NOVEMBER 17, 2025

www.orbitai.com.mx

WELCOME TO ORBITAI

NEWSLETTER

Artificial Intelligence Is Rewriting the Grammar of Medicine: The Rise of Delphi-2M

Artificial Intelligence is redefining the “grammar” of modern medicine. A recent Nature publication introduces Delphi-2M, a model that goes far beyond predicting isolated diseases—it models the entire natural history of human health. This is more than a predictive tool: it is a generative model capable of learning how health conditions evolve, compete, and interact across a lifetime.
At OrbitAI, we explore why Delphi-2M represents a prototype of a true foundation model for the future of healthcare.

Read More on our LinkedIn profile

- **Precision Engineering: A New Architecture of Time**

The research team adapted a GPT-2–style architecture to handle the continuous and complex nature of clinical data. Its core innovations revolve around time:
State Tokens
The model operates with more than 1,200 tokens representing high-level diagnoses (ICD-10) and key health factors such as BMI or smoking status.
Continuous Age Encoding
Instead of discretizing time, Delphi-2M uses continuous sine-cosine encodings to capture clinical chronology with exceptional precision.
Dual Output Heads
One head predicts what health event will occur; the other predicts when it will occur.
This dual structure enables actionable long-term risk anticipation—something traditional models struggle to achieve.

- **The Challenge: Regulated but Fragmented Data**

Validated on a cohort of 1.9 million individuals, Delphi-2M sets a new benchmark for long-range clinical forecasting.

State Tokens
The model operates with more than 1,200 tokens representing high-level diagnoses (ICD-10) and key health factors such as BMI or smoking status.

Continuous Age Encoding
Instead of discretizing time, Delphi-2M uses continuous sine-cosine encodings to capture clinical chronology with exceptional precision.

Dual Output Heads
One head predicts what health event will occur; the other predicts when it will occur.
This dual structure enables actionable long-term risk anticipation—something traditional models struggle to achieve.

Statistical Power and Long-Horizon Prediction

Validated on a cohort of 1.9 million individuals, Delphi-2M sets a new benchmark for long-range clinical forecasting.

- Strong overall performance: Predicts the next health event with an average AUC near 0.77.
- Long-term stability: Maintains high accuracy even when forecasting up to 10 years into the future, with only a slight performance drop.
- Critical vulnerability detection: Identifies high-risk trajectories with exceptional precision.

This level of statistical performance allows clinicians to estimate disease burden and understand how health conditions unfold over time in ways previously impossible.

Generative AI and Privacy by Design

One of Delphi-2M's most groundbreaking capabilities is its ability to generate synthetic health trajectories.

Synthetic Clinical Data
Just as an LLM generates text, Delphi-2M creates realistic medical sequences without exposing any real patient data.

High Utility
A secondary model trained only on synthetic data preserved strong performance—demonstrating that privacy and innovation can truly coexist.

Read More on our LinkedIn profile

LLMs in Aviation

The Next Frontier: LLMs in Aviation

The rise of Large Language Models (LLMs) marks the next stage of evolution. These models go beyond text processing—they interpret technical manuals, incident reports, maintenance logs, and safety protocols. Key applications include:

- Quality inspection: automated analysis of reports and images of structural defects.
- Predictive maintenance: correlating historical failures with sensor patterns.
- Pilot training: generating adaptive simulations tailored to real performance profiles.
- Regulatory compliance: intelligent summarization of ICAO, EASA, and FAA regulations.

Yet this power comes with new risks: misinterpretation, bias propagation, data exposure, and cybersecurity vulnerabilities.

The EASA (2024) has already emphasized that LLMs must remain "human-supervised" in critical systems—following the human-in-the-loop principle.

According to Reuters (2024), 83% of Chinese enterprises have adopted generative AI, compared to a global average of 54%, highlighting both the rapid adoption curve and the geopolitical race for technological leadership.

AI, Sustainability, and Air Safety

AI not only optimizes operations—it is also accelerating the decarbonization of air transport.

Boeing and Rolls-Royce, for instance, use AI to model fuel efficiency and engine wear, achieving up to 20% reductions in kerosene consumption through real-time optimization (Boeing Sustainability Report, 2024).

Additionally, the integration of AI into control towers and dynamic flight routing has reduced air congestion, potentially cutting 43 million tons of CO_2 emissions annually if implemented at scale (IATA, 2025).

Safety also benefits: predictive maintenance algorithms detect micro-fractures before they become critical, and AI-assisted vision systems enhance landing precision in low-visibility conditions.

From Theory to Trust: Testing AI in Production

AI is neither magic nor marketing. Its value arises only when deployed in real environments under human oversight and governance.

Organizations that implement robust ontologies, data-audit mechanisms, and verified LLMs move beyond experimentation to measurable operational impact. As Pascal TEA summarizes:

"The success of AI in aviation does not depend on having more data—it depends on data that truly means something."

Perhaps more than any other sector, aviation proves that trust, traceability, and production-grade testing are the true engines of artificial intelligence. And on this flight toward the future, ontology is the runway.

Read More on our LinkedIn profile

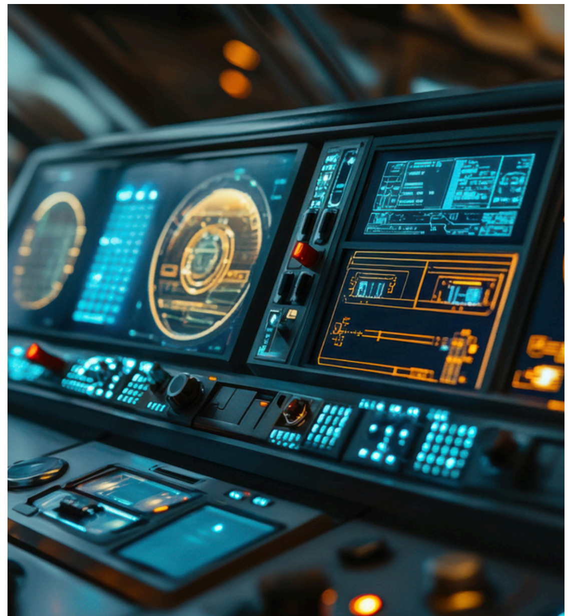

Aviation and IA

Contact us!

Website: https://www.orbitai.com.mx/
Email: info@orbitai.com.mx
Telephone: 811-629-26-58
Linkedin: Orbitaiwithus