Graph LLMs: AI's next frontier in banking and insurance

pwc

The real question is who isn't investing in AI—and at what cost?

In an era where data is the new oil, the banking and insurance sectors are sitting on untapped reservoirs of data to be harnessed using Artificial Intelligence (AI). Global AI investment in banking and insurance sector is projected to rise significantly. In banking sector alone, GenAI spending is estimated to reach approximately US$85 billion[1] by 2030 – a 1,400% increase from 2024's US$6 billion – driven by a clear focus on efficiency and customer experience enhancements. This acceleration could generate up to $2 trillion[2] in revenue by 2028, boosting banking sector profitability by an estimated 9% through productivity gains. Similarly, AI investments in the global insurance industry are projected to grow nearly 60% by 2027, reaching $15.9 billion[3].

AI is attracting significant investments across the Middle East, with governments and private enterprises allocating substantial resources to AI R&D and innovation hubs. These investments are aligned with national strategies, like Saudi Vision 2030 and the UAE's National AI Strategy 2031.

Despite this substantial investment and advancements in the technology, traditional AI and Machine Learning (ML) models typically treat data points in isolation and often face limitations when handling complex, relational data inherent in the financial services and insurance sectors.

The rise of Graph Language Models (Graph LLMs) promises to uncover the nuanced relationships that drive deeper insights and more accurate predictions.

In this article, we explore how Graph LLMs are transforming the banking and insurance industries by boosting operational efficiency, enhancing customer experiences, and driving revenue growth. We explain core concepts, strategic implementation approaches and key recommendations for executives to harness the transformative potential of Graph LLMs.

Key insights

Graph LLMs enable hyper-personalised customer interactions, enhancing loyalty and service satisfaction.

Real-time, relationship-aware dynamic pricing models offer competitive and precise financial products.

Advanced pattern recognition capabilities provide robust security against sophisticated fraud schemes.

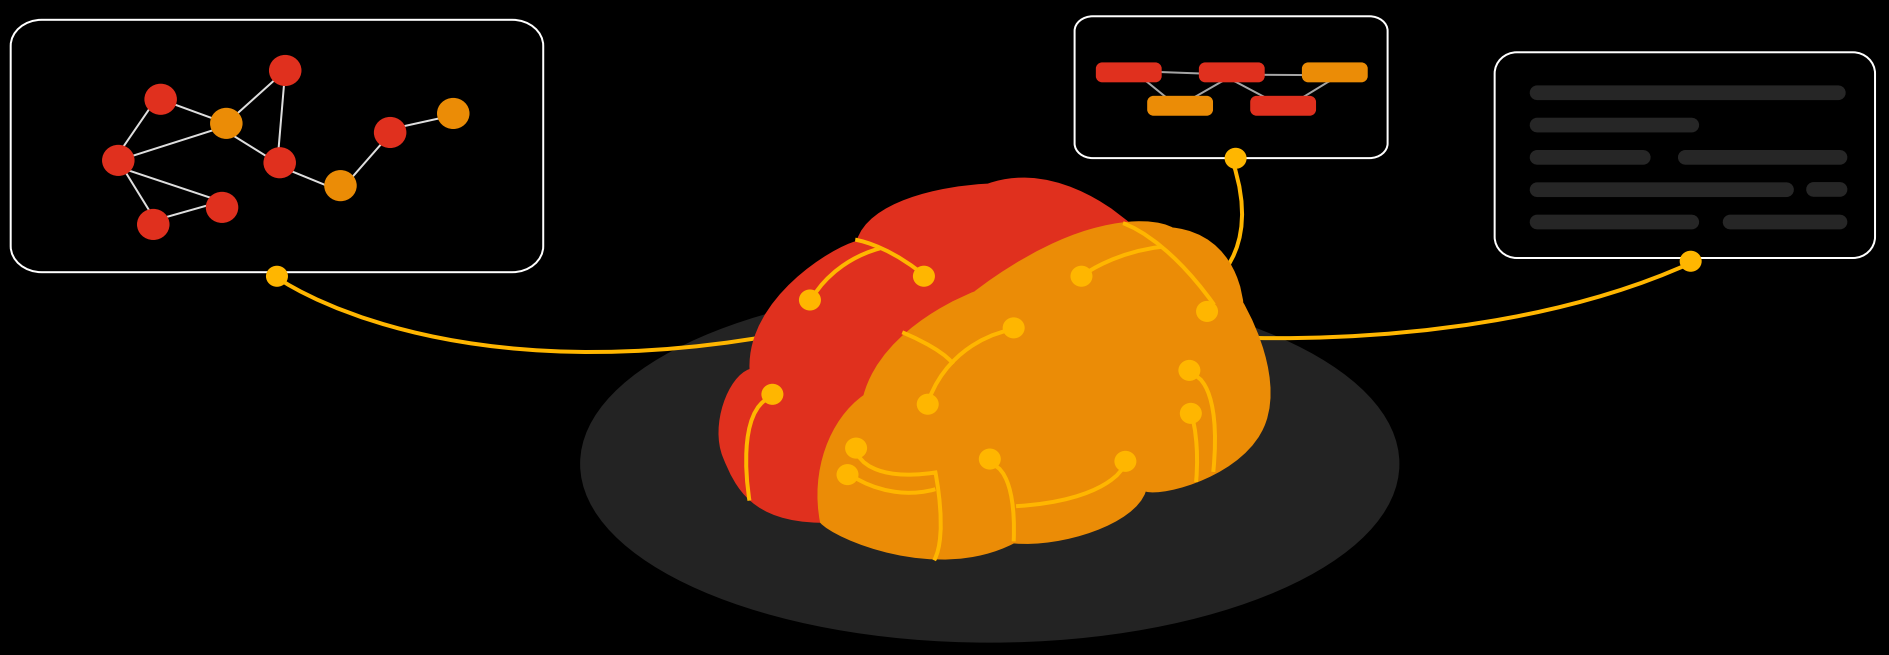

What are Graph LLMs?

Graph LLMs refer to the integration of Graph Neural Networks (GNNs) with Large Language Models (LLMs) to enhance the understanding and processing of data that has both textual and relational (graph-structured) components. This hybrid approach leverages the strengths of LLMs in natural language understanding and the capabilities of GNNs in modelling relationships between entities.

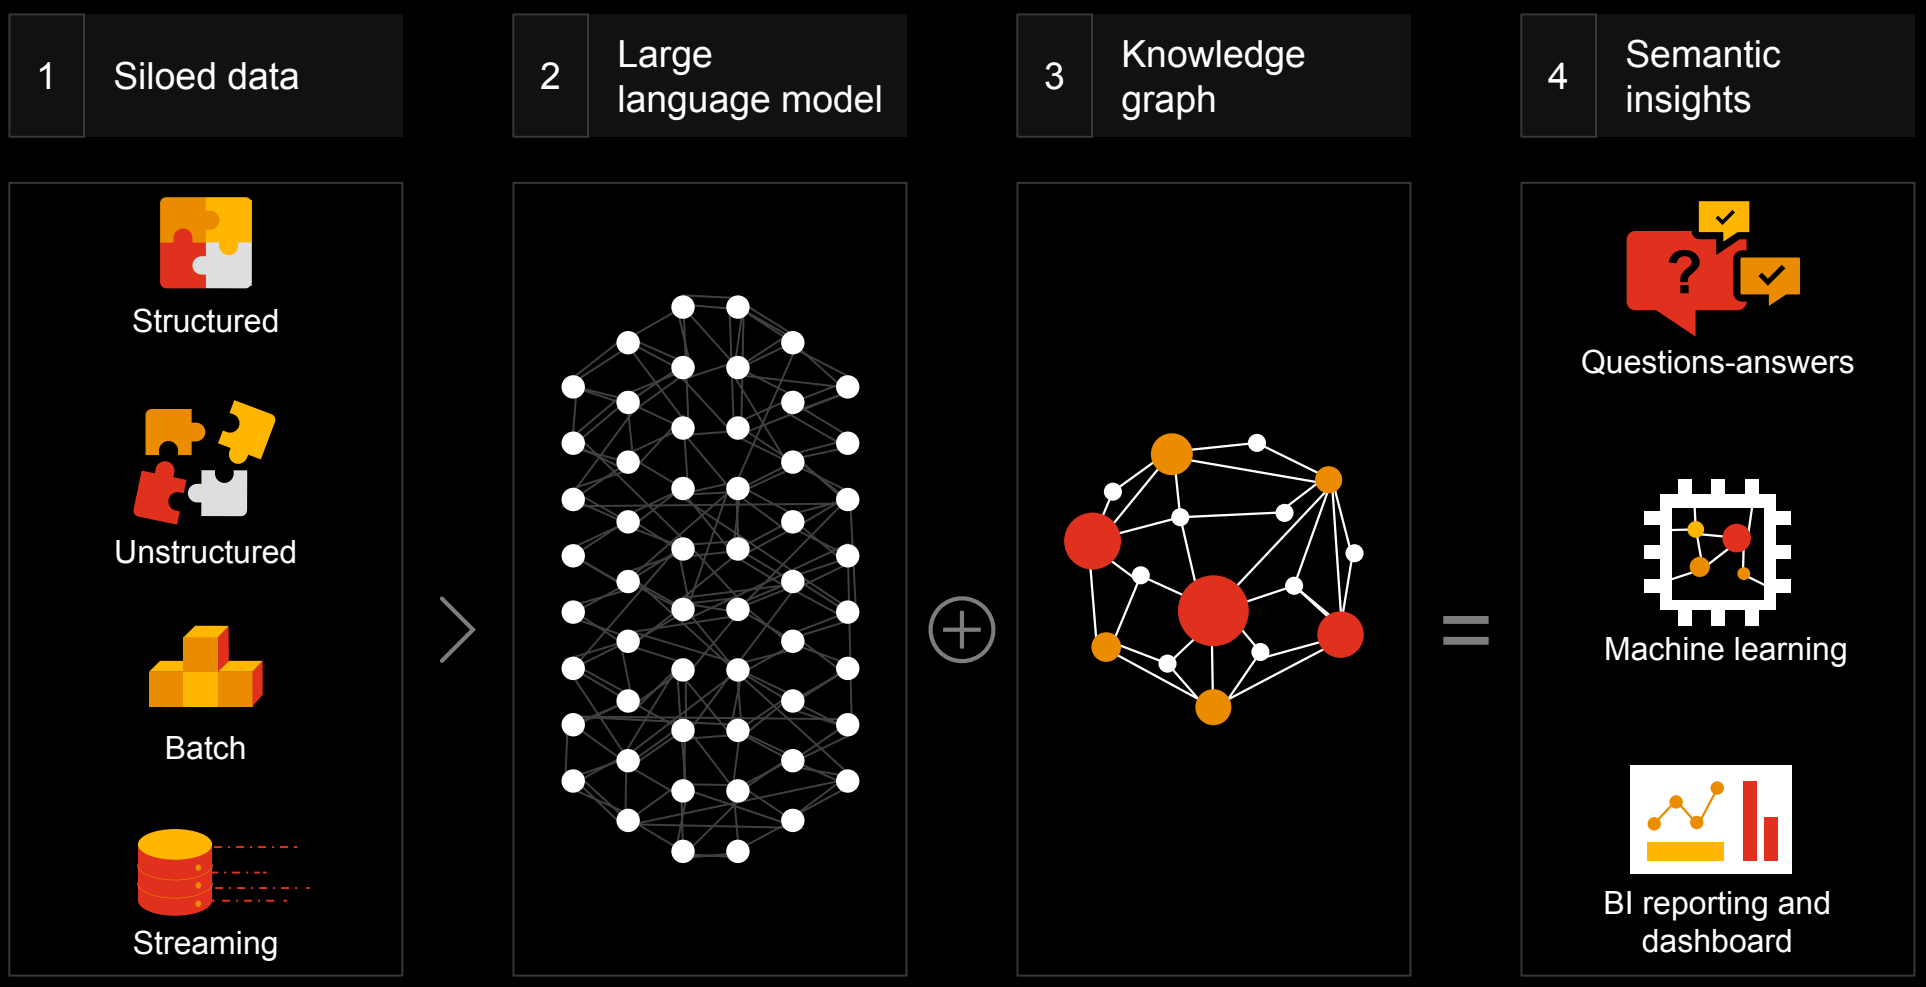

Key concepts

Large Language Models (LLMs)

Trained on vast textual datasets, LLMs like GPT-4, BERT, and LLaMA excel at understanding context, generating human-like text and performing various language tasks.

Graph Neural Networks (GNNs)

GNNs are designed to operate on graph-structured data, where entities are nodes connected by edges. They effectively capture relationships and interactions between entities in networks such as social graphs, knowledge graphs and molecular structures.

Integration of graphs and LLMs

Graph-augmented language models

Incorporate graph embeddings or representations into LLMs. Enable the model to consider relational information alongside textual data.

Knowledge graphs and language models

Use knowledge graphs to provide structured information that enhances language understanding. Improve tasks like question answering by providing context about entities and their relationships.

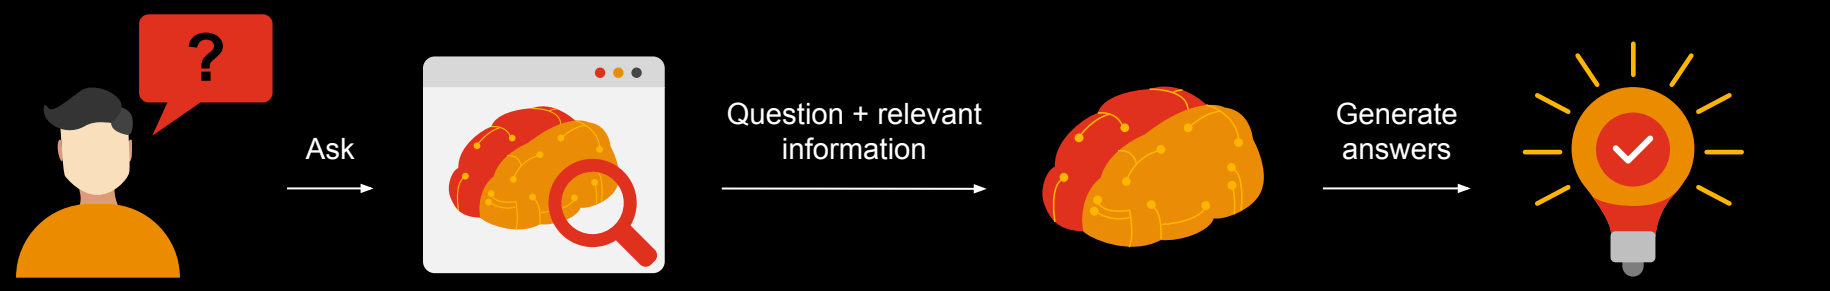

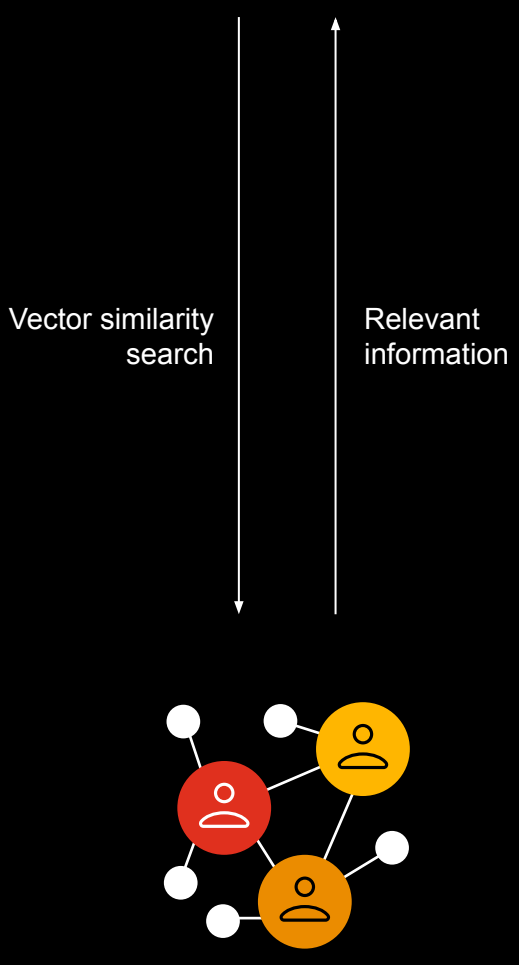

Applications

Question-answering systems

Enhancing LLMs with graph data allows for better comprehension of complex queries involving multiple entities. Improves accuracy in providing precise answers.

Knowledge graph completion

Predicting missing links or nodes in a knowledge graph by leveraging textual descriptions and relational data.

Natural language understanding

Better entity recognition and disambiguation by considering relationships in a graph. Improved semantic understanding for tasks like summarisation and translation.

Recommendation systems

Combining user interaction graphs with textual data to provide personalised recommendations.

Discover the nuances in data

Graph LLMs combine the prowess of GNNs with the linguistic capabilities of LLMs leveraging graph databases. They excel in handling relational data and contextual nuances, making them ideal for sectors where relationships are as valuable as the entities themselves.

Outpacing tradition: The unmatched edge of Graph LLMs

1

Relational data handling

Graph LLMs can model complex relationships between data points, uncovering hidden patterns.

2

Contextual nuance

They understand context at a granular level, improving the accuracy of insights.

3

Business insights

By mapping relationships, they provide actionable intelligence for specific operational units.

4

Efficient deep searches

Queries that require deep joins across several tabular data can be slow and resource intensive in traditional relational databases. Graph databases behind Graph LLMs offer natural language queries for business users and excel in traversing relationships, offering faster query responses for business users.

5

Flexibility and agility

The rigid schema of relational databases make it difficult to adapt to change. Graph databases behind Graph LLMs offer a more flexible schema, which is easier to evolve as business needs change.

6

Intuitive data modelling

Graph databases allowed for a more intuitive representation of real-world scenarios, making it easier for business users to understand and work with their data.

The Graph LLMs and Agentic AI connection

The rise of agentic AI introduces a new dimension to Graph LLMs. While Graph LLMs excel at analysing and understanding relationships in data, agentic AI can take it further by autonomously acting on these insights. Instead of merely providing predictions or recommendations, Agentic AI can make real-time decisions, prioritise actions, and even execute tasks independently. This capability enables systems to adapt and optimise themselves without human intervention, as outlined in our recent PwC Middle East report.

Working together, Graph LLMs and Agentic AI are capable of creating powerful, self-sustaining systems that not only understand the data but can also take meaningful actions that drive business outcomes. This synergy is poised to revolutionise industries, providing businesses with smarter, more autonomous solutions.

Breaking barriers: How Graph LLMs surpass traditional models

Traditional LLMs have taken us far, enabling chatbots, virtual assistants and automated reporting. However, they operate in isolation, treating data points as unrelated entities. Graph LLMs overcome this limitation by understanding and leveraging the relationships between data points – such as, customers, transactions, policies and claims – in a networked, interconnected manner.

This relational intelligence is a game changer. It allows for more accurate predictions, nuanced customer interactions and robust fraud detection mechanisms. In a world where data is the new oil, Graph LLMs act as the refineries, unlocking its full value.

Here are some specific ways in which Graph LLMs surpass traditional models:

Handling relational data

Graph LLMs excel at modelling networks – whether social connections, transaction histories, or organisational structures – enabling them to detect fraud rings, identify cross-selling opportunities and optimise supply chains.

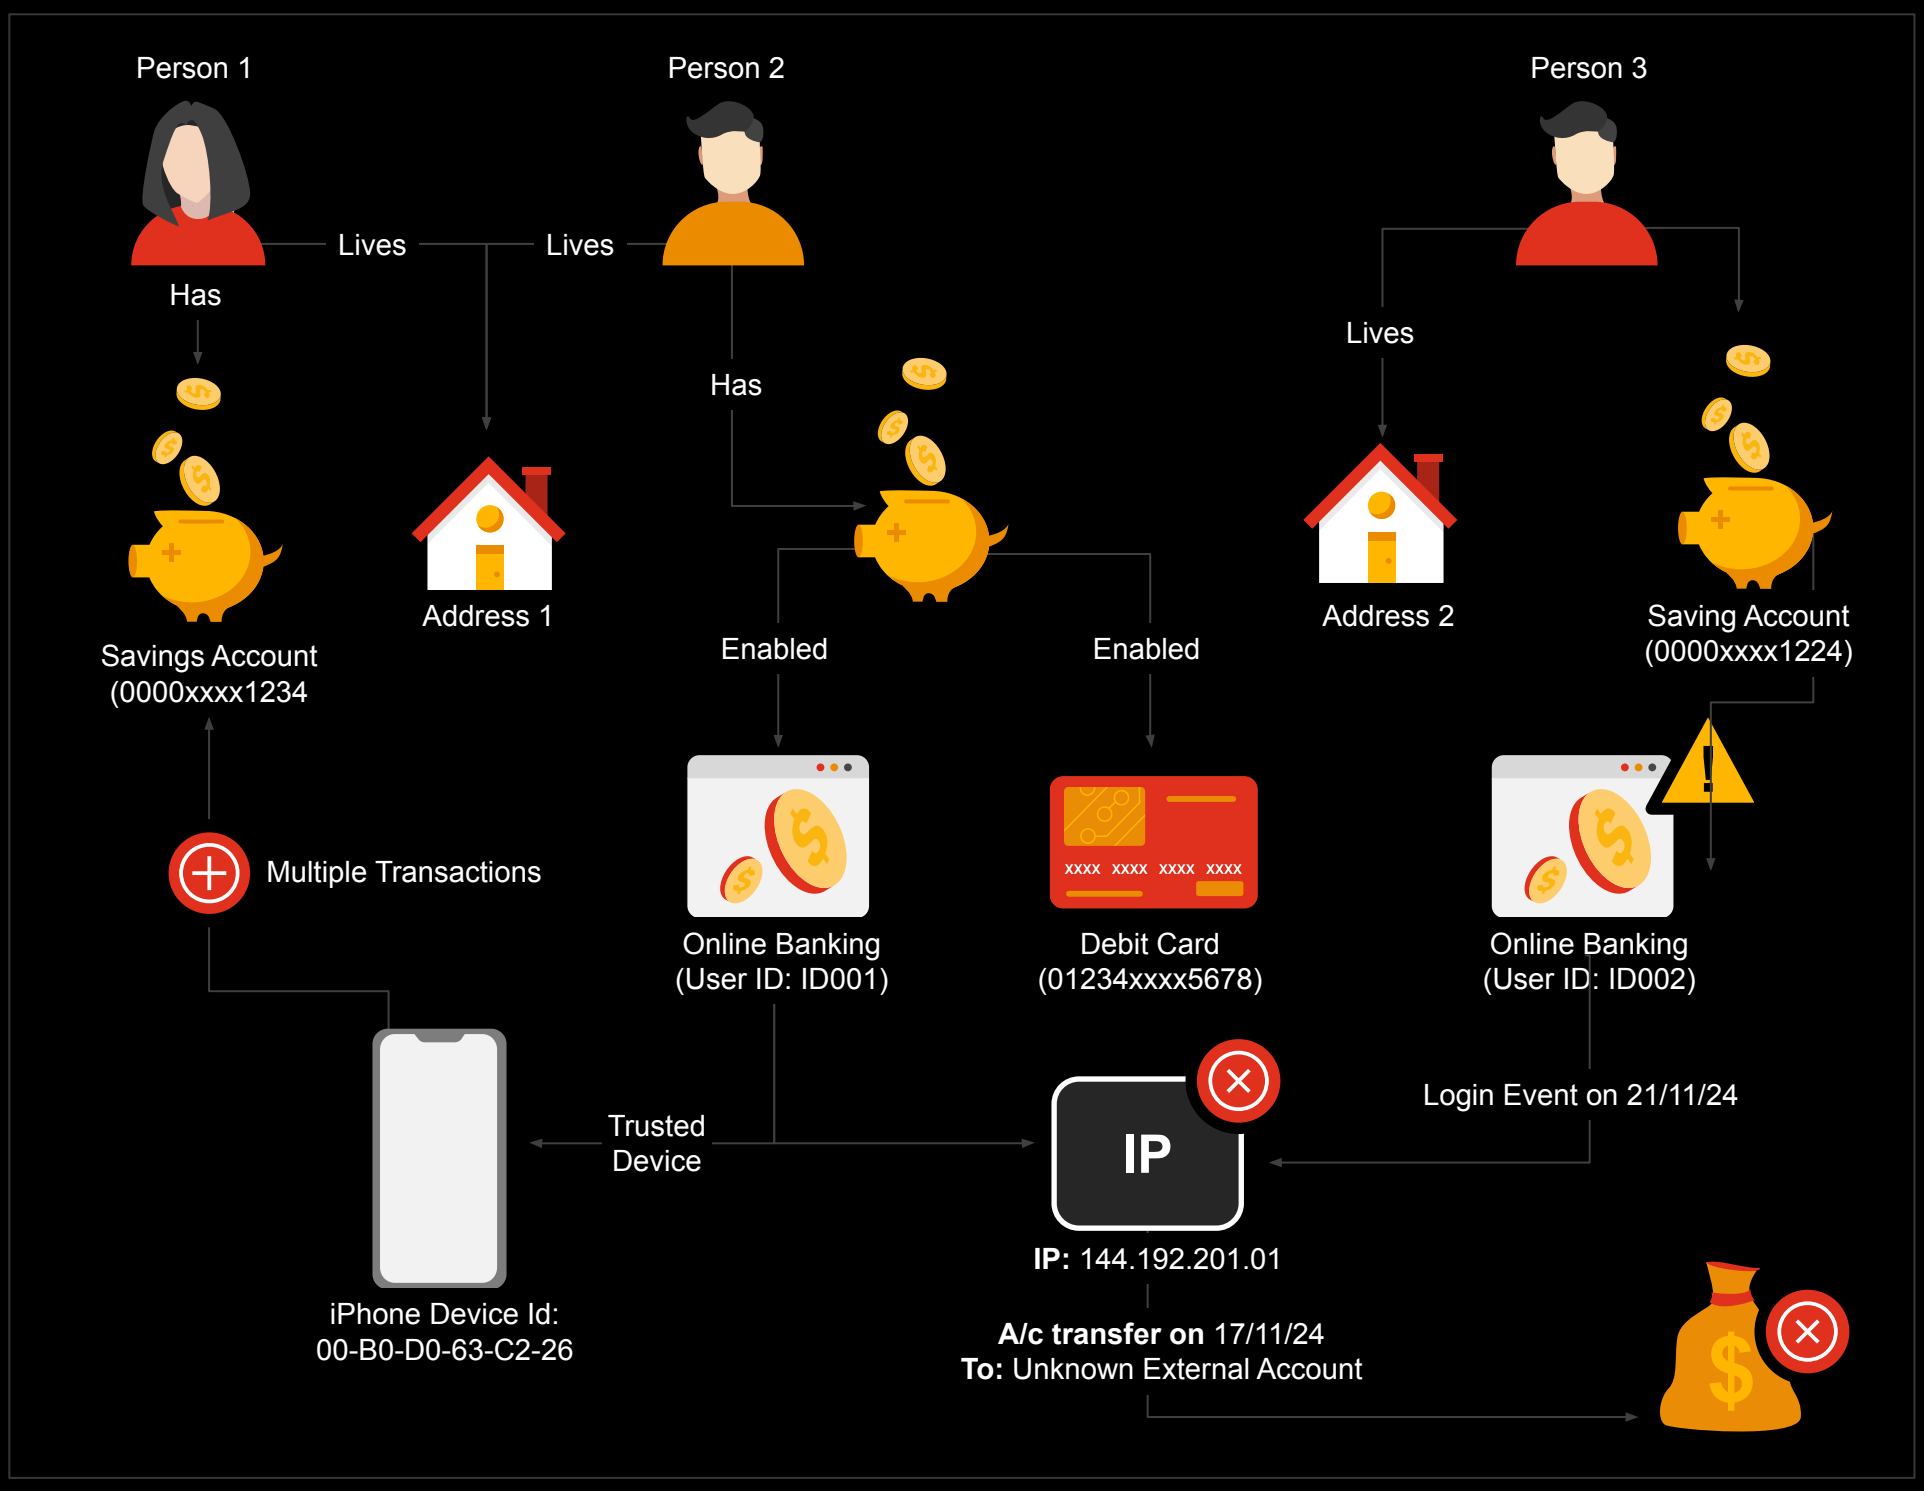

Contextual nuances

Understanding context is crucial in language and data interpretation. Graph LLMs capture semantic relationships, improving the accuracy of natural language processing tasks such as sentiment analysis and customer feedback interpretation.

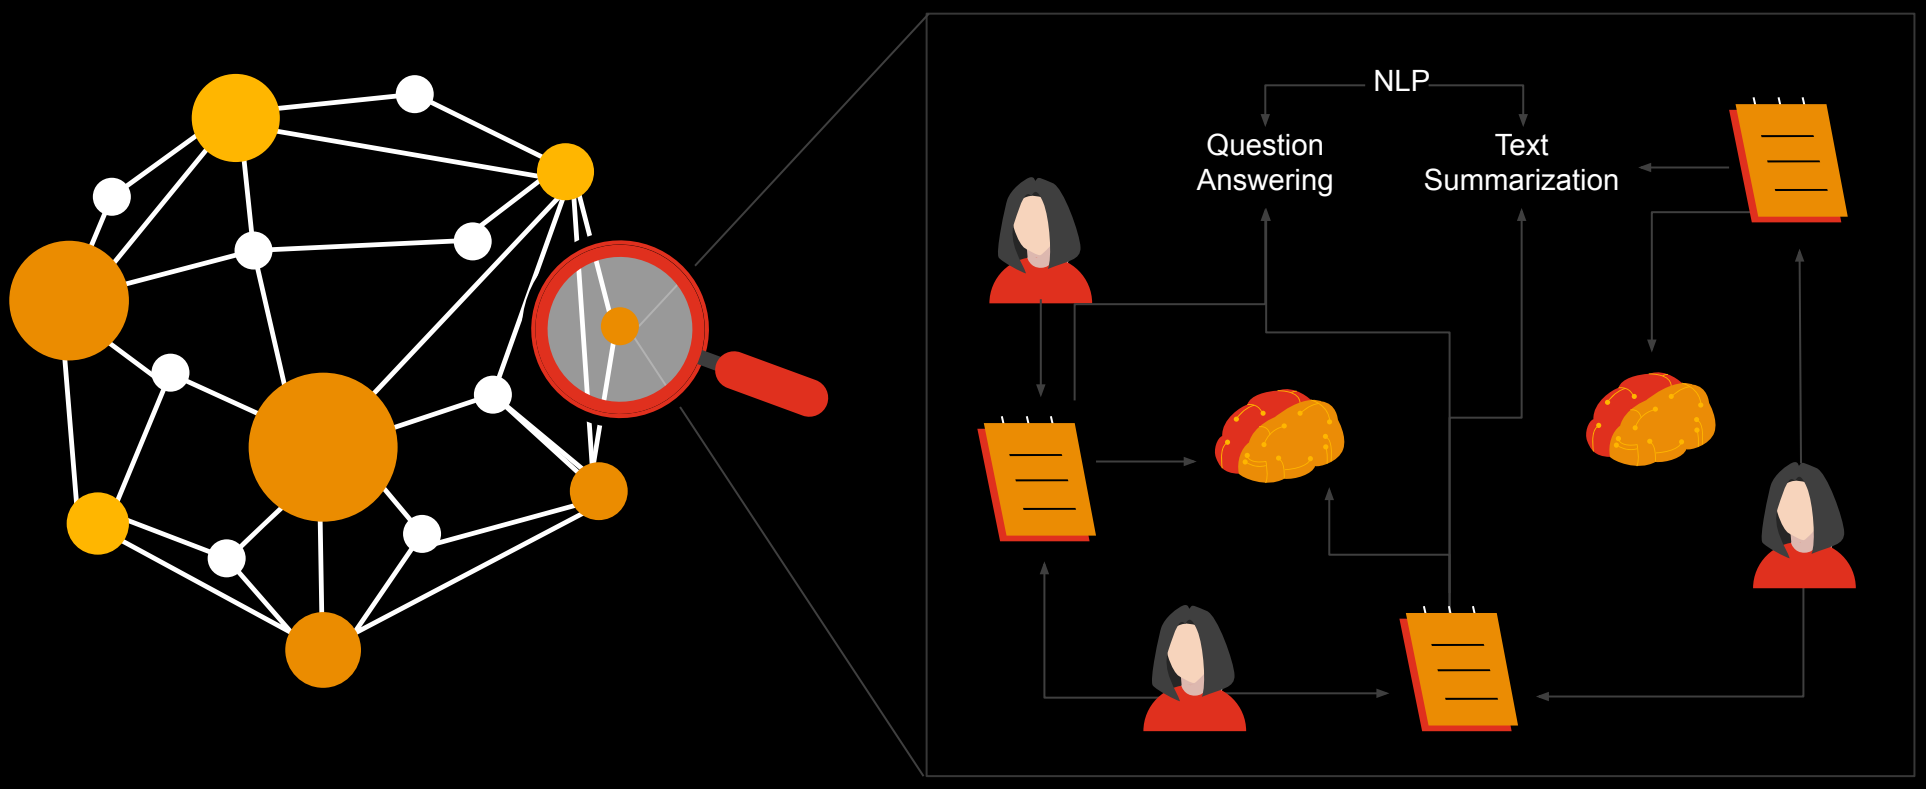

Uncovering complex business insights

By analysing interconnected data, Graph LLMs uncover insights often missed by siloed data analysis. They enhance scenario planning, risk assessment and strategic forecasting with higher precision.

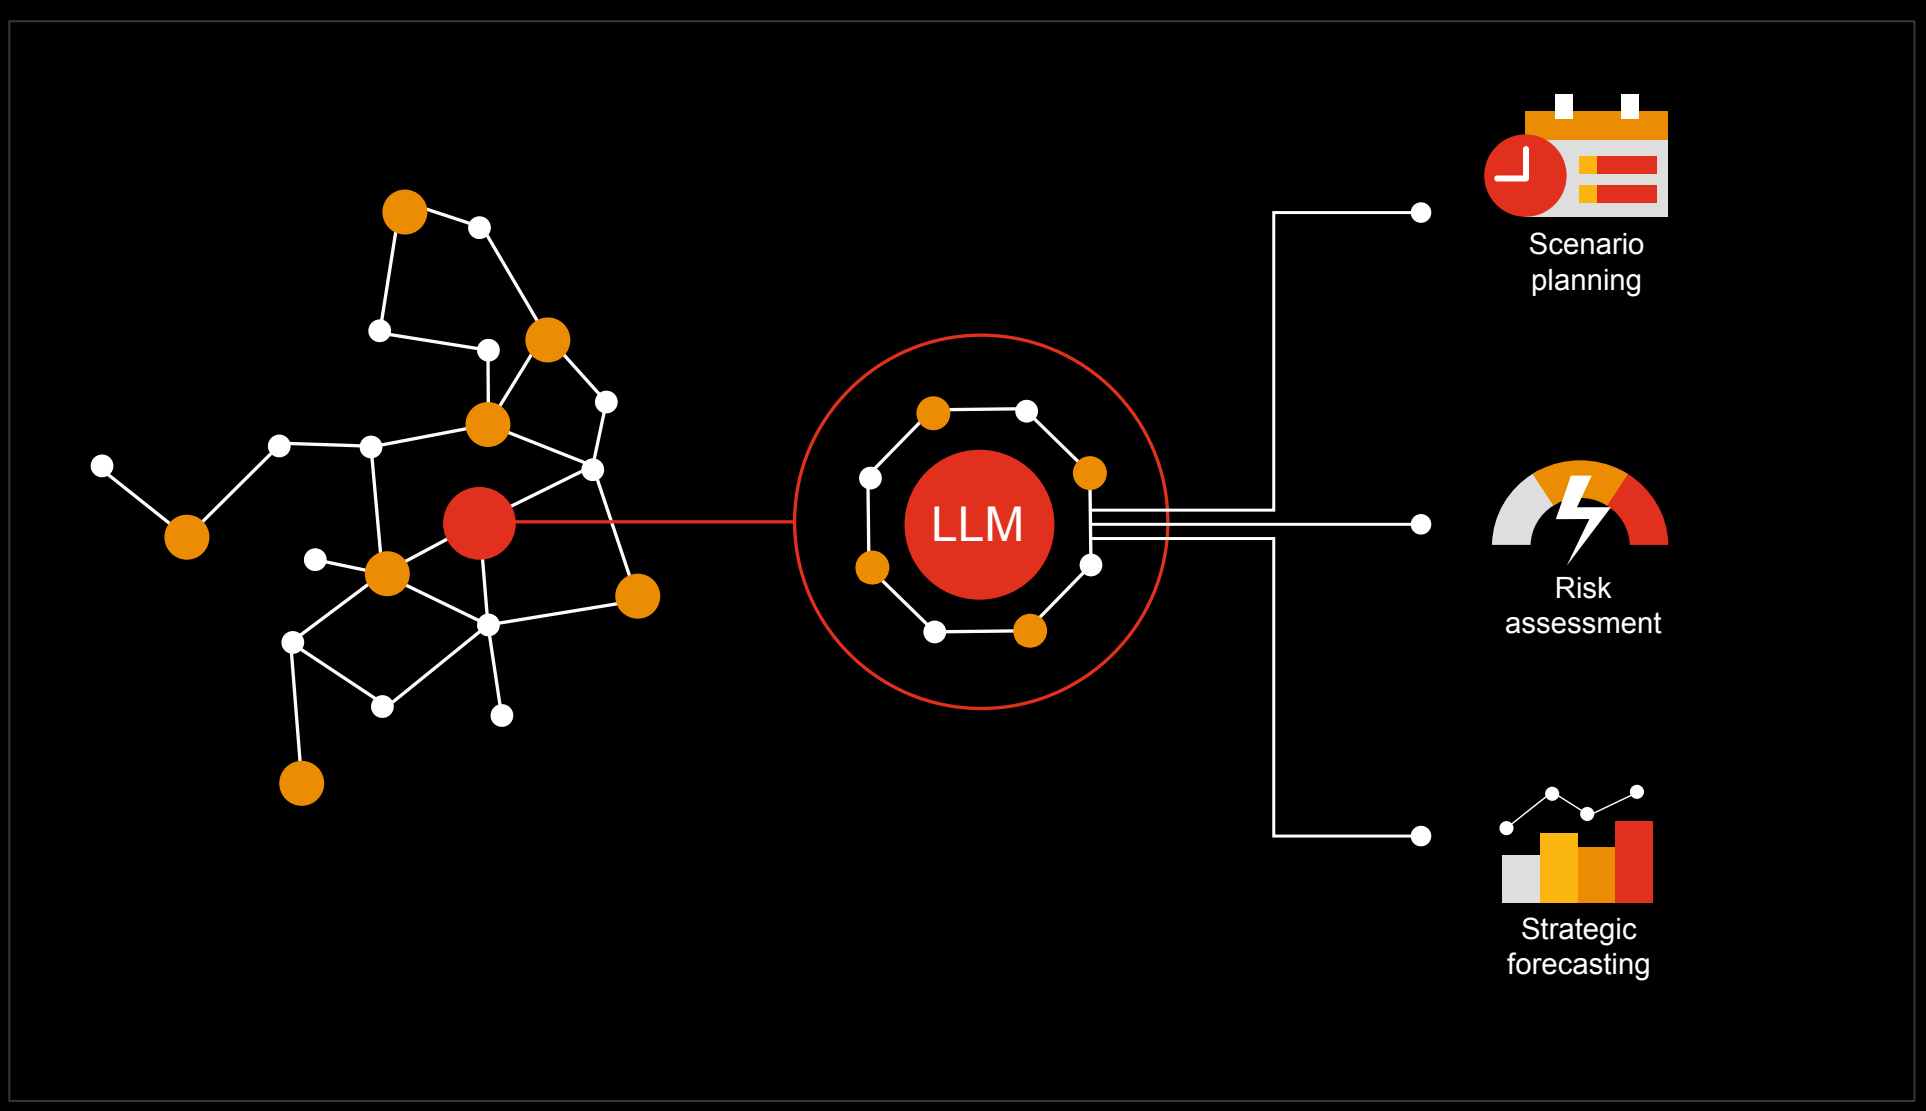

Graph LLMs: The competitive advantage

By integrating the structural analysis of GNNs with the linguistic depth of LLMs, Graph LLMs bring unprecedented clarity and efficiency to various business functions across the insurance and banking sectors.

Personalising customer experience like never before

In the era of Graph LLMs, every customer feels like they're your only customer.

The status quo

Personalised banking and insurance services are often limited to basic demographic segmentation.

The Graph LLM advantage

Graph LLMs enable hyper-personalisation by considering not just the customer but their entire network – social connections, transaction histories, and even real-time behavioural data. This leads to tailored product recommendations, proactive service offerings, and enhanced customer loyalty.

Dynamic pricing models

In a world of constant change, Graph LLMs ensure your pricing stays ahead of the curve.

The status quo

Traditional pricing in banking and insurance relies heavily on static risk assessments and historical data. Interest rates, premiums, and policy terms are often fixed without accounting for real-time changes in customer behaviour or market conditions. This rigidity can lead to mispriced products, either overcharging customers or exposing institutions to undue risk.

The Graph LLM advantage

With Graph LLMs, institutions can implement dynamic pricing models that adjust in real time based on a multitude of interconnected factors. These models consider not just individual customer data but also their network – social influences, market trends, and behavioural patterns. This leads to pricing that's more accurate, competitive, and personalised.

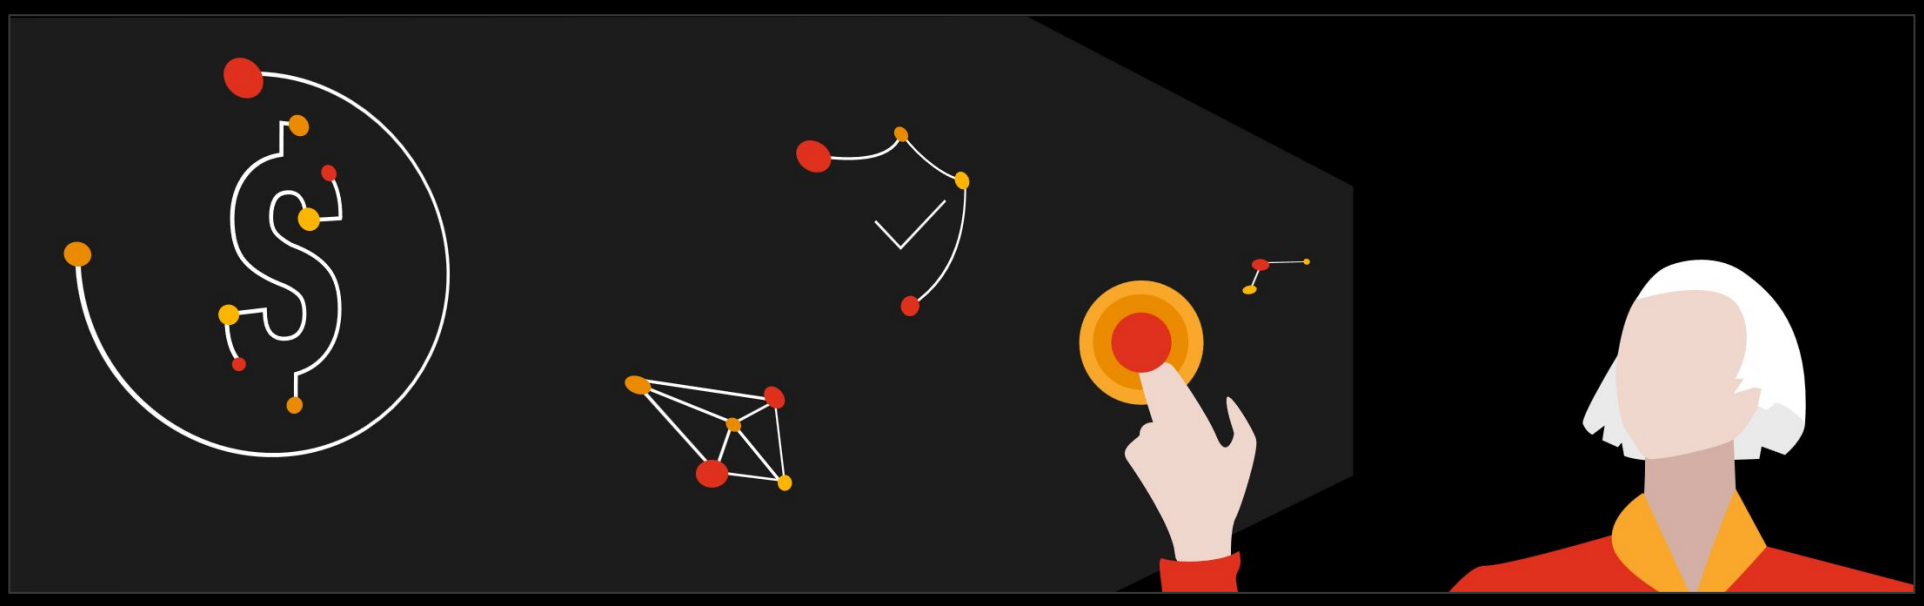

Revolutionising credit scoring and underwriting

Graph LLMs turn credit scoring from a snapshot into a living portrait.

The status quo

Credit scoring models rely heavily on historical data and often fail to accommodate unconventional data sources.

The Graph LLM advantage

By incorporating alternative data points – such as social media activity, payment networks, and peer relationships – Graph LLMs provide a more holistic and accurate assessment of creditworthiness and risk. This opens doors to underbanked populations and niche markets.

Automated claims processing

Graph LLMs transform claims processing from a bottleneck into a bullet train.

The status quo

In the current landscape, insurance claims processing is often a sluggish and labour-intensive ordeal. Adjusters manually sift through paperwork, verify details across disparate systems, and sometimes even conduct in-person assessments. This not only delays settlements but also opens the door to human error and fraudulent claims slipping through the cracks. Customers are left waiting, frustration mounts, and trust erodes.

The Graph LLM advantage

Graph LLMs can revolutionise this process by automating claim verification and decision-making. By analysing relationships between policyholders, past claims, incident reports, and external data sources in a graph structure, these models can quickly validate the authenticity of a claim. They can spot anomalies and patterns indicative of fraud that traditional linear models might miss. The result? Faster claim approvals, reduced operational costs, and a significant boost in customer satisfaction.

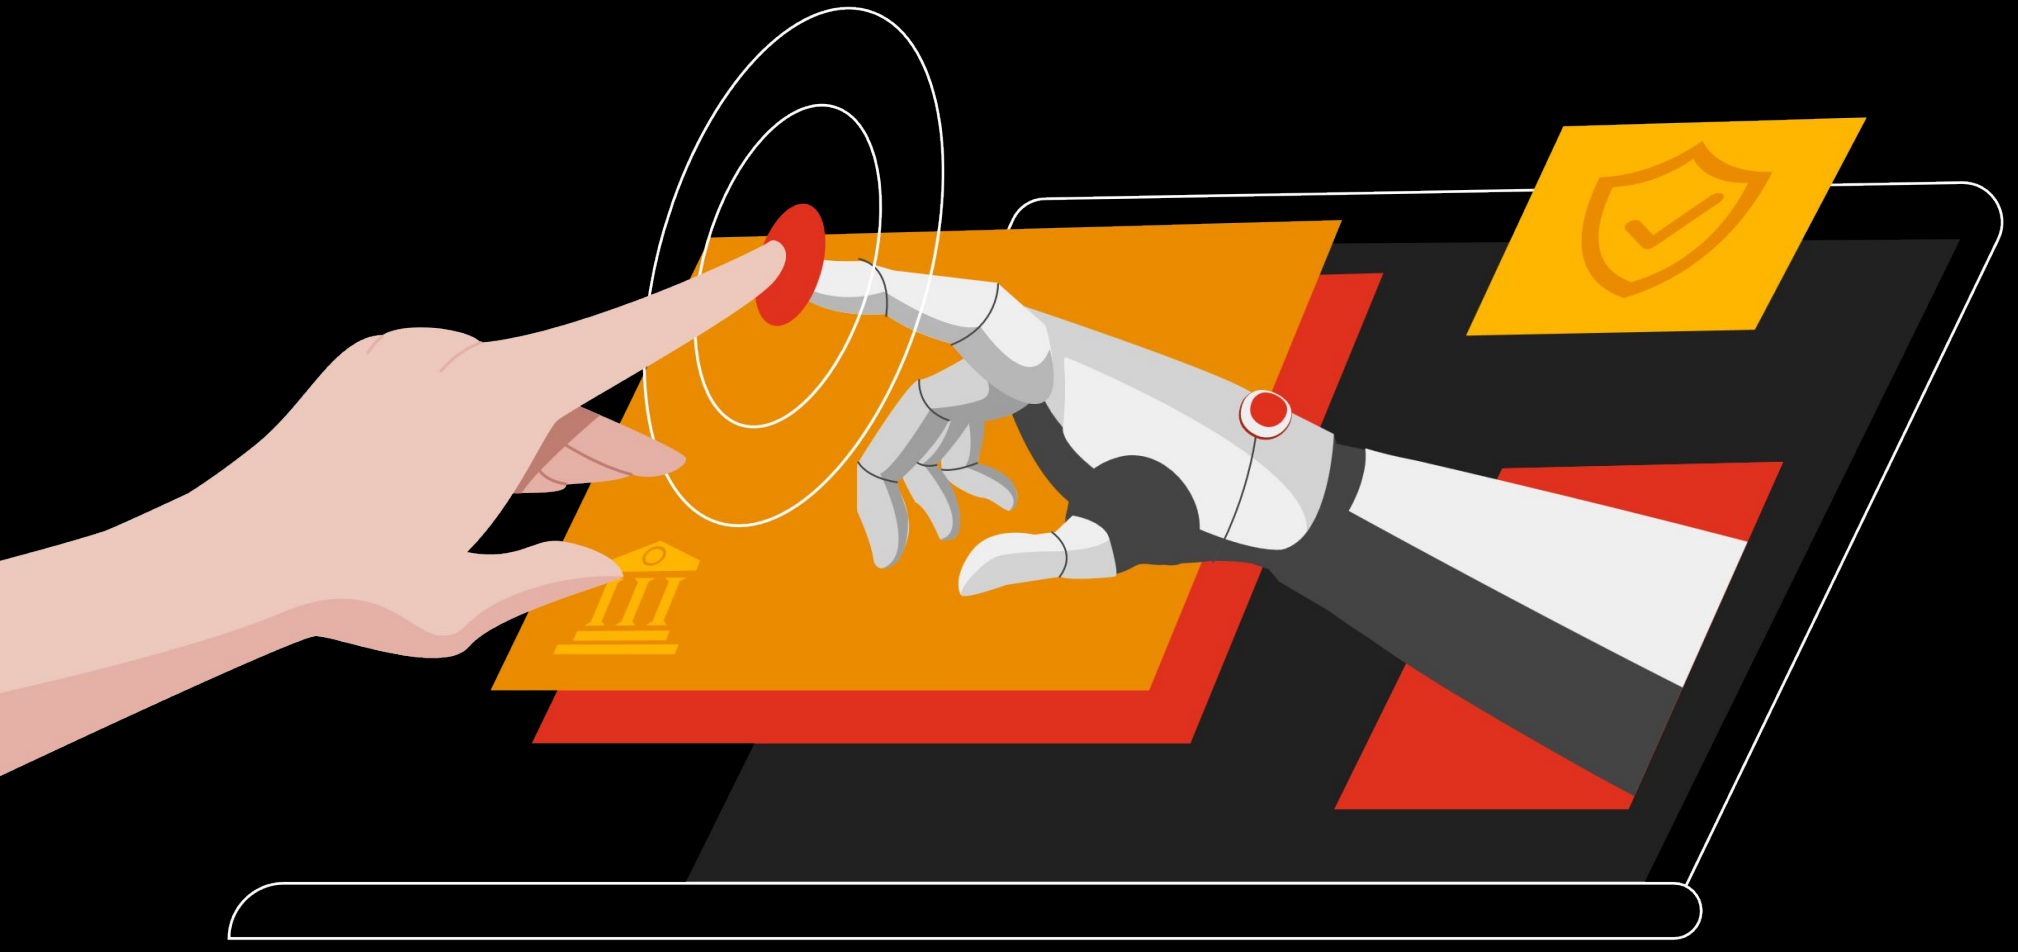

Transforming fraud detection and risk management

Graph LLMs don't just catch the fraudsters you know – they expose the ones you don't.

The status quo

Banks and insurers lose billions of dollars annually to fraud. Traditional AI models detect anomalies but often miss sophisticated schemes that exploit network relationships.

The Graph LLM advantage

By mapping and analysing the intricate web of relationships between entities, Graph LLMs can identify hidden patterns indicative of fraud. For instance, they can detect collusion between seemingly unrelated accounts or policies, flagging risks that linear models overlook.

Enhancing regulatory compliance and reporting

Stay ahead of regulators without breaking a sweat – Graph LLMs have got your back.

The status quo

Compliance is a resource-intensive necessity, fraught with the risk of costly penalties for non-compliance.

The Graph LLM advantage

Graph LLMs can automate the complex task of compliance by understanding regulatory requirements in the context of an organisation's entire data ecosystem. They can flag potential compliance issues in real time and streamline reporting processes.

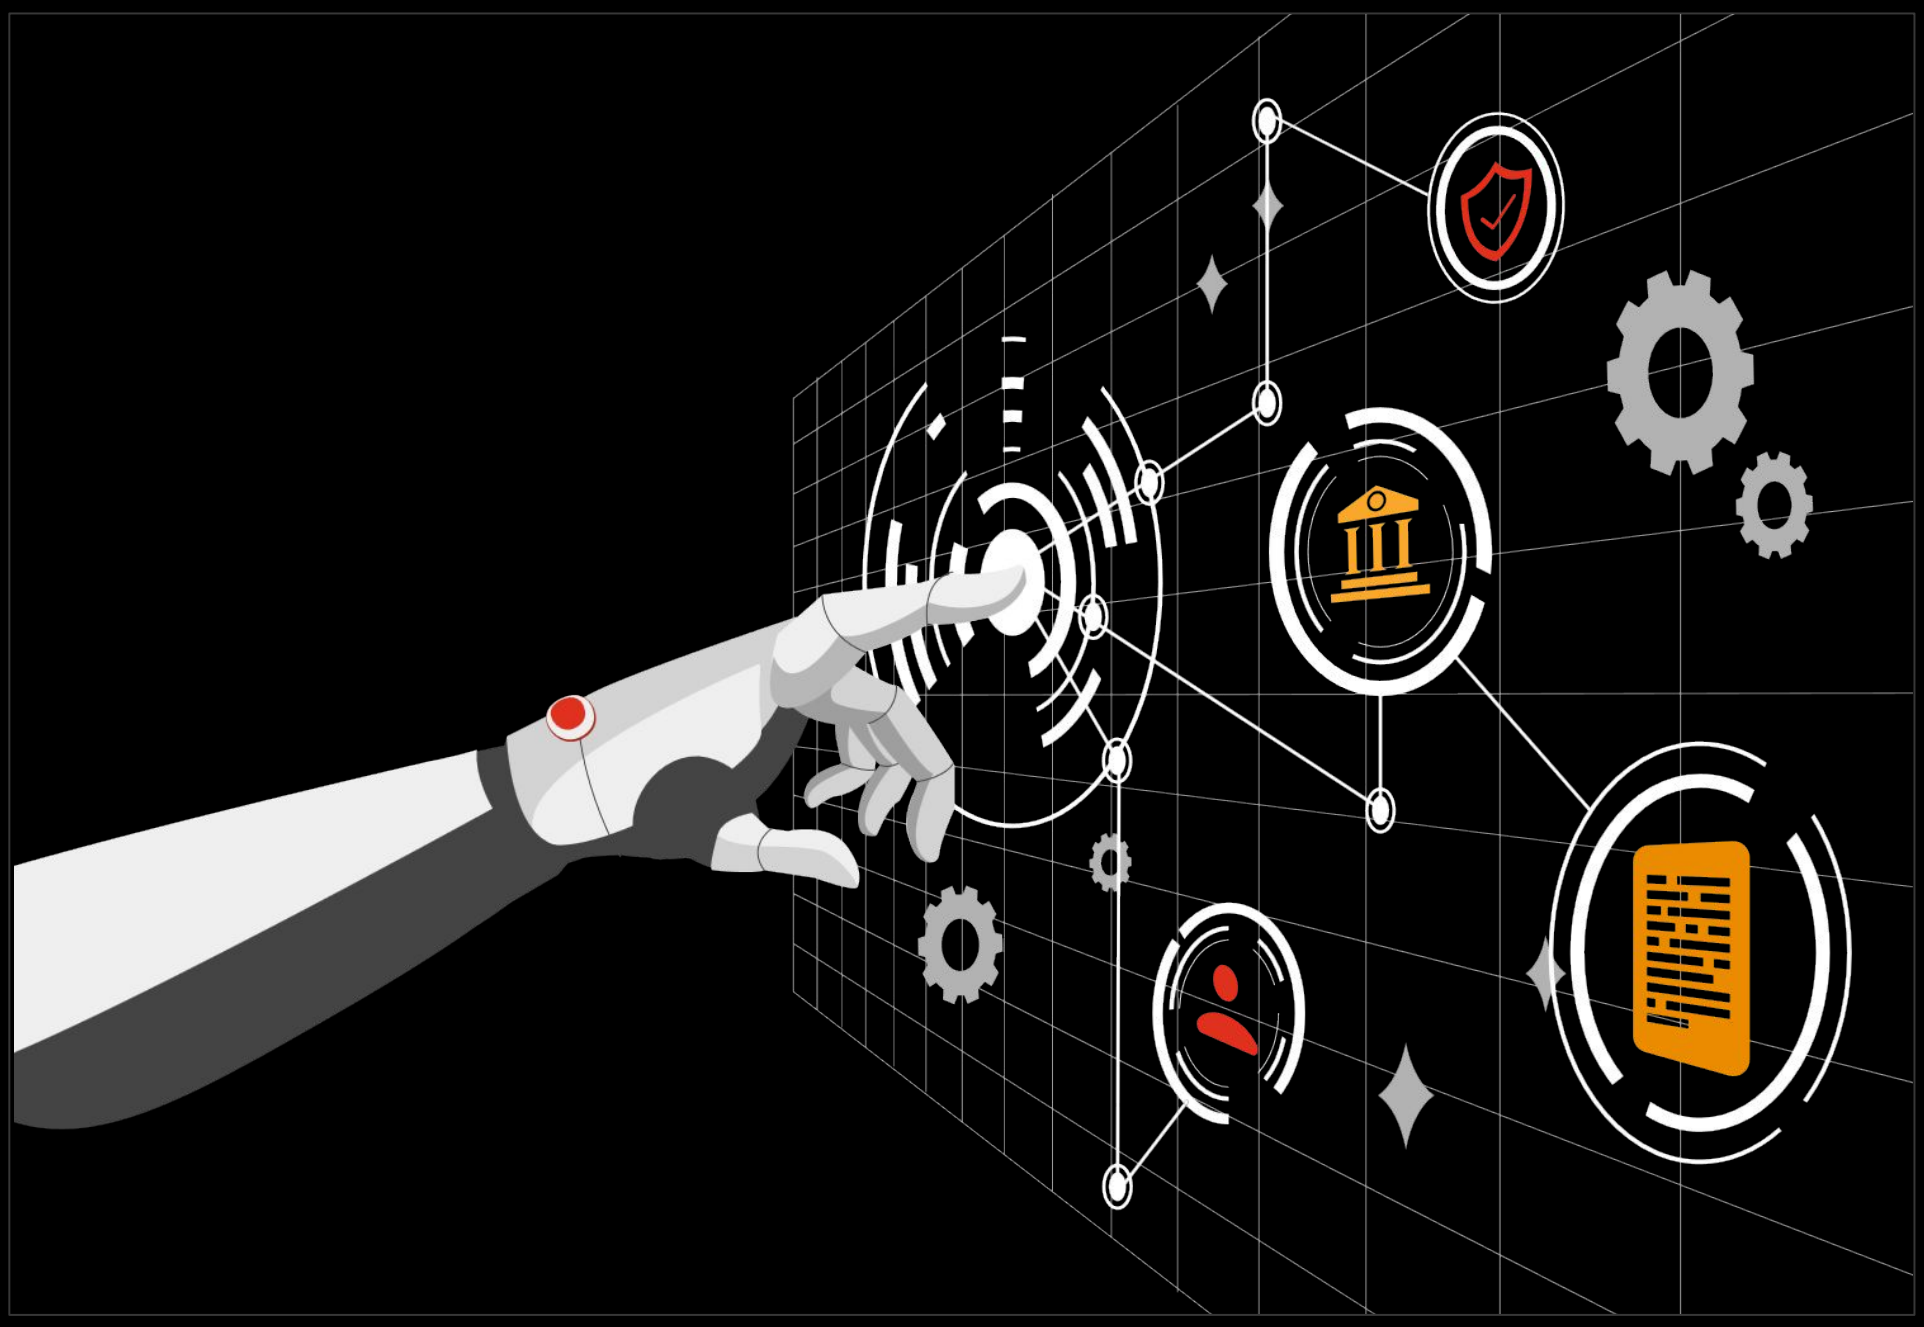

Real-world success stories

Boosting operational efficiency

Efficiency unleashed: Graph LLMs in action

Operational efficiency is the backbone of profitability. Graph LLMs streamline processes by automating complex tasks and providing predictive analytics.

Example:

BNP Paribas Personal Finance, a major player in retail financing in Europe, implemented a graph-powered fraud detection system[4]

Technology stack

Custom graph-powered model and internal data lakes

Financial impact

Reduced fraud losses by 20%

Non-financial benefits

Improved customer trust and regulatory compliance

Enhancing customer experience

Personalisation perfected: Transforming customer journeys with Graph LLMs

In a saturated market, customer experience is a key differentiator. Graph LLMs personalise interactions by understanding customer relationships and preferences.

Example:

Zurich Insurance uses knowledge graphs to enhance claims processing and improve risk insight[5].

Technology stack

Knowledge graph, advanced data analytics and framework.

Financial impact

Reduced claim processing time and improved operational efficiency.

Non-financial benefits

Enhanced fraud detection, faster claims resolution, and a more streamlined customer experience.

Driving revenue growth

Revenue reimagined: The strategic impact of Graph LLMs

Revenue growth stems from intelligent decision-making and market responsiveness. Graph LLMs provide insights that drive strategic initiatives.

Example:

Leading Financial software company Intuit, is using Graph LLMs for risk and fraud events detection and improving financial service operations[6].

Technology stack

proprietary LLMs, graph database, real-time analytics.

Financial impact

50% improvement in detecting fraud events.

Non-financial benefits

Better decision-making, and integrated data across systems.

Key GNN Tools: Open-source vs. commercial tools

A new generation of open-source and commercial GNNs are empowering Graph LLMs in the banking and insurance sectors:

Open-source GNNs

PyTorch geometric[7]

A library for deep learning on graphs within PyTorch.

A Features of PyTorch Geometric

GNN support

Enables easy implementation of GNNs for various tasks

Efficient data handling

Optimised for graph-structured data and supports sparse matrices and batch processing

Model library

Provides pre-built models for node and graph classification, link prediction, etc.

PyTorch integration

Fully integrates with PyTorch for GPU acceleration and deep learning workflows

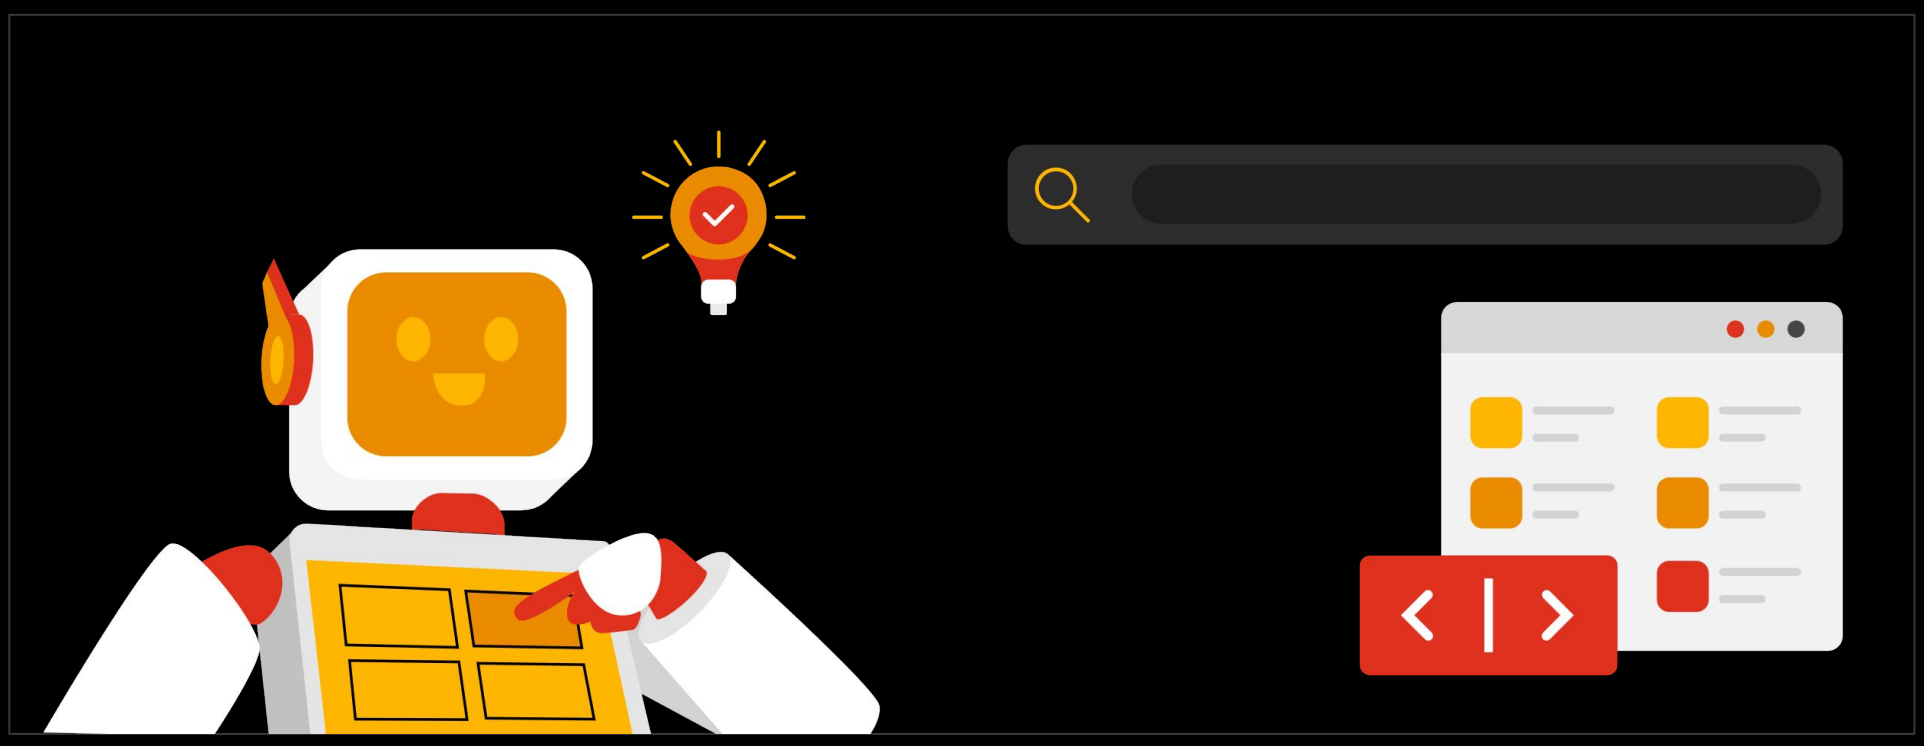

DGL (Deep Graph Library)[8]

A scalable, high-performance Python library for deep learning on graphs.

Features of DGL

GNN framework

DGL is a Python library designed for deep learning on graph-structured data, supporting various graph-based models

Scalability

Optimised for scalability, DGL efficiently handles large graphs and distributed computing

Flexible and extendable

Supports multiple backend frameworks like PyTorch, TensorFlow, and MXNet, allowing flexibility in model development

Wide range of applications

Suitable for applications such as node classification, link prediction, and graph classification

StellarGraph[9]

Machine Learning (ML) on graphs for link prediction, classification, and more.

Features of StellarGraph

Graph ML

StellarGraph provides tools for implementing machine learning on graph-structured data, supporting tasks like node classification, link prediction, and graph classification

Rich algorithms

The library includes a variety of graph neural network algorithms such as GCN, GraphSAGE, and GAT, tailored for different graph learning tasks

Scalable and efficient

It is designed to handle large-scale graph data, optimised for both performance and scalability

Integration with TensorFlow/Keras

StellarGraph is compatible with TensorFlow and Keras, making it easy to integrate with existing deep learning workflows

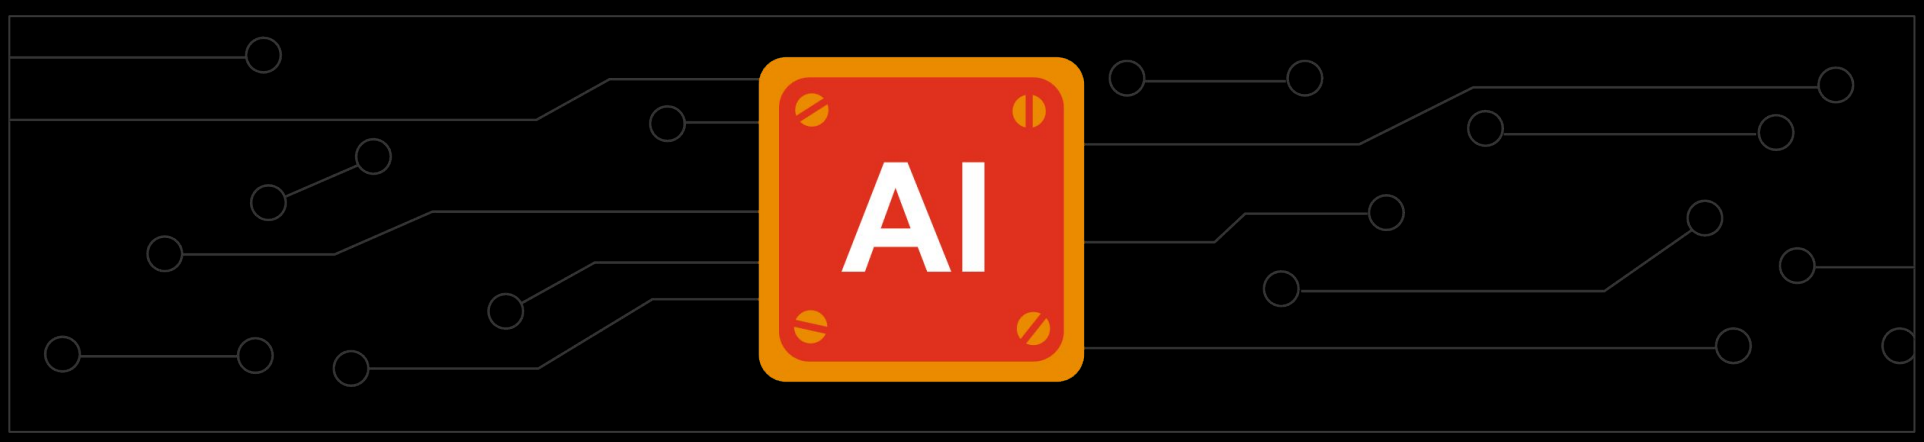

Commercial solution providers

Neo4j[10]

A platform for graph analytics and machine learning.

GenAI integration

Enhances graph analysis with generative AI capabilities

Advanced vector indexing

Improved support for high-dimensional vectors and similarity searches

Cloud integration

Expanded cloud object storage support for data handling and operations

Enhanced query flexibility

Offers more powerful querying features for better performance and flexibility in graph manipulation

TigerGraph[11]

A scalable graph database for enterprise applications.

Advanced scalability

TigerGraph 3.2 introduces cross-region replication and query workload management to support large-scale, mission-critical applications, ensuring smooth operations in both public and private cloud environments

Robust security

Provides advanced security features such as user-defined roles and fine-grained access controls, ensuring secure and compliant data management

Advanced querying and developer tools

Offers enhanced querying capabilities with flexible parameters, improved query language features, and faster batch query processing, empowering developers to create efficient graph-based solutions

Kubernetes integration

Full native Kubernetes support for efficient deployment, management, and scaling of graph-based applications, making it easier to integrate with modern cloud-native ecosystems

Microsoft Azure Cosmos DB with graph API[12]

A globally distributed, multi-model database service.

Global distribution

Cosmos DB offers seamless graph data management across multiple regions, ensuring high availability and low-latency access globally.

Gremlin query support

Leverages the Gremlin query language to perform graph traversals and analyse relationships efficiently within a distributed environment.

Elastic scalability

Automatically scales throughput and storage based on demand, making it suitable for both small and large applications.

Multi-model database

Integrates with other Cosmos DB APIs, such as SQL, MongoDB, and Cassandra, allowing users to handle different types of data within a single platform

Blueprint for success: Your guide to Graph LLM implementation

1

Identify use cases

Focus on areas with relational complexities – fraud detection, customer segmentation and risk assessment.

2

Choose the right tools

Select tools that align with your existing technology stack and scalability needs.

3

Prepare data

Ensure data quality and integrate siloed data sources to create comprehensive graphs.

4

Pilot and scale

Start with pilot projects to demonstrate value before scaling across the organisation.

5

Invest in talent and partnerships

Hire experts in graph theory and AI, or partner with consultants or tech firms specialising in Graph LLMs.

The essential Graph LLM checklist

1

Data quality and structure

Ensuring high-quality, well-structured data is crucial for the effectiveness of both the graph database and the LLM. The data must accurately represent the relationships and entities it models. These can be validated by business process owners. Banking and insurance organisations need to ensure the data meets regulatory requirements and customer risk assessments accurately.

2

Effective training

The LLM should be effectively trained or fine-tuned on domain-specific data to align with the graph database's schema and use cases. This ensures that the model can generate accurate and relevant queries and responses.

3

Integration depth

Deep integration between the LLM and the graph database is necessary for seamless operation. This includes the ability of the LLM to understand and generate Graph Query Languages (GQLs) based on natural language input.

4

Scalability

The system must be scalable to handle growing data volumes and increasingly complex queries without a significant drop in performance.

5

User experience

The interface should be user-friendly, allowing business users to interact with the graph database using natural language without needing to understand the underlying graph query language.

6

Performance monitoring

Continuous monitoring of the system's performance is essential to ensure that the integration delivers the expected speed, accuracy, and efficiency. Monitoring should focus on compliance check efficiency and data retrieval times, especially for high-stakes financial queries.

7

Security and privacy

Implementing robust security measures to protect sensitive data and ensuring compliance with privacy regulations is critical, especially when dealing with personal or proprietary information. Specific attention should be paid to financial data protection, customer privacy laws, and regulatory compliance. The EU General Data Protection Regulation (GDPR) and Payment Card Industry Data Security Standard (PCI DSS) are two well-known examples.

8

Customisation and flexibility

The ability to customise the LLM's responses and the graph database's behaviour to fit specific business needs and domains is a key factor for success.

9

Continuous improvement

The system should be designed to allow for continuous improvements and updates, both in the LLM’s capabilities and the graph database’s features.

10

Foster innovation, adapt, and evolve

Encourage a culture of innovation within your organisation by promoting rapid experimentation and cross-functional collaboration. Prepare your workforce to adapt your strategies and processes as the technology evolves and new opportunities arise.

The way forward

For the banking and insurance sectors, Graph LLMs far more than another technological advancement; they're offering a paradigm shift. They offer a rare trifecta: untapped operational efficiency, enhanced customer experience and significant revenue growth. By understanding and leveraging the complex relationships inherent in financial data, banks and insurance firms can unlock a new level of operational excellence, product innovations and customer satisfaction.

The impact of this technology goes beyond short-term gains; it establishes a foundation for long-term resilience and competitive advantage in an increasingly data-centric industry.

Those who embrace Graph LLMs have the potential to lead the industry, while those who don’t risk playing catch-up in a data-driven world.

In the high-stakes game of banking and insurance, Graph LLMs are the ace up your sleeve.

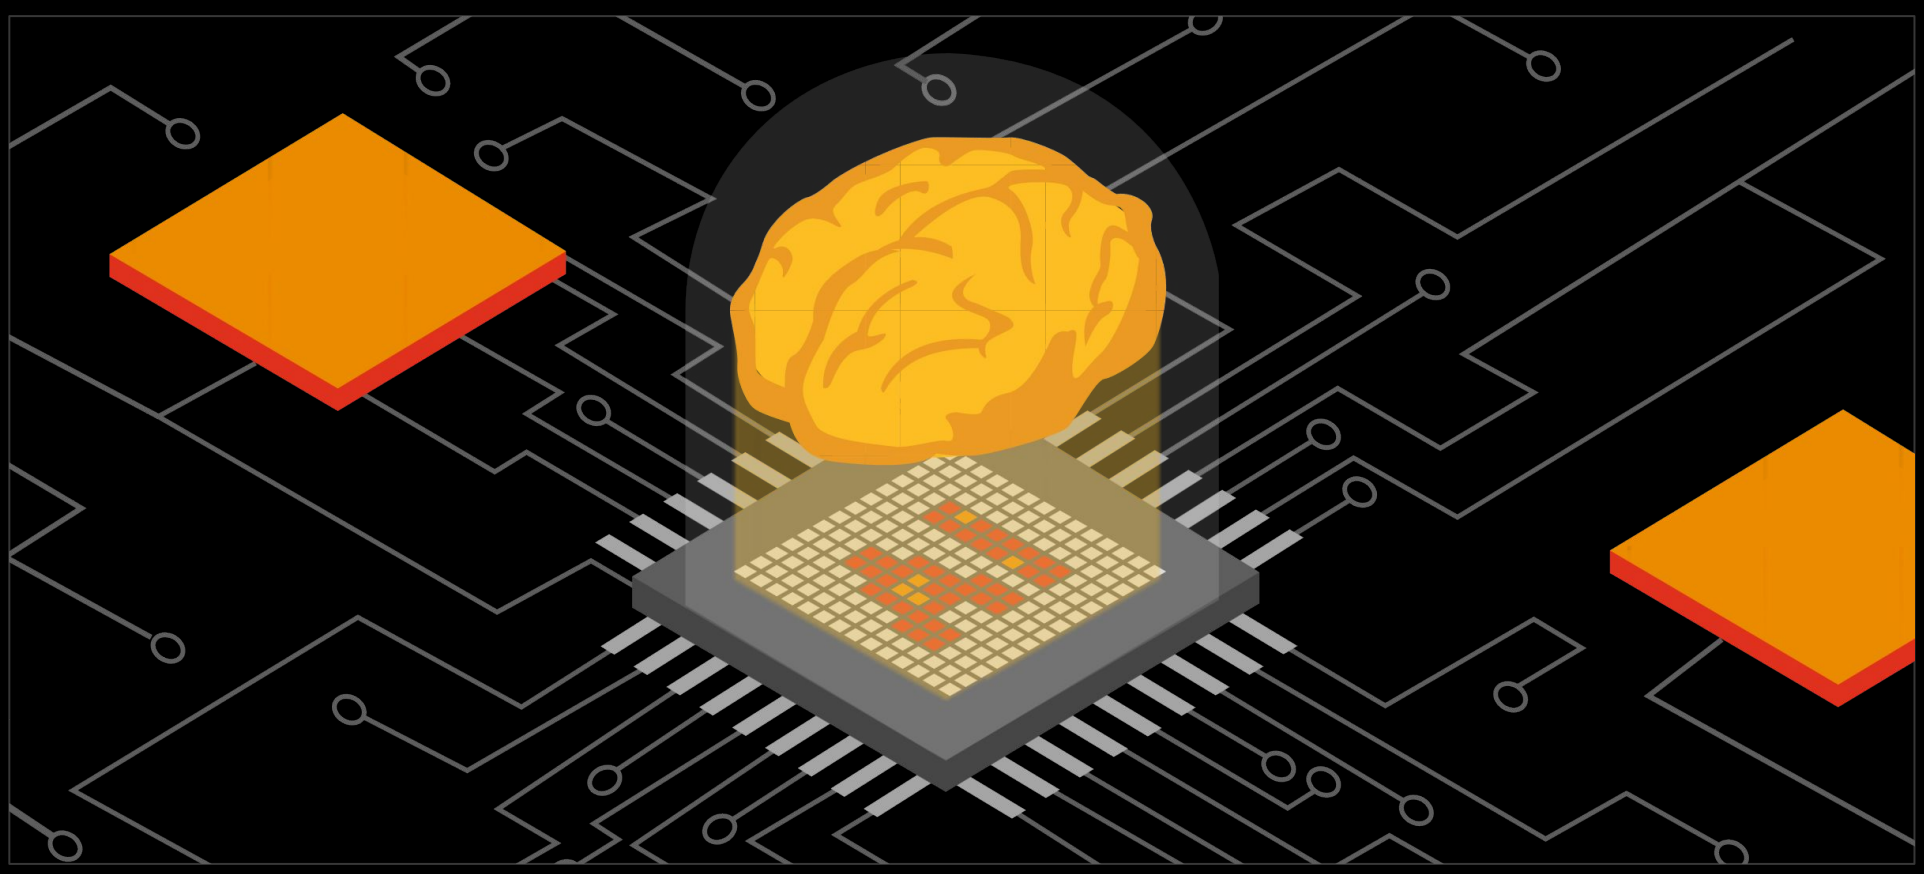

References

1. Generative AI Spending from Banking Industry to Grow by over 1400% by 2030, as Banks Seek to Scale AI to Revolutionise Business Models
2. Citi Publishes New Report on AI in Finance
3. Compare AI Software Spending in the Insurance Industry, 2023-2027
4. BNP Paribas Personal Finance - Graph Database & Analytics
5. Neo4j + Zurich Insurance Switzerland Case Study
6. Intuit - TigerGraph
7. PyTorch Geometric
8. GitHub - dmlc/dgl: Python package built to ease deep learning on graph, on top of existing DL frameworks.
9. StellarGraph - Machine Learning on Graphs
10. Neo4j
11. TigerGraph
12. Azure Cosmos DB - NoSQL and Relational Database

Get in touch

Akif Kamal

Partner
Technology Consulting
PwC Middle East
akif.kamal@pwc.com

Dr. Mohammad Tanvir Ansari

Director
Technology Consulting
PwC Middle East
mohammad.tanvir.ansari@pwc.com

Kaushal Chapaneri

Senior Associate
Technology Consulting
PwC Middle East
kaushal.chapaneri@pwc.com

About PwC

At PwC, our purpose is to build trust in society and solve important problems. We're a network of firms in 149 countries with more than 370,000 people who are committed to delivering quality in assurance, advisory and tax services. Find out more and tell us what matters to you by visiting us at www.pwc.com.

Established in the Middle East for over 40 years, PwC Middle East has 30 offices across 12 countries in the region with around 12,000 people. (www.pwc.com/me).

Disclaimer

The tools and technologies outlined in this report are not an exhaustive list. It's important for organisations to consider factors relevant to their own context before deciding to invest in Graph LLMs.