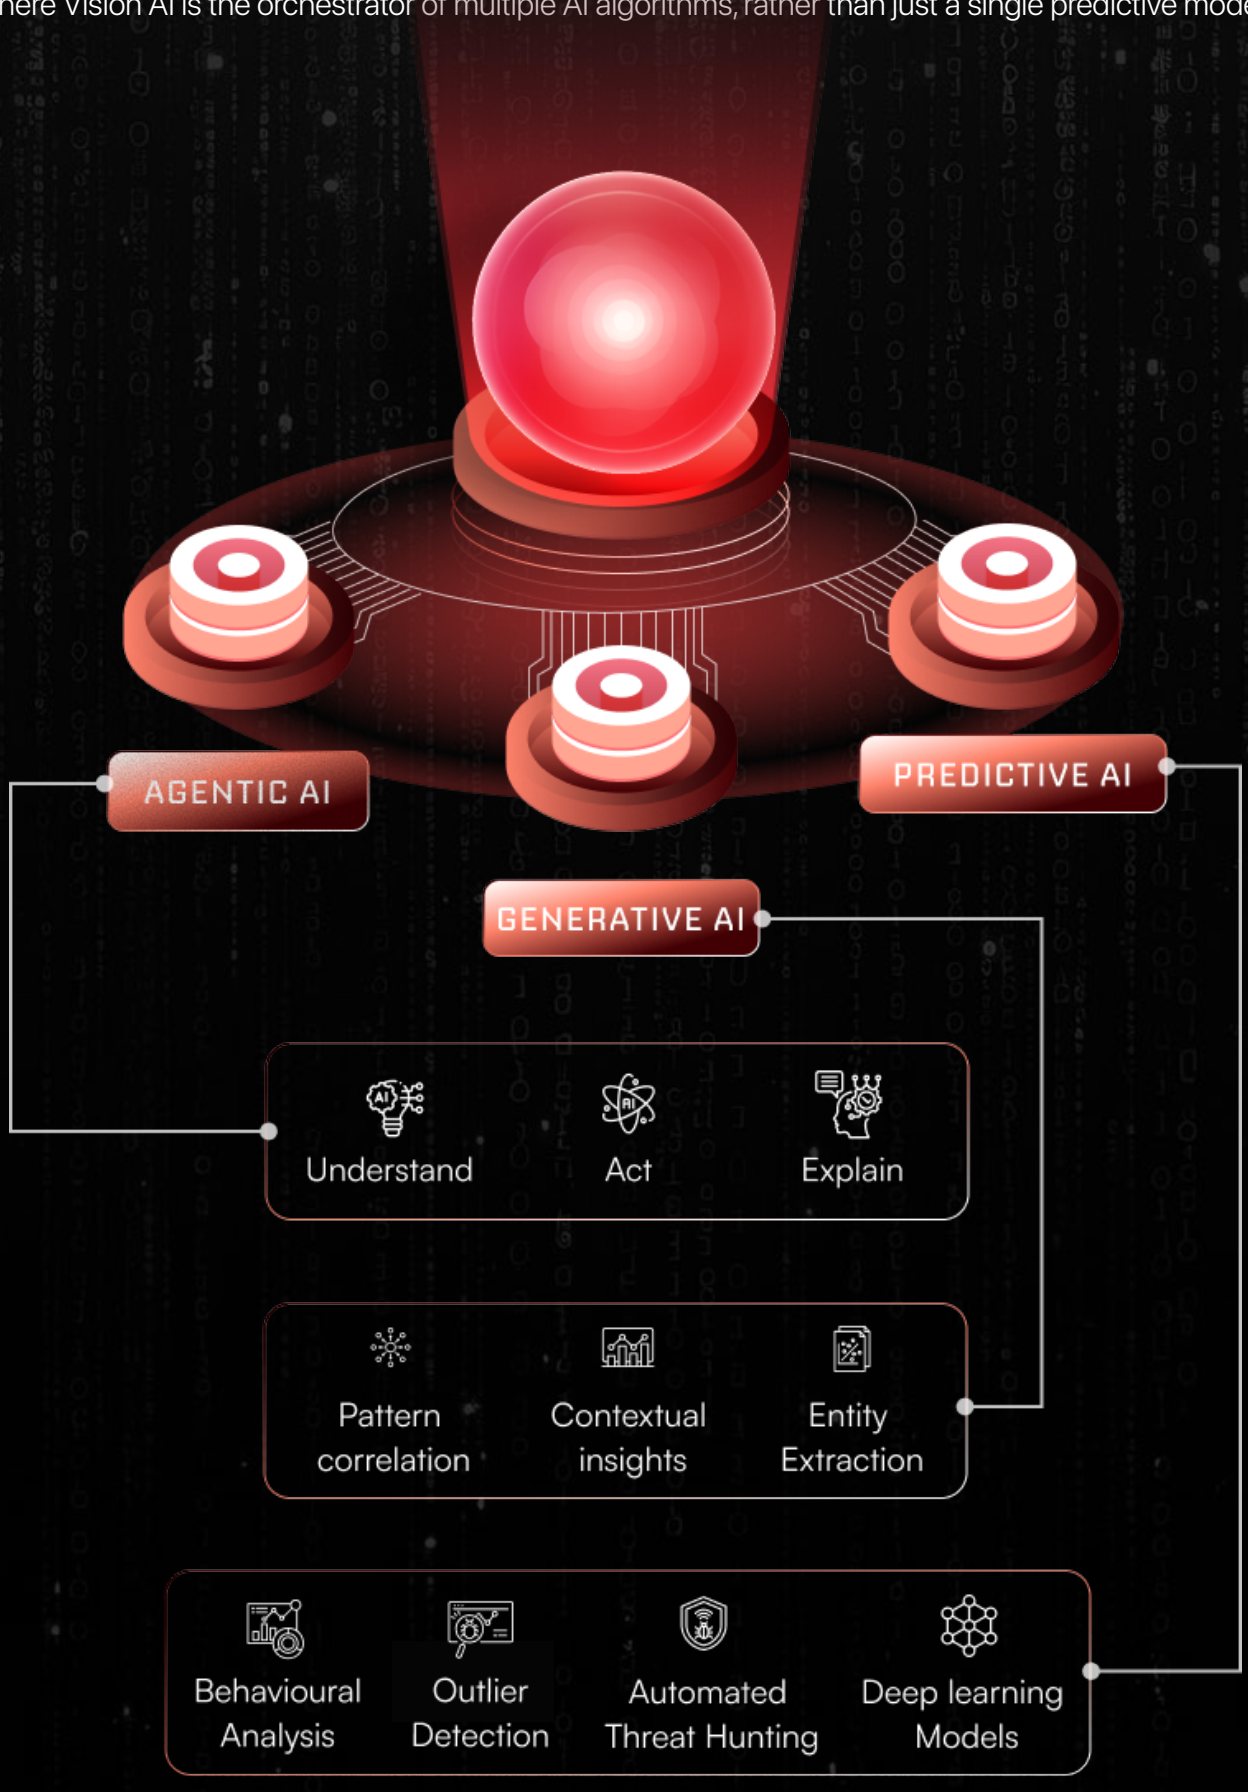

- **Behavioral Analytics:** Continuously analyzes billions of flows, sessions, and logs to surface subtle deviations in user, entity, and device behavior.

- **Outlier Detection:** Leverages unsupervised anomaly detection and grouping based on common topics. For instance, in a collection of cybersecurity incident reports, terms like "exfiltration," and "command and control" could be part of a 'Malware operations topic'. Words that don't align with any of the topics are considered 'Outliers'.

- **Analyst Feedback Integration:** Just as topic modeling uses patterns in words to group them into topics, analyst feedback provides patterns of judgment, which are essentially the decisions analysts repeatedly make when classifying events. By capturing these, the model continuously refines itself.

- **Supervised Learning Pipelines:** Unlike topic modelling which organizes related words in a cluster, supervised learning uses labeled examples and historical context to build precise categories to determine what is malicious.

Vehere Vision AI fuses **Large Language Models (LLM)** with real-time and historical data to deliver precise, explainable, evidence-backed answers to analyst questions such as "Why was this host flagged?" or "Show all east-west traffic initiated by user Jane.Doe after 11 pm". All outputs are **role-aware, explainable, and auditable**, keeping analysts in control while accelerating decisions.

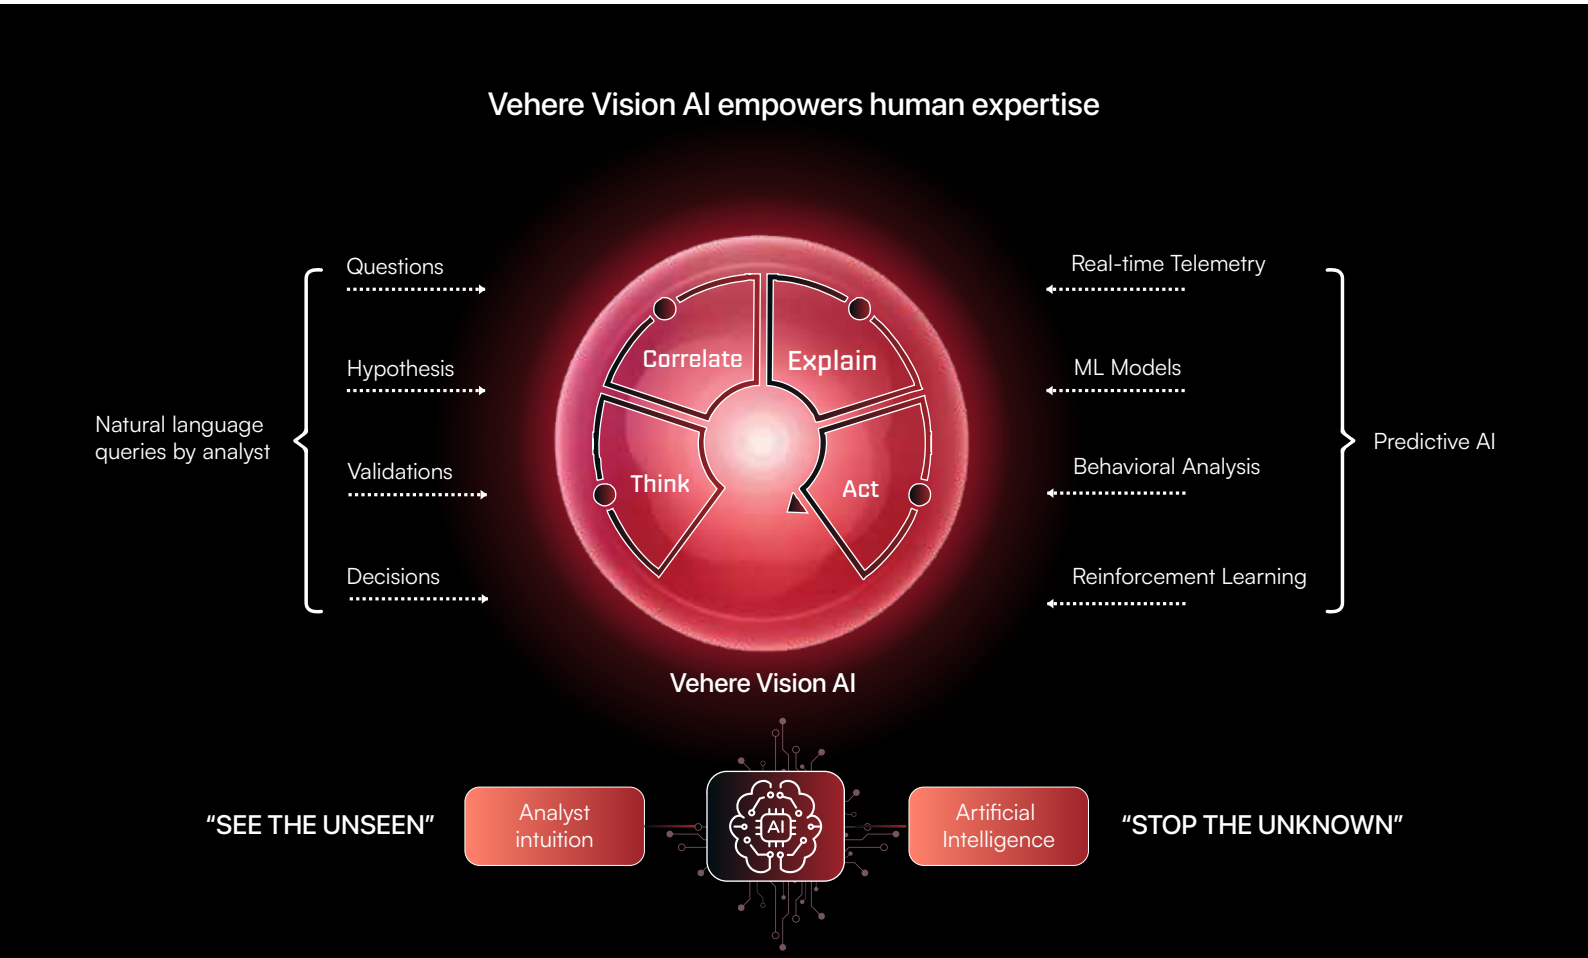

Figure 1.8: The diagram illustrates how Vehere Vision AI integrates analyst intuition and predictive AI. Analysts interact with the system through natural language queries while Vehere Vision AI draws on machine learning models. By thinking, correlating, explaining, and acting,Vehere Vision AI bridges human expertise with AI powered insights to support both understanding and action

Vehere Vision AI: Your force multiplier

Vehere Vision AI embeds intelligence across the entire workflow

1) Natural Language Intelligence:
Enables analysts to query in plain language

Example: An analyst types: "Why was this host flagged yesterday?"

Vehere Vision AI: Responds with a concise summary: "The host initiated 43 failed RDP logins to multiple servers within 10 minutes, deviating from its baseline behavior."

2) Rule Automation and Reinforcement Learning:

- Automatically generates contextual detection rules from observed patterns.
- Refines its detection logic with analyst feedback.

Example: Repeated suspicious SMB connections observed across several entities.

Vehere Vision AI: Automatically drafts a detection rule: "Flag outbound SMB connections from finance subnet to any external IP." Analysts can validate and deploy in clicks.

3) Advanced Protocol and Entity Understanding:

- Decodes complex network protocols in real time, surfacing anomalies
- Maps relationships between users, devices, and applications to reveal hidden attack paths.

Example: During analysis of encrypted traffic, Vision AI detects anomalous DNS tunneling patterns.

Vehere Vision AI: Decodes the protocol usage and signals: "Outbound DNS traffic contains unusually long TXT records consistent with data exfiltration attempts."

4) Correlation and Threat Storytelling

- Connects weak, disparate signals (from network flows, IOCs, identity events) into coherent narratives.
- Vehere Vision AI highlights lateral movement, insider risks, and anomalies within peer groups.

Example: A weak identity alert, a strange flow pattern, and a flagged IOC appear separately in logs.

Vehere Vision AI: Stitches them into a narrative: "User J.Doe authenticated from a new geolocation, accessed 500 MB of files over SMB, and connected to an IP associated with known C2 infrastructure. Hence, it is likely lateral movement with exfiltration."

5) Response Acceleration and Strategy Refinement

Recommends next steps, playbooks, and mitigation strategies aligned to frameworks like MITRE ATT&CK. Generates incident timelines and executive-ready summaries, cutting investigation and reporting time.

Example: A ransomware event is detected.

Vehere Vision AI: Recommends "Quarantine host HR-PC-21, disable user account j.doe, and block SMB traffic to 10.10.45.0/24. Similar tactics observed in MITRE ATT&CK T1486 (Data Encrypted for Impact)".

6) Continuous Adaptation and Learning

Employs reinforcement learning and feedback loops to refine detection models with evolving adversary tactics.

Example: Analysts mark a recurring detection as benign (legitimate backup traffic).

Vehere Vision AI: Learns from the feedback, refines its model, and stops surfacing those false positives; reducing noise over time.

7) Analyst Empowerment at Scale

Transforms the way analysts work by embedding intelligence directly into their workflows. Through natural language interactions, analysts can query the system in plain English and receive evidence-backed answers.

Example: A junior analyst queries Vision AI: "Show me unusual file access by interns this week."

Vehere Vision AI: "Intern group accessed HR payroll files for the first time on Sept 11 which is abnormal for this peer group." The junior analyst escalates confidently with evidence in hand.

Mapping Vehere solutions' capabilities to the AI models that power them

The following section illustrates how Vehere solutions' core capabilities align with the AI algorithms:

Analysis of high-volume flow and session data in real time: Hybrid Deep Neural Network

At the heart of Vehere NDR is a hybrid Deep Neural Network (DNN) architecture that combines Apache Spark in memory data processing and supervised DNN for network anomalies detection. This dual approach enables the system to process high-volume flow and session data in real time, uncovering sophisticated attack patterns that span both immediate and long-term behaviors.

DGA and malicious Domain Detection: Tree Based Algorithm

Domain Generation Algorithms (DGAs) are used by Advanced Persistent Threats (APTs) to create disposable, malicious domains for Command-and-Control (C2) communications. Vehere's tree based algorithm detects DGAs by learning linguistic and statistical patterns to accurately flag suspicious domains in live DNS traffic.

DNS Anomaly Detection: Unsupervised clustering algorithm

Vehere leverages an unsupervised clustering algorithm, a probabilistic topic-modeling technique, to establish thematic patterns of normal DNS resolution behavior. By grouping domain activity into "topics," it identifies deviations that may indicate stealthy malware beaconing or covert domain tunneling without relying on predefined threat signatures.

Vehere applies the unsupervised algorithm to network telemetry to learn latent structures of threat behaviors. This enables the platform to surface subtle, evolving attack patterns that would remain invisible to manual inspection.

Adapted from Liu, J., Wang, H. and Zhang, Y. (2025), 'A deep dive into latent Dirichlet allocation (LDA) and its applications on recommender systems

Insider Threat detection: Entity Behavior Analysis

The EBA engine builds a 7-14 day behavioral baseline for each network entity monitoring communication pattern, resource access patterns and time-based activities. It flags anomalies such as unauthorized access, or compromised endpoint behavior, reducing false positives through contextual correlation.

Data hoarding and exfiltration detection: Tree based algorithms

To detect large, anomalous data transfers; often precursors to data exfiltration, Vehere applies tree-based unsupervised algorithms that isolate abnormal patterns by comparing real-time transfer volumes against historical baselines for each entity.

Explainable AI (XAI):

Vehere integrates explainability frameworks that allow SOC analysts to see why an alert was triggered. For example, highlighting that a spike in outbound encrypted traffic at atypical hours was anomalous compared to baseline peer behavior.

Instead of a cryptic probability score, Vehere provides analysts with clear, contextual reasoning mapped to MITRE framework.

Language Identification (LID): Transformers and Deep learning models

The LID module detects the language of audio sessions, transcribes the content, and translates it into English. It leverages linguistic pattern recognition and neural machine translation to assist analysts in multilingual investigations.

Speaker Identification (SID): Voice embedding models:

SID uses deep neural embeddings and supervised classifiers to match recorded voices with pre-trained samples, assigning confidence scores to each match. This enables precise identification of individuals in recorded communications.

VISIBILITY	DETECTION	RESPONSE
Device Clustering	Flow and Session Analysis	Root Cause Analysis
Critical Asset Prioritization	Domain Threat Detection	Detector-specific Root cause analysis
Behaviour Profiling, High Risk User identification	Entity and User Behaviour Analytics	Agentic correlation
Traffic Pattern Identification	Dynamic File Analysis	Intelligent Prioritization
Peer Group Modelling	File & Exfiltration Analysis	Autonomous responses

Accelerating analyst efficiency with the power of Vehere Vision AI

In an ideal world, every Security Operations Center (SOC) would have time and resources to analyze every alert in depth. The reality, however, is that most teams are not only stretched thin, but often staffed with varying levels of experience, tasked with doing more in less time.

Vehere Vision AI accelerates workflows by augmenting less-experienced analysts, enabling faster detections and accurate assessment of event severity, so teams can focus on responding rather than drowning in triage.

Vehere Vision AI	Embedded intelligence layer that unifies generative, predictive and agentic AI to turn raw signals into contextual insights and autonomous actions; working as a co-analyst for the SOC
Threat Severity and Risk Scoring	Threat severity is an indicator of the urgency of a threat and its impact. Risk scores are assigned to anomalies, helping prioritize high-risk incidents and reducing alert fatigue
MITRE ATT&CK Mapping	Highlights which MITRE ATT&CK tactic, technique or procedure a detection relates to
Customized Reporting	Personalized graphical and statistical representations by selecting relevant data fields, applying filters, and organizing information to align with their unique requirements
Autonomous Responses	Trigger autonomous, policy-based actions through Vehere Vision AI
Agentic Correlation	Enriches raw alerts with session data and live threat intelligence automatically
Language ID and Speaker ID	LID automatically recognizes the spoken language or dialect to enable accurate processing and translation. SID detects and distinguishes individual speakers using voice patterns
Mass profiling	Analyzes bulk or target data, builds profiles, extracts and maps key entities and their relationships
Behavior Profiling	Learns normal user and device patterns to detect deviations such as access from unusual IPs, data exfiltration to unusual destinations, or logins at atypical hours.
Forensic Analysis	Detailed evidence with session reconstruction to support post-incident investigations and root cause analysis

> “We are driven by a mission to defend enterprises and nations by eliminating blind spots and ensuring a safer digital environment. We continue to harness hybrid AI to pioneer advanced detection capabilities that disrupt adversaries and prevent them from achieving their objectives”

For more information, please contact us at **sales@vehere.com**