SymphonyAI
SENSA

WHITE PAPER

SymphonyAI Sensa: Discovering the whole truth

Introduction

SymphonyAI Sensa's technology is a true AI-powered discovery platform for uncovering criminal behaviors and risks hidden in vast amounts of customer, transaction, correspondent and digital finance data. Our hybrid approach of cutting-edge, unsupervised and supervised machine learning, combined with traditional rules and analytical flexibility, energizes a new generation of fully explainable financial crime detection and decisioning.

This paper will walk through the philosophy behind Sensa, how it's built and what makes it the best financial crime discovery engine in the industry.

Our philosophy

Transparent, nimble, explainable, scaled, effective

Risk and compliance teams must meet the requirements of regulators while protecting their enterprise from attacks and abuse. We built Sensa with these users in mind. It is tailor-made to deeply understand transaction, customer, wallet and DeFi data to clearly identify behaviors that potentially threaten the business…all while maintaining privacy, auditability and transparency of a highly regulated industry.

The challenges to effectively balance and solve

Realizing effectiveness in financial crime fighting is complicated and multi-dimensional. There are many unique challenges to this line of business, including:

- **Balance of ROI:** Three dimensions must be optimized – the cost of operations, the failure to be effective with regulators and the repercussions of hindering a highly customer-focused competitive enterprise.
- **Long-range behaviors:** Financial criminals' behavior are opaque, deep, exploit weaknesses in the organizational culture and change gradually over time to avoid detection.
- **Data imbalance:** The vast majority of customers, transactions and wallets are not suspicious;

only a tiny minority show behavior worthy of investigation. Finding this "needle in a stack of needles" is a huge analytical challenge.

- **Behavior opacity:** Behavior and interaction networks have complex properties making them difficult to understand – a weakness effectively exploited by criminal endeavors.
- **Expert knowledge:** There is a vast amount of expert knowledge needed but rarely industrialized in the investigation process; it also is challenging to incorporate this into a machine learning system.
- **Conflicting assessment:** Different investigators may come to different outcomes using the same data. This can be due to their skill levels, workload, training, bias, etc. and can make any signal in the data very noisy**.**

We focus on creative and pragmatic ways to overcome these issues to create explainable transparency that fits seamlessly with existing workflows, processes and systems.

Sensa component workflow

Sensa supports the different components and functionality with a financial crime detection focus.

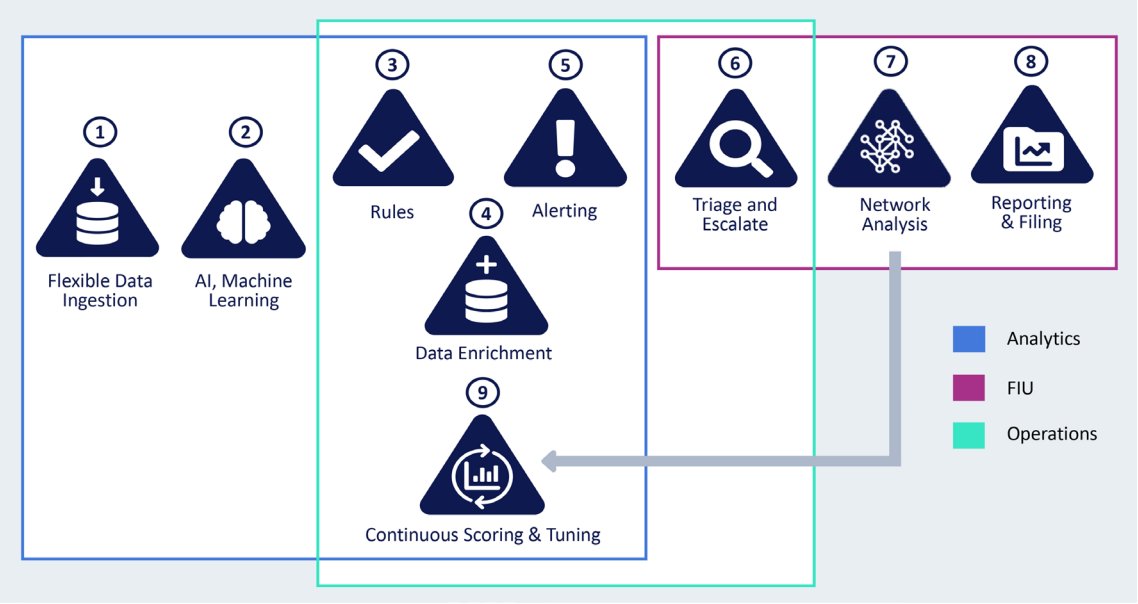

Figure 1: Sensa Functional Flow

Figure 1 shows how Sensa provides value throughout the financial crime prevention lifecycle. There are nine key aspects:

1. **Flexible data ingestion**
No new data - maximize the "informational yield" from existing data without the burdensome data integration and homogenization projects of the past.

2. **AI and machine learning**
Comprehensive analytics capability to allow Sensa and customer models to be built, integrated and deployed in real time.

3. **Rules**
Support existing rules and typologies; quickly create and test new rules via an intuitive UI.

4. **Data enrichment**
Optimize third-party data usage and speed up investigations with additional insights.

5. **Alerting**
Control how models and rules are used to alert across your entire enterprise.

6. **Triage and escalate**
Streamline alerts using risk-based scoring, prioritization, queues and customizable workflows.

7. **Network analysis**
Network and relationship exploration are put at the center of a modern and intuitive investigation process that provides a holistic risk view of your customers.

8. **Reporting and filing**
Built-in KPI reporting and business-focused dashboards with easy integration with downstream systems including case filings.

9. **Continuous scoring and tuning**
Proactive and ongoing evolution of risk scoring and auto-tuning strategy recommendations.

Fit for purpose detection: maximizing the right approach for the right problem

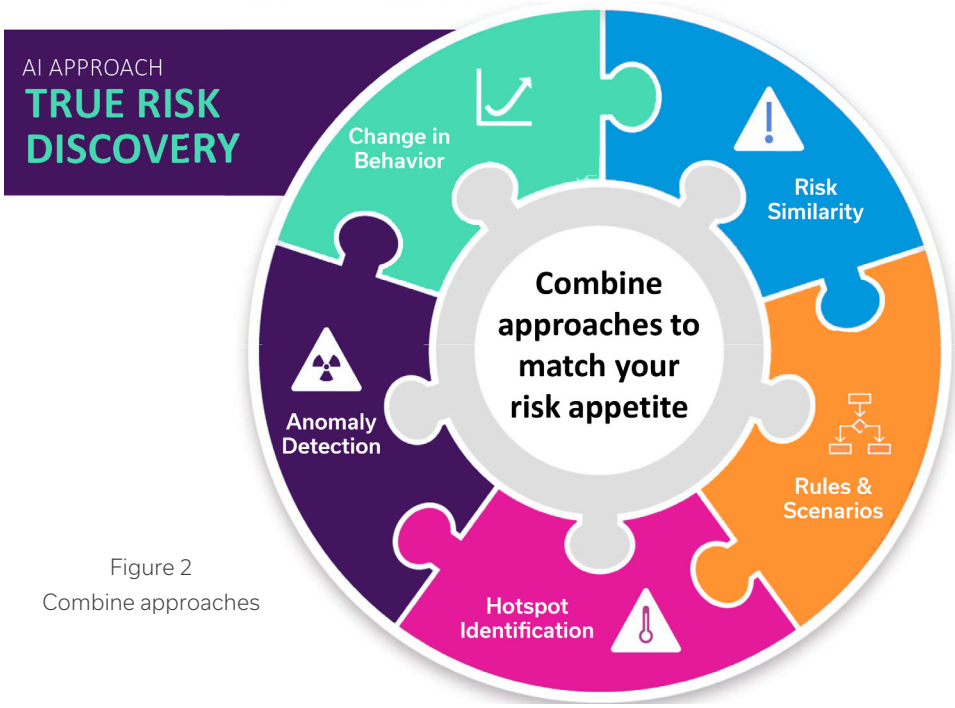

Figure 2
Combine approaches

A combined, hybrid approach to machine learning

While cutting-edge machine-learning methods, data models and scalable deployment platforms are great, they aren't worth much without models that serve an actual business purpose. We recognize that an effective AML strategy requires more than one type of model detecting one specific type of behavior. Instead, we have built a suite of models spanning the entire AML space. Each model can provide risk scores for every entity, allowing for fast and effective investigations. These models include:

- **Risk Similarity Model:** This model compares the behavior of an entity with the behavior of entities that have been alerted and escalated in the past. Compared to a rules-based approach on its own, the Risk Similarity Model has been used to reduce the false-positive rate of AML alerts significantly. We can achieve this by combining the latest machine learning techniques with GraphML and TDA methods.

- **Anomaly Detection:** This model compares the behavior of an entity with its peers in a fine-grained, behavioral map. Using multiple techniques and the composite of these different

"lenses," we can find anomalous behavior in a very sophisticated yet intuitively explainable and visual way. Our anomalies have specific names in English.

- **Change in Behavior:** This model looks at how an entity's behavior changes over time and provides key behavioral indicators. Many other change-in-behavior techniques do not place the changes in context and do not provide much insight. Knowing whether a change in behavior is normal is extremely important to avoid generating noise.

- **New Behavior Evolution / Hotspots:** This model discovers suspicious behaviors missed by existing systems. This could be due to sophisticated adversarial behaviors of the money launderer, an incorrectly set Transaction Monitoring System (TMS) alert threshold or many other reasons. This model also detects where decision errors may have taken place, with an investigator incorrectly closing a case instead of escalating it.

These models work in tandem with each other, dynamically mapping to your risk appetite or focus on exposure. For example, the Risk Similarity model can reduce the false positive rate of the AML investigation system, thus freeing up the investigator's time to study new entities that were uncovered by the new behavior detection model — which have a significantly higher likelihood of indicating money laundering activity.

The Sensa platform is also designed with scalability in mind. While these are example models within the system, it is possible to add many more – models you create, third-party models, challenger models, or rules and models built with the Sensa modeling "Factory."

Explainable AI

Sensa models are built from the ground up with explainability in mind. Superior model performance alone is not enough. For our models to be trusted and used as part of a business workflow, all decisions need to be explainable and understandable to the broader business. This runs true from our behavioral discovery and alerts to the individual models themselves. No matter how sophisticated a model is under the hood, we strive to make the results and associated explanations as clear and understandable as possible.

Whenever a model outputs a score, the User Interface (UI) clearly shows which features, data and activity were used to reach that assessment. These are presented in simple, business-friendly language and compared to its peer group and the wider population. It's highly intuitive and can be used by an investigator to kick off a deeper dive investigation.

Accelerate time to value

Traditional Business Intelligence (BI), current machine learning and all applications demand rigid data management components in development and deployment, often consuming over 80% of the project time and cost.

With Sensa, the operational data model is "discovered" as an accurate reflection of the truth within your data…not some imposed, abstract data model. As part of the deployment, you'll discover what your data has to tell you and how your company is really operating.

The result? Accurate, real data mapping. A data model that genuinely reflects your firm's operations, market, customer environment and time-to-value discovery, at a fraction of the time and cost of anything else.

Network analysis and graph machine learning

Graph structures are a cornerstone of machine learning in Sensa, and they have many unique and interesting properties that make them challenging to analyze.

First off, edge distributions are typically power-law distributed. This means that a small number of nodes are involved in the vast majority of edges. In the context of financial transactions, you may find that a utility company is involved in many transactions…orders of magnitude greater than the number of transactions of the average individual. As a result, graphs become difficult to analyze, leading to numerous counter-intuitive results. For example, the "friendship paradox" states that within a social network, on average, a person is less famous than their friends. To analyze data of this nature, we therefore need more sophisticated methods and must simply and quickly combine multiple "analytical lenses" to increase resolution and fidelity.

We do this in a few ways. First, we statistically analyze the graph itself. We look for the localized structure of neighborhoods within the graph to create features that can be used within other machine-learning algorithms. To do this effectively, we have dedicated considerable time to finding an optimal set of features to use – features that come as part of the reusable Feature Store provided at deployment. These features also have real-world interpretability, increased model performance and aid in model explainability.

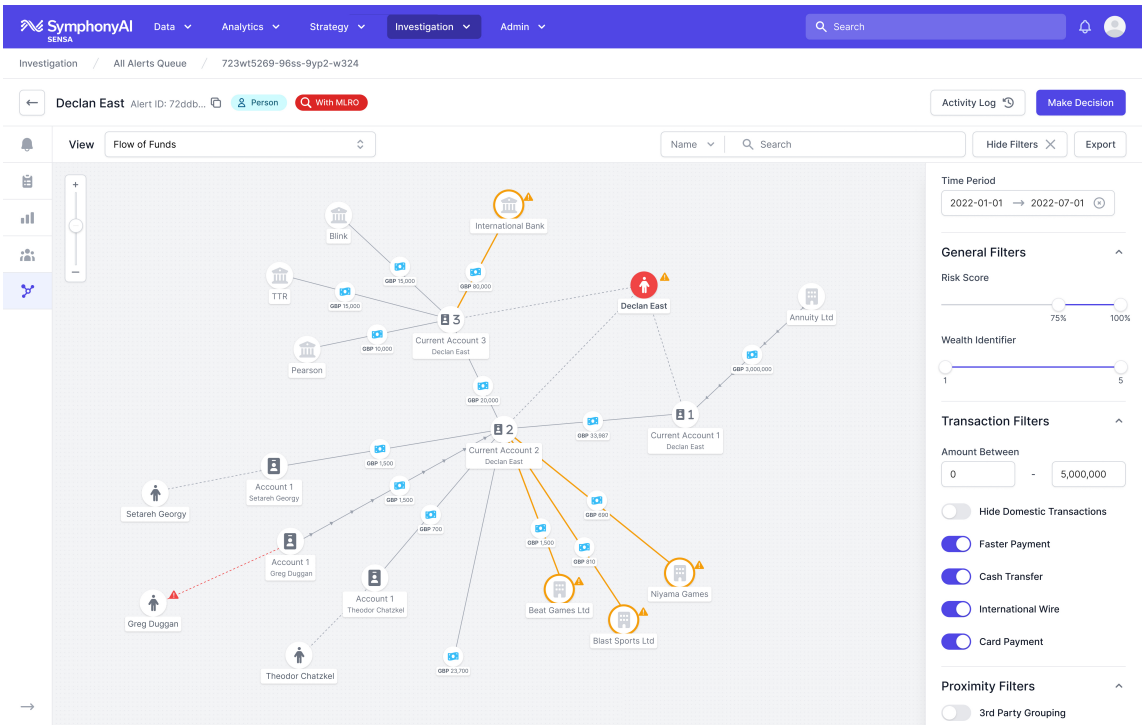

Figure 3

We also use graph representation methods extensively by taking a graph structure and embedding it in a vector space. The representations are used to ensure that both the local and global relationships of the transaction graph are consistent. This allows machine-learning algorithms to understand and use them. In other AI platforms, the information contained here would never be uncovered and would effectively be "thrown away" at the outset. For money laundering applications, this is clearly problematic, because laundering is a money flow through a complex network.

Finally, we also make use of several native, graph machine-learning algorithms where summary statistics and representations are not used. Instead, the machine-learning algorithms work directly on the graph structure itself. These are cutting-edge machine-learning methods that produce impactful results.

What about Topological Data Analysis (TDA)?

Topological Data Analysis is still a powerful tool to quickly visualize risk, and it remains a selection within Sensa's analytical components. It enables our clients to map different "lenses" to increase fidelity on hidden risks and attacks.

For example, there may be more than one "natural segmentation" in complex data sets. We typically look for a nested, hierarchical, behavioral description of customers. This allows us to identify both large-scale, "broad-brush" behavior as well as highly complex behaviors. This is a highly intuitive concept. For instance, in the real world, your behavior is likely similar in some respects to those who live on your street, but also equally similar to those within your town, city, state, country or even continent.

In recent years, there has been a flurry of research papers showing the power of combining TDA with machine learning and deep-learning AI methods for uncovering areas where machine learning "fails." At Sensa, we are ahead of the curve in this respect and have been combining TDA, GraphML and supervised and unsupervised machine learning for years. The key, though, is focusing the right technique or combination of techniques to maximize the "yield" of information from your data.

Our technology

The four main pillars that make up our technology are:

1. **Cloud-native**
2. **Security at design**
3. **Dynamic scalability**
4. **Data-driven**

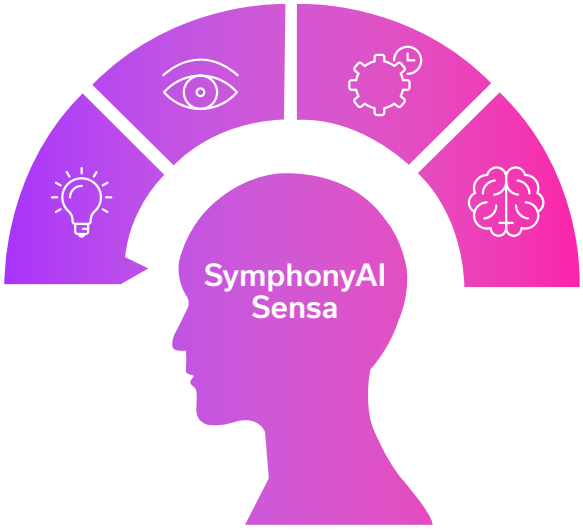

Figure 4

Being cloud-native:

Gives it the ability to run on any cloud or on-prem depending on the client's existing infrastructure. Our product uses modern concepts, technologies and frameworks, making SensaAML™ cloud-agnostic.

Security:

Was never an afterthought but lies at the core of building our platform. We gave it the ability to anonymize data without interfering with the ongoing data science efforts. The zero-trust model keeps the data secure. Available integration with identity and access management (IdAM) makes it a breeze to secure and use.

Dynamic scalability:

We use a domain-orientated microservices architecture (DOMA) comprised of subsystems made of independent microservices orchestrated within a Kubernetes cluster. This ensures that the way Sensa is deployed is incredibly flexible, resilient and supports far more dynamic use of resources e.g., scaling out infrastructure only when needed.

Data-driven:

The more data we consume, the more insights we can find resulting in event-driven ingestion and processing. The system data repository grows organically with data demands. Based on a less-schema model, we use both time series and graph databases to ingest data and use machine learning to map entity relationships dynamically. Our ability to map these relationships allows our customers to graphically traverse the networks of their customers' relationships aided by Sensa's supervised and unsupervised machine-learning predictions.

Sensa uses a common application software (CAS) framework: solutions are dynamically configured without the need to modify the core code of the platform. So, you get faster deployment, agility and nimbleness as requirements change and a seamless, simple integration with any existing workflows or systems.

Summary

Enterprises that adopt AI and GraphML will outperform enterprises that don't. Regulators have asked Financial Institutions (FIs) to experiment and innovate towards a risk-based approach rather than the old and failing rules-based approach. Artificial Intelligence will define the winners and losers over the coming years.

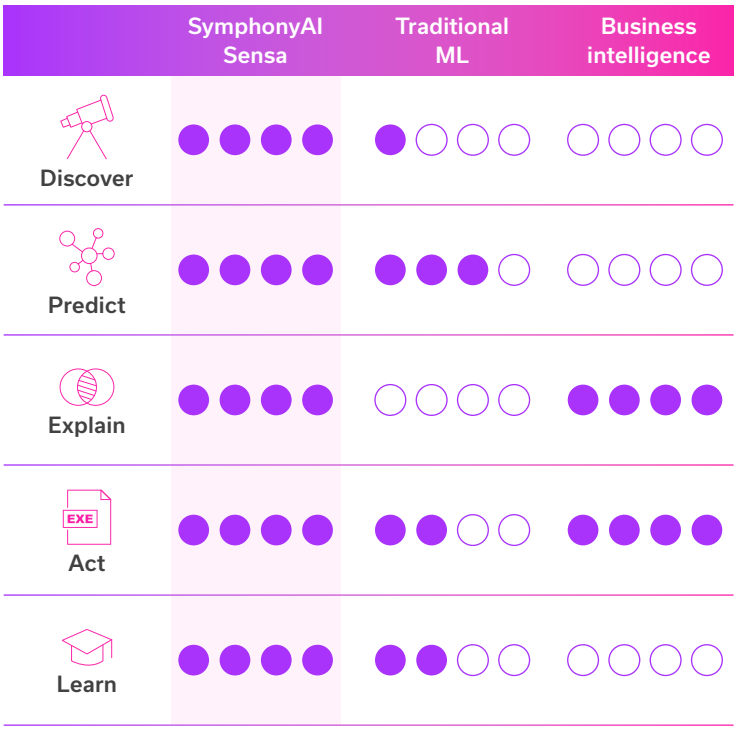

Figure 5: Sensa analytical capabilities

As a company, Sensa has pioneered the development of this next-gen platform. The four pillars outlined at the start of this paper: cloud-native, security at design, dynamic scalability, and data-driven, represent the necessary features for a next-gen, future-proof platform. A commitment to establishing an AI-driven AML strategy requires a commitment by every risk department — however, the payoff for those FIs making that investment is immense.

About SymphonyAI Sensa

SymphonyAI Sensa provides unsurpassed crime detection services to financial institutions worldwide. Powered by a proprietary combination of advanced AI and machine learning, Sensa discovers criminal activity routinely unseen by other detection systems. Customers have achieved up to an 81x risk-to-alert ratio improvement, 263% increase in SARs and as much as a 77% reduction in false positives. Coupled with predictive insights, operational efficiencies are gained by identifying genuine anomalies from the customers' current data set. Sensa is user-friendly and delivers fully explainable results to satisfy regulators. Learn more at www.symphonysensa, LinkedIn or Twitter.

About SymphonyAI

SymphonyAI is building the leading enterprise AI company for digital transformation across the most important and resilient growth industries, including retail, consumer packaged goods, financial services, manufacturing, media, and IT service management. SymphonyAI businesses have many leading enterprises as clients in each of these industries. Since its founding in 2017, SymphonyAI has grown rapidly, approaching 2,000 talented leaders, data scientists, and other professionals. SymphonyAI is an SAIGroup company, backed by a $1 billion commitment from successful entrepreneur and philanthropist Dr. Romesh Wadhwani.