Case Study: Entity-Event Knowledge Graph for Powering AI Solutions (Montefiore)

PEER & PRACTITIONER RESEARCH Published 12 March 2021 - ID G00744262 - 6 min read

Chief Data and Analytics Officer Research Team

Initiatives: Artificial Intelligence

AI solutions are often hindered by fragmented data and siloed point solutions. Montefiore's data and analytics leader used semantic knowledge graphs to power its AI solutions and achieved considerable cost savings as well as improvements in timeliness and the prediction accuracy of AI models.

Montefiore HEALTH SYSTEM, INC.

- Company Name: Montefiore
- Industry: Healthcare
- Headquarters Location: New York City
- Revenue: $6.3 Billion (2019)
- Employees: 34,082 (2019)

Overview

Cost-effective, just-in-time AI solutions have been extremely difficult to build because of a lack of interoperability among thousands of data and applications systems. Typical solutions, such as data marts and siloed point solutions with third-party vendors that use multiple data models, worsen the already fragmented data and analytics landscape in modern enterprises.

The data and analytics leader at Montefiore enabled the development of the company's advanced analytics and AI applications by creating a knowledge graph that provides an integrated view of its data from various business applications and source systems. Montefiore's approach to data modeling using an entity-event knowledge graph has increased the accuracy and timeliness of its predictions for diseases such as acute respiratory distress syndrome from COVID-19 and led to lower costs, better care and more lives saved.

Solution Highlights

- Montefiore uses an entity-event knowledge graph to model enterprise data.
- Algorithms and knowledge bases enable Montefiore to build and continuously update its knowledge graph.

Challenge

Although enterprise data has enormous potential to unlock value for business, it is hard to tell what data has what value without understanding the context of the business problem at hand. There is often a large value gap between enterprise data and concrete business problems the organization must solve using advanced analytics and AI (see Figure 1). Data and analytics leaders struggle to close this gap in a way that allows them to apply enterprise data for concrete business problems without creating analytical silos and point solutions.

Figure 1. The Value Gap Between Data and Business Problems

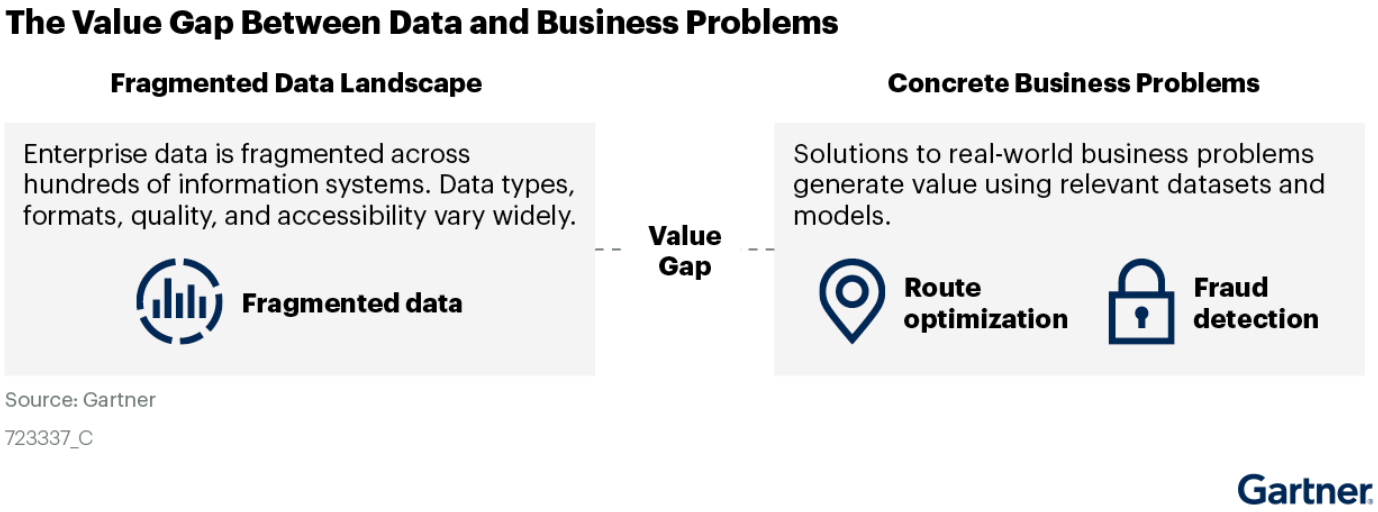

Many data and analytics leaders continue to ask what they can do to organize their data to power a variety of advanced analytics and AI applications in a repeatable and reusable way. Businesses operate in functional silos, which results in a fragmented enterprise data landscape that is not conducive to that end. And applying point solutions and siloed data marts bought from vendors do not help because these tools rarely achieve an acceptable return on investment (see Figure 2).

Figure 2. The Cost of Data Lakes, Warehouses and Analytics Siloes to the Enterprise

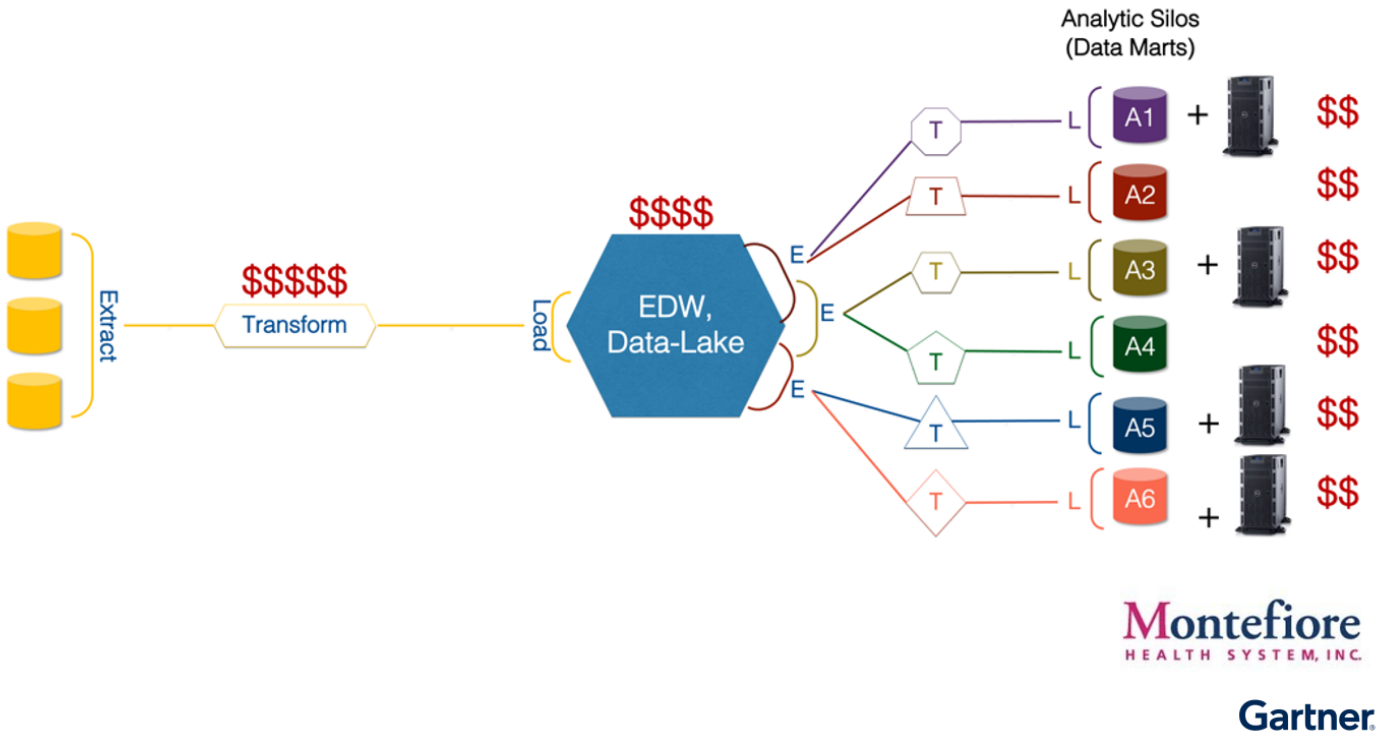

Source: Montefiore

Business Context

Montefiore has more than 3,000 hospital beds in 10 hospitals in the Bronx, Westchester and the Hudson Valley in New York. The data and analytics leader, Dr. Parsa Mirhaji, leads 23 FTEs at the Center for Health Data Innovations at Montefiore. The team has already built multiple AI applications that are currently in production, with more AI models in the pipeline.

Solution Overview

To prepare data for AI, Montefiore created a data and analytics platform called the Patient-Centered Analytic Learning Machine (PALM). The platform provides a unified approach to organizing data for advanced analytics solutions and AI applications in the enterprise (see Figure 3).

Figure 3. Montefiore's Patient-Centered Analytic Learning Machine

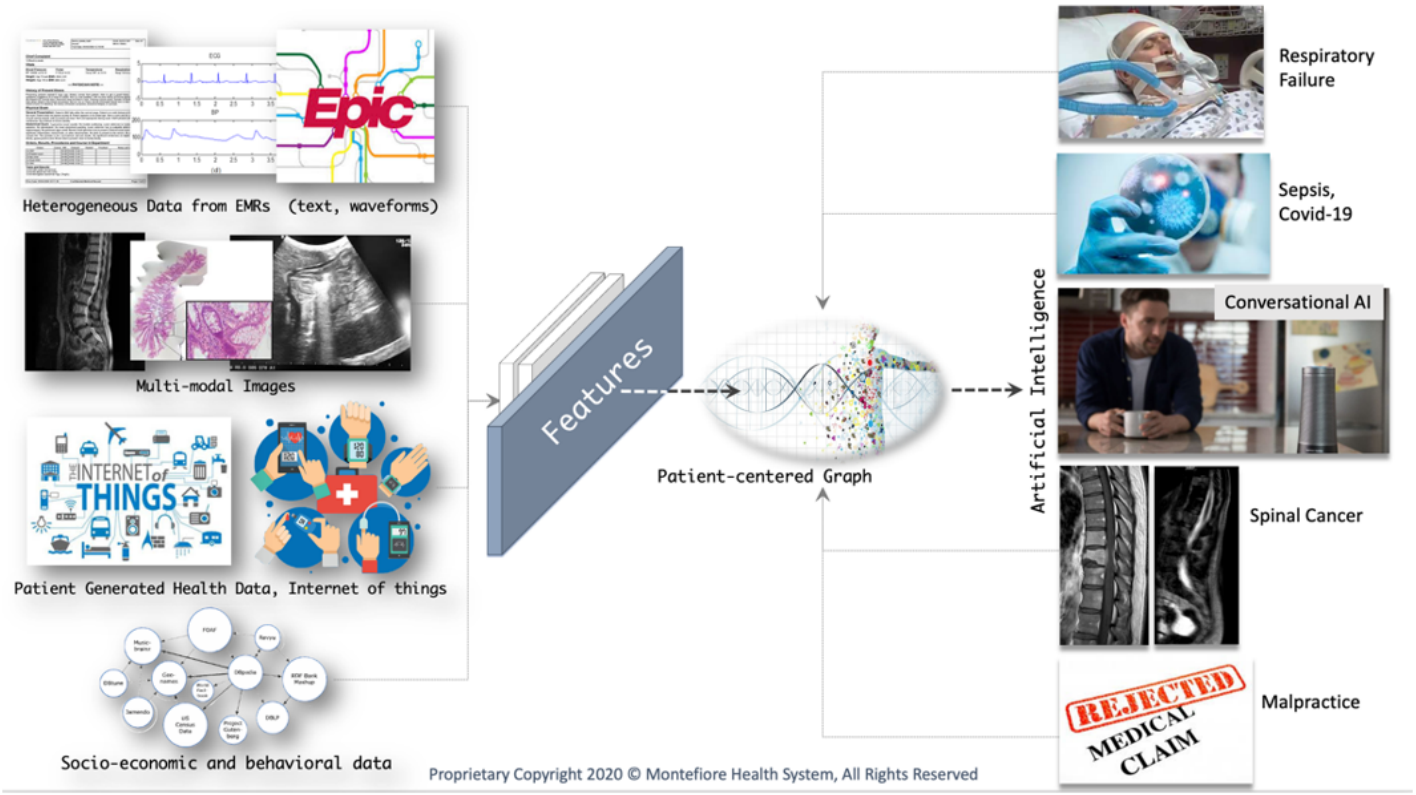

Source: Montefiore

Montefiore's PALM uses the following components:

- Knowledge graph with integrated taxonomies and ontologies
- Visual modeling and rules layer
- Inference engine
- Graph database
- Relational-to-graph conversion ETL
- Integration with Spark and machine learning libraries
- Continuous integration and feedback processing

Entity-Event Knowledge Graph Data for AI

Montefiore organizes its data with an entity-event schema that can be used and reused for many solutions.

To address the issue of a fragmented data landscape, Montefiore adopted an open-standard knowledge graph approach to organizing data for the building of advanced analytics and AI solutions. The reason for this choice is that knowledge graphs provide a composable data model that can be flexibly built and extended for many use cases. All the relevant source data for 10 hospitals were organized in an efficient graph schema called an entity-event data model. In this schema, a patient is an entity, and each element of care or experience the patient receives is an event.

The building blocks of such a database are triples such as “Patient X — has received — Medication A” or “Patient Y — was diagnosed with — Disease B.” This format can be extended to many more entity types, such as doctors, drugs and hospitals. Triples can have attributes that could be used to cater to different usage contexts. The model also incorporates temporal elements to enable users to see how things change over time (see Figure 4).

Figure 4. Entity-Event Knowledge Graph

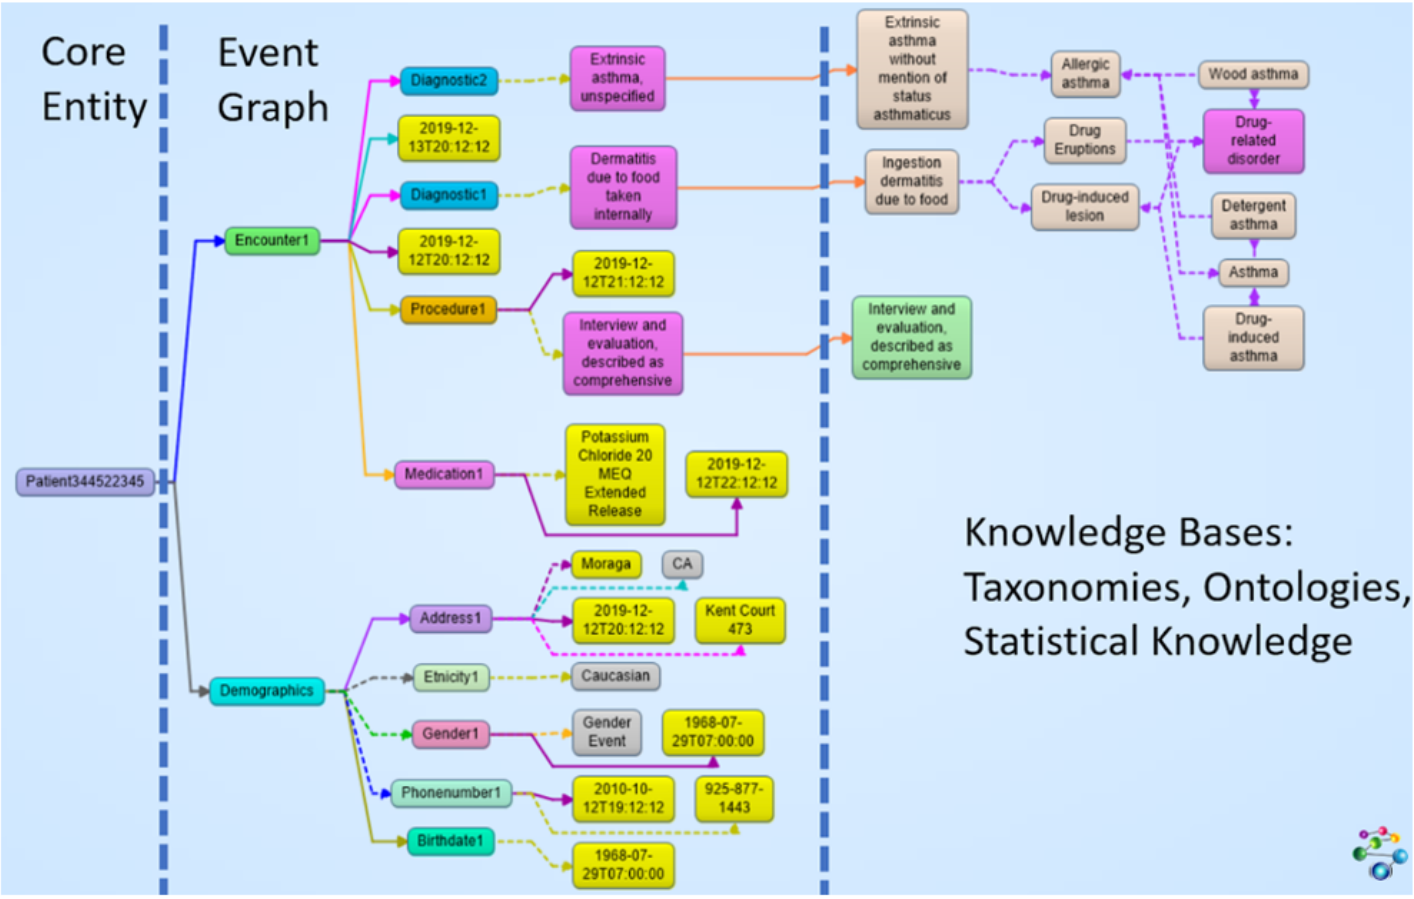

Montefiore HEALTH SYSTEM, INC.

Gartner

Source: Montefiore

In the case of healthcare, a single symptomatic observation — such as a biometric reading or an interpretation from a care provider — can be a discrete data point. Regardless of its specificity or interpretive nature, that same observation can be attributed to dozens of disease management and treatment scenarios. Each of those potential combinations is important, so Montefiore logs each as a triple in the knowledge graph. This enables users to find all possible connections in the knowledge graph. The same applies to expert care providers with specialties and more.

Algorithms and Knowledge Bases to Build the Knowledge Graph

Montefiore builds and continuously updates its knowledge graph using algorithms and business ontologies.

While identifying the right schema for organizing the integrated data is a foundational step, building a knowledge graph requires additional work. The data and analytics team at Montefiore take the following three steps to build its knowledge graph for powering its AI solutions:

1. Extract relevant source data and load it into the knowledge graph using the source metadata and the knowledge bases (taxonomies and ontologies) the data and analytics team has developed over time.
2. Transform knowledge graph data on demand to generate features for just-in-time analytics and AI applications.
3. Extract analytics results and byproducts from AI applications and load them into the knowledge graph. Every part of the workflow used to build a model (for example, who did it, what data was used and when it was run) gets stored back into the same graph as triples to allow for learning from feedback (see Figure 5).

Figure 5. Building the Patient-Centered Analytics Learning Machine

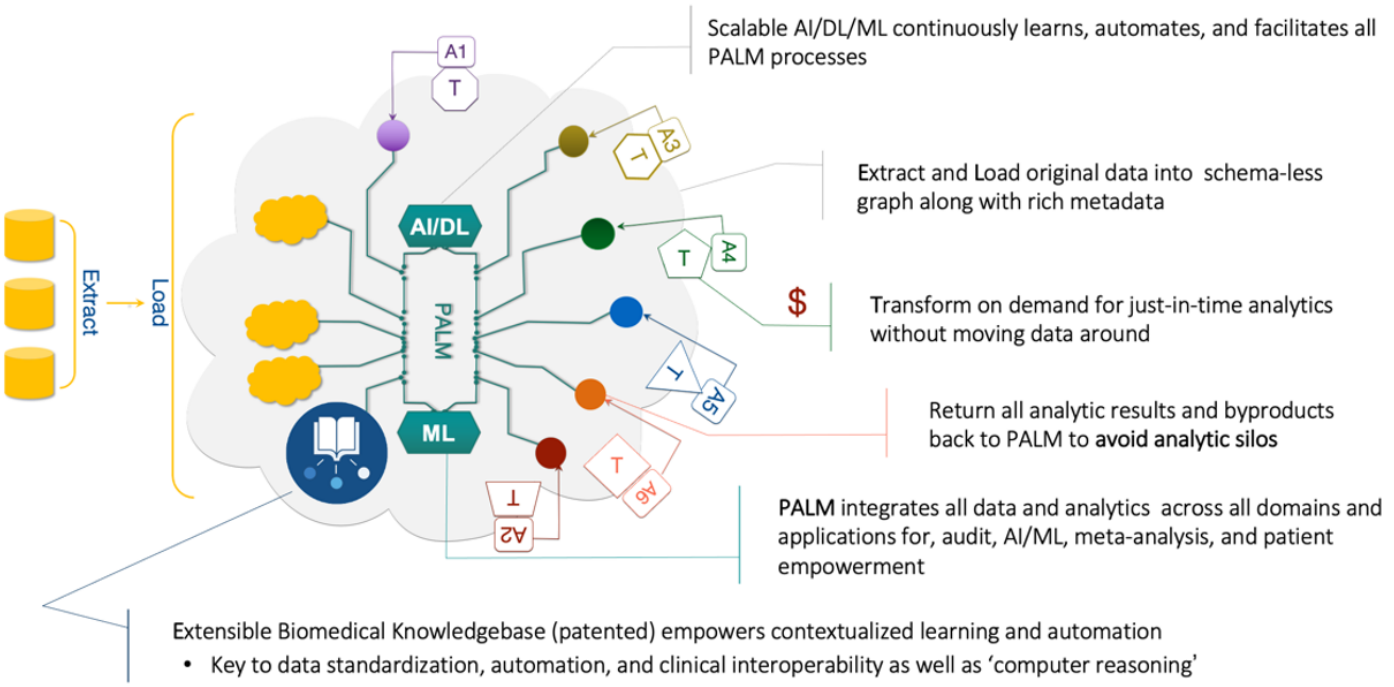

Source: Montefiore

Results

One benefit of the PALM is the dramatically improved accuracy and timeliness of predictions for diseases such as acute respiratory distress syndrome from COVID-19. Earlier and more accurate predictions lead, in turn, to lower costs, better care and more lives saved (see Figure 6).

Figure 6. The Benefits of the Patient-Centered Analytic Learning Machine

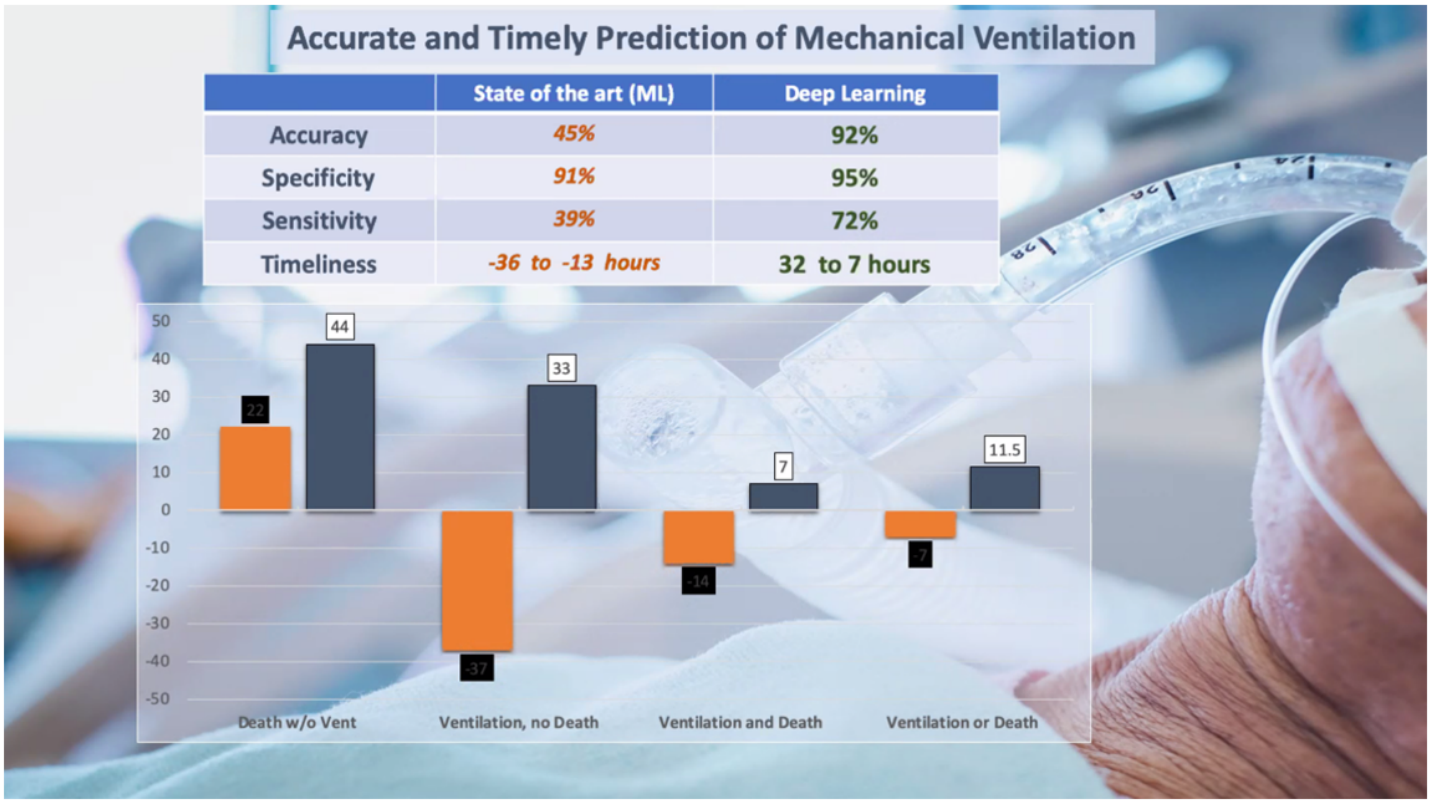

Source: Montefiore

The platform is currently being used to develop new models for finding similar patients, detecting alternative treatment pathways and constructing a 360-degree view of the patient.

About This Research

We developed this case study to describe Montefiore's approach to creating a unified knowledge graph to make complex data and analytics systems interoperable. The case study is based on interviews with Dr. Parsa Mirhaji, director of the Center for Health Data Innovations at Montefiore and Albert Einstein College of Medicine, and his collaborator Jans Aasman, CEO of Franz Inc.

Recommended by the Authors

Case Study: Data and Analytics Monetization With Knowledge Graphs and AI (Turku City Data)

Case Study: Answering Critical Business Questions With Graph Analytics (Jaguar Land Rover)

Graph Steps Onto the Main Stage of Data and Analytics: A Gartner Trend Insight Report

Infographic: An Intelligent Composable Business Demands a Data Fabric

Creating the Composable Healthcare Organization for Healthcare and Life Science CIOs

Tool: Healthcare and Life Science CIOs Executive Presentation for Composable Data and Analytics

Tool: Healthcare Provider CIO Executive Presentation for a Composable Digital Health Initiative

---

Is this content helpful to you?

YES NO

Recommended For You

Use Stay Conversations to Improve Engagement and Retention

Summary Translation + Localization: Cool Vendors in Autonomous Vehicle Systems

Market Insight: How Sensors Drive New Interactions in the Future Connected Home

Toolkit: Measuring Internal Satisfaction With Your Supply Chain COE

Summary Translation: Ignition Guide to Creating a Digital Twin MVP