

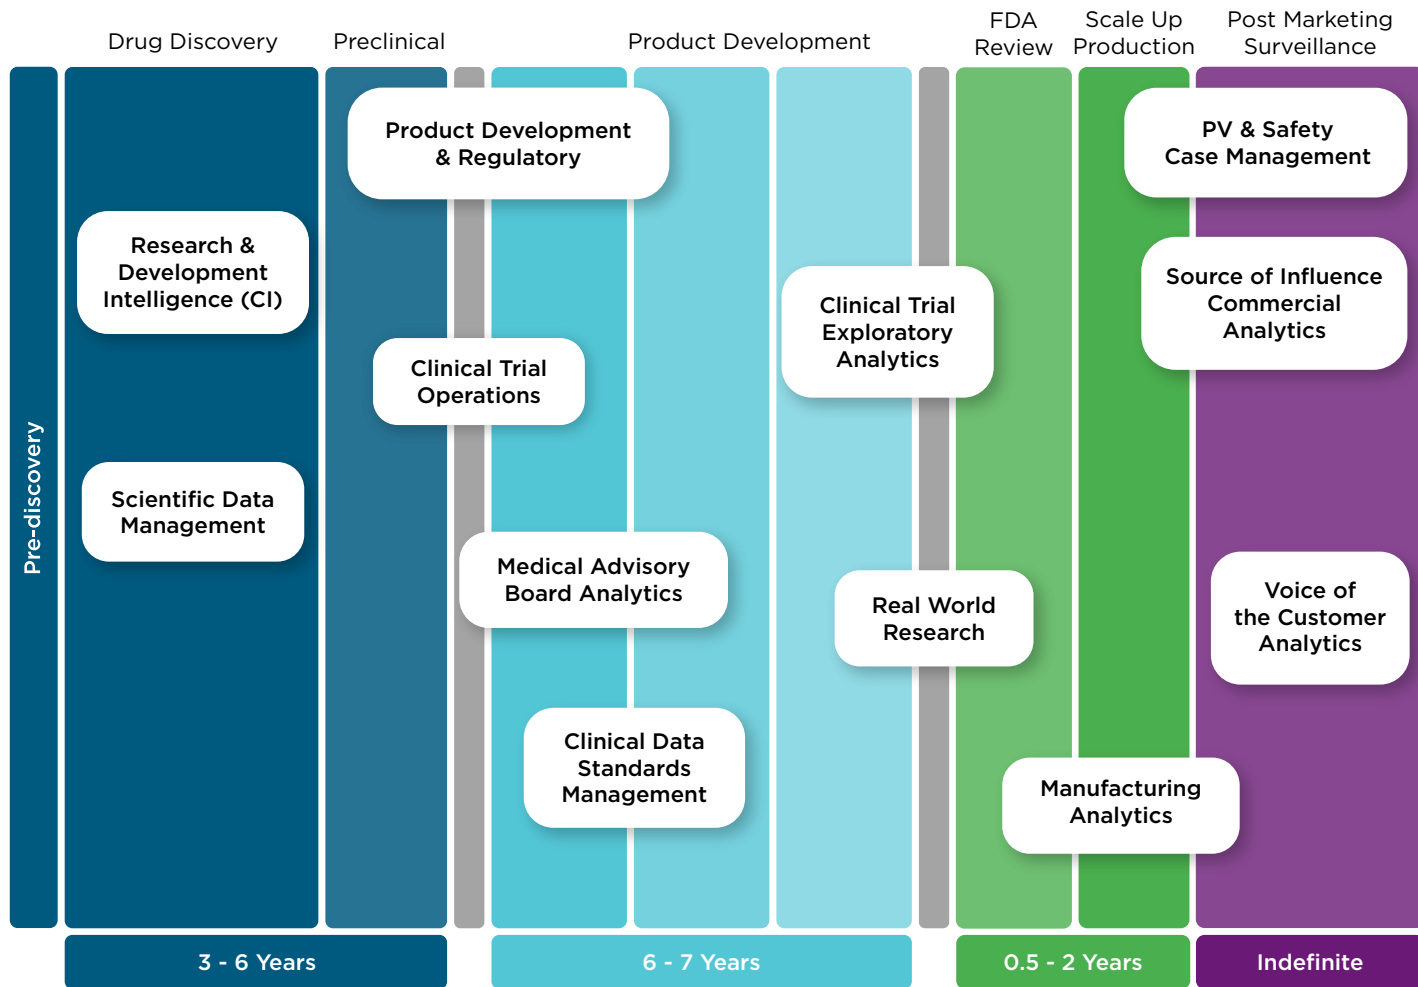
5 Use Cases of Knowledge Graphs in Drug Development & Manufacturing

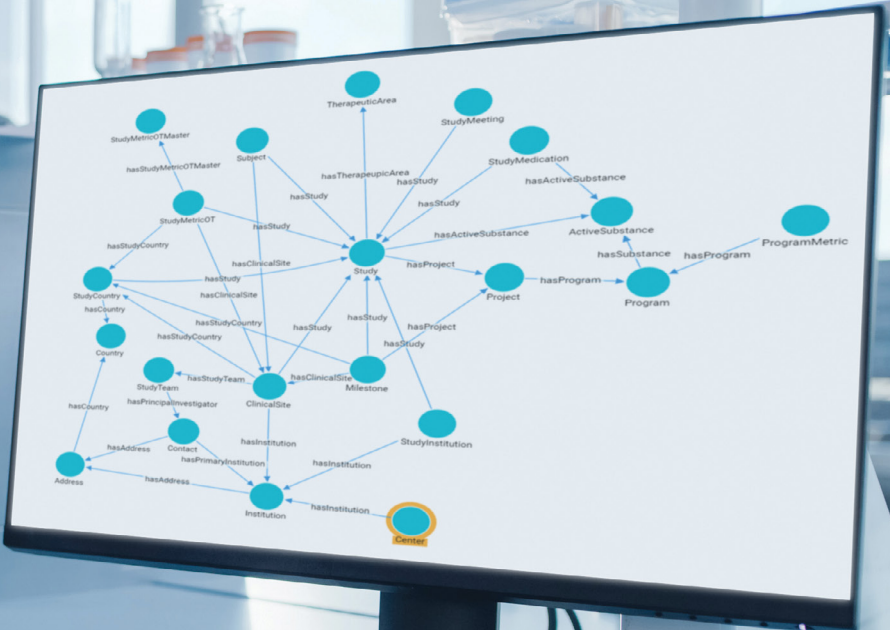
Introduction

In pharmaceutical development, data is the foundation of research and innovation. From executives to researchers and developers, everyone is motivated by the necessity of this progress, understanding its impact on the world. After all, billions of people rely on these drugs to change their lives in big ways and small.

If data is the foundation of innovation, then siloes are its chief obstruction. Data is stored in a variety of formats, both structured and unstructured alike. Accessing this data is difficult, using legacy systems with little integration and interoperability with open-source and commercial solutions. Existing tools lack scalability and flexibility; they offer no sustainability, and no opportunity to meet growing and changing business demands.

Knowledge graphs offer a bright alternative to existing structures for pharmaceutical development. In this booklet, we explore several use-cases across the drug development and manufacturing process to highlight the ways knowledge graphs transform data and derive real business value. From its semantics layer that turns raw data into actionable insight, to its connectivity and accessibility, knowledge graphs enable agile data integration. Companies can develop safer products faster and still meet stringent regulatory requirements.







Knowledge graphs in pharmaceutical drug development connect data silos and unravel complex relationships, turning fragmented information into cohesive insights that accelerate innovation and enhance precision in manufacturing. They complement existing infrastructure investments to make data AI-ready, paving the way for more efficient and informed decision-making.

Caroline Phares

Vice President of Healthcare and Life Sciences, Altair

Revolutionizing Cancer Research with Knowledge Graphs

The Challenge


Cancer research requires integrating diverse datasets—genomic sequences, treatment histories, and clinical trials—often stored in siloed formats across platforms like TCGA, GENIE, and MSK-IMPACT. These inconsistencies hinder data exploration, slow research, and complicate treatment evaluation.

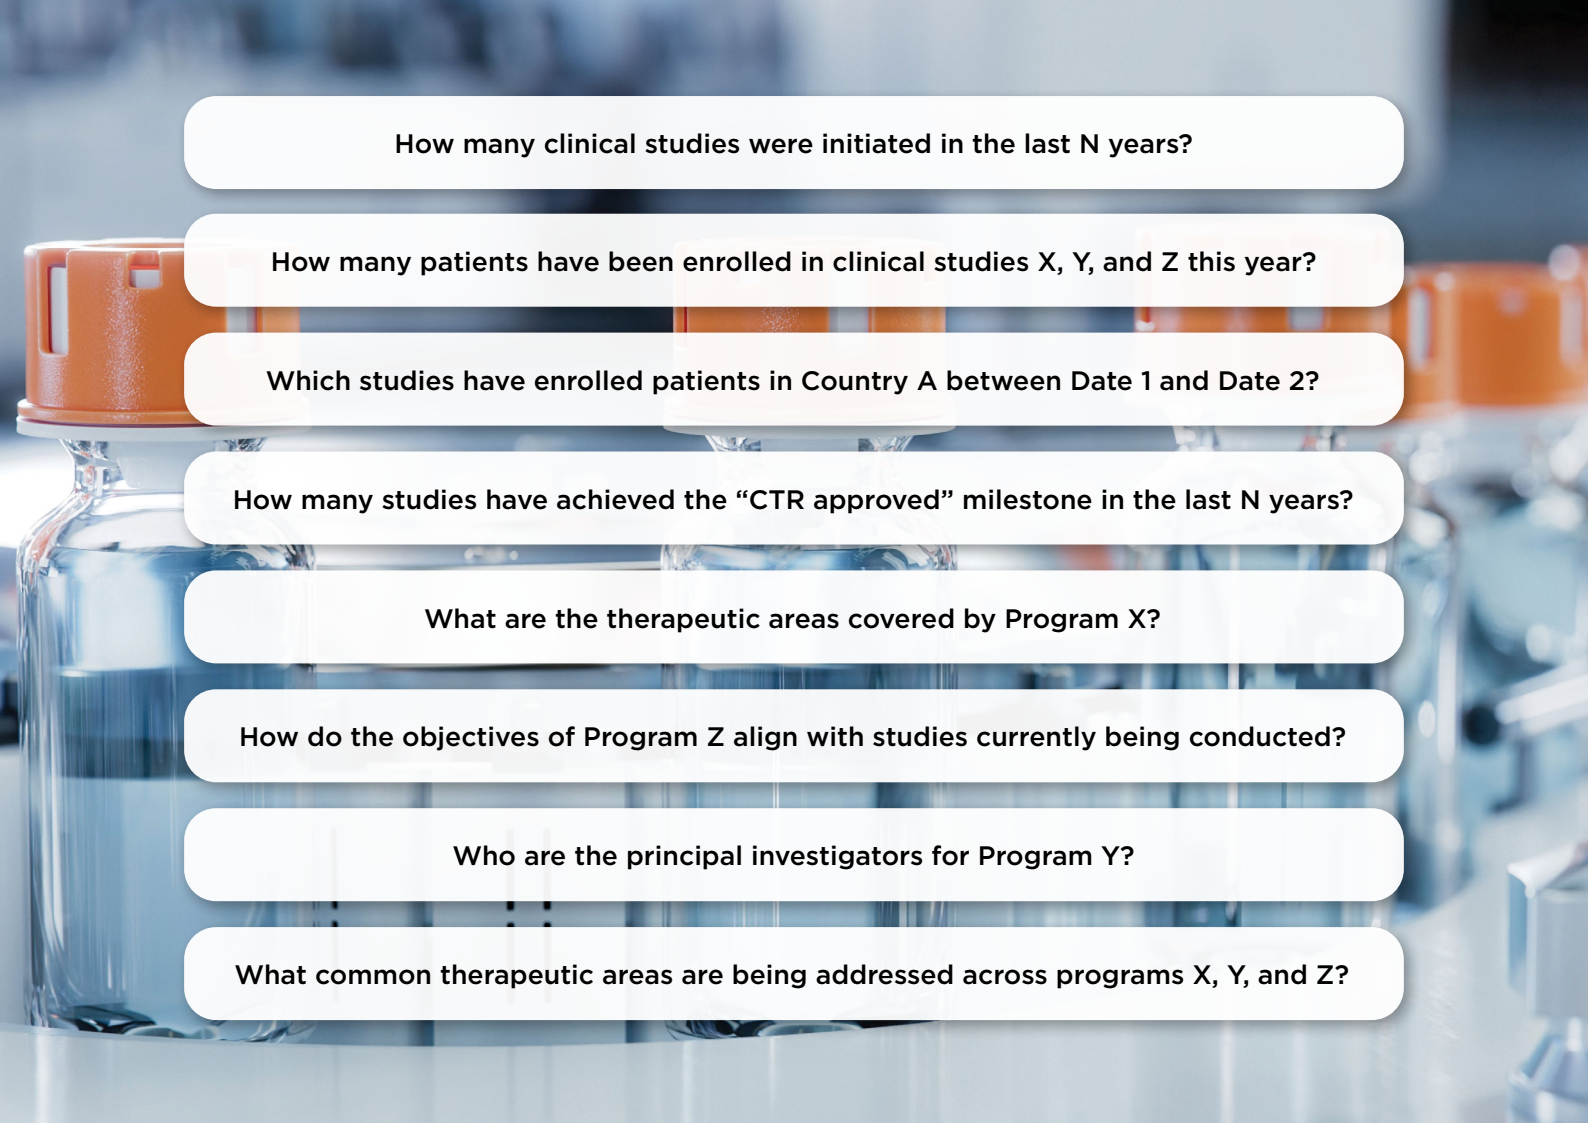
The Solution

Knowledge graphs unify disparate oncology data into a semantically enriched network. By mapping raw data to standardized ontologies, they offer a cohesive view of cancer-related information. Researchers can explore this graph visually or via tools like Jupyter Notebooks, enabling intuitive, context-rich analysis.

The Business Value

Knowledge graphs eliminate manual searches across disconnected systems, accelerating research and improving clinical decisions. They reduce time-to-discovery, enhance treatment insights, and support precision oncology with a scalable, future-ready architecture.



The background of the image shows a blurred laboratory setting with several glass vials or bottles. One vial in the foreground on the left has a prominent orange cap. The overall color palette is cool, with blues and greys, and the text boxes are white with rounded corners.

How many clinical studies were initiated in the last N years?

How many patients have been enrolled in clinical studies X, Y, and Z this year?

Which studies have enrolled patients in Country A between Date 1 and Date 2?

How many studies have achieved the “CTR approved” milestone in the last N years?

What are the therapeutic areas covered by Program X?

How do the objectives of Program Z align with studies currently being conducted?

Who are the principal investigators for Program Y?

What common therapeutic areas are being addressed across programs X, Y, and Z?

Turning Clinical Trial Data into Actionable Insights

The Challenge

Clinical trial operations data is scattered across databases, spreadsheets, and unstructured systems. Answering basic questions often requires technical expertise and manual effort, slowing down decisions and creating inefficiencies.

The Solution

Knowledge graphs bring order to this complexity by connecting data into a unified, semantic layer. With natural language ontologies and integration across data types, users can ask questions in plain language and receive accurate, contextual answers. Combined with large language models, this enables powerful, interactive analytics without the risk of AI hallucinations.

The Business Value

Knowledge graphs cut through data silos, enabling faster insights, reducing IT workload, and empowering teams to make informed decisions without writing code. They improve data quality, support compliance, and scale across use cases—while keeping sensitive data secure.

Connecting Patient Safety Data While Preserving Anonymity

The Challenge


Pharmacovigilance draws on diverse sources—clinical trials, drug labels, narratives, and global affiliates—often in inconsistent formats and systems. Traditional tools can't easily harmonize this data, and safety experts spend more time wrangling it than analyzing it. Complicating matters further, some data must remain siloed to protect patient anonymity and meet regulatory standards.

The Solution

Knowledge graphs integrate structured and unstructured data while respecting data boundaries. By modeling relationships semantically, they support advanced signal detection, risk analysis, and reporting—even when sensitive data must stay partially separate.

The Business Value

Knowledge graphs reduce manual effort, accelerate analysis, and empower safety teams with self-service tools. This results in better data quality, faster reporting, improved compliance, and ultimately safer, smarter decisions.








From Trials to Distribution: The Full Picture

The need for accessible, integrated data does not disappear once products move to manufacturing, but shifts. Pharmaceutical manufacturing brings its own datasets, like recipes, storage, supply chain management, regulatory compliance, impact analysis, adverse effects records, patient records, substance tracking, and more. It is vital to operational efficiency that teams across the organization have access to this data to make informed, data-driven decisions.

Knowledge graphs empower teams to discover new, deeper connections from pre- to post-manufacturing data. This solution bridges the gap, offering an agile, sustainable solution that grows with the business, providing a holistic view of enterprise data.



A Fully Traceable Batch Genealogy

The Challenge


Pharmaceutical companies often rely on outdated systems to store data, but they lack critical integration with one another, leading to siloed data. Current solutions are often unsustainable and unscalable, leading to domain experts' time spent tracking data down.

The Solution

Knowledge graphs solve data accessibility challenges by integrating data from a variety of sources, both structured and unstructured alike. Using this technology, pharmaceutical manufacturers can transform existing outdated systems to something scalable, sustainable and interoperable with open-source technologies to meet every business demand.

The Business Value

Knowledge graphs transform batch genealogy, offering a collaborative, multiuse solution to trace a batch from all angles. Pharmaceutical companies can meet regulatory standards and audit requirements more easily. Data becomes connected under a single source of truth for a better understanding of manufacturing processes.





Standardizing Access with IDMP Adoption

The Challenge

Pharmaceutical companies struggle with Identification of Medicinal Products (IDMP) adoption because of the complexity, quality, and quantity of enterprise data. Additionally, there are little to no implementation standards guiding IDMP adoption – coupled with a shortage of technical talent to oversee it in the industry. Without such standards, companies could struggle with impact analysis, tracking harmful substances, and pharmacovigilance.

The Solution

Knowledge graphs provide a detailed and integrated look into enterprise data across the drug development process. Analysts can use knowledge graphs to pinpoint manufacturing shortages based on the information gathered at the earliest sign, or identify which products contain harmful substances across the supply chain. Researchers can deploy knowledge graphs to break down silos storing data on drug adverse effects and identify possible interactions earlier.

The Business Value

Knowledge graphs make drug data accessible, offering deep connectivity of an organization's data, leading to safer products on the market.

Conclusion

Data does not have to exist in isolation, hindering innovation and slowing progress for pharmaceutical manufacturers. This is a small selection of existing use-cases that grows every day as more teams adopt knowledge graphs.

Across the drug development and manufacturing process, knowledge graphs are transforming data – and how teams access it. From researchers and developers at the cutting edge, to executives overlooking entire businesses, knowledge graphs provide a new way to extract insight and confidently make data-driven decisions. Companies can use this technology to grow with the business to meet changing demands and regulatory requirements. Knowledge graphs empower researchers to discover complex relationships between data points; and at the highest levels, executives can make data-driven decisions to steer their organizations.

Teams can derive real business value and actionable insight using knowledge graphs at every step in pharmaceutical development.





Learn more: altair.com/knowledge-graphs

Changing tomorrow, together.



See why Altair was recognized as
A Leader in the 2025 Gartner Magic
Quadrant for DSML platforms.

