Mantis: A Simulation-Grounded Foundation Model for Disease Forecasting

TECHNOLOGY NUMBER: 2026-079



OVERVIEW

Mantis is a plug-and-play, mechanistically interpretable AI platform that delivers accurate, long-range infectious disease forecasts—outperforming expert-trained models without relying on real-world data for training.

- **Universal forecasting:** Delivers reliable, out-of-the-box predictions across diseases and regions, regardless of data availability.
- **Mechanistic insight:** Provides transparent, epidemiologically grounded explanations for each forecast, supporting actionable decision-making.

BACKGROUND

Reliable forecasting of infectious disease is critical for public health readiness—informing hospital staffing, vaccine distribution, and interventions during outbreaks. Yet most current models need large volumes of disease-specific data and frequent expert tuning, making them impractical during new epidemics or in regions where data is scarce. Existing approaches often struggle to generalize, offer only short-term forecasts, or operate as "black boxes" that public health leaders cannot trust. As outbreaks and data constraints continue to pose global challenges, there is a growing market need for forecasting tools that are accurate, rapid, and explainable—especially in resource-limited settings.

Technology ID

2026-079

Category

Software
Software & Content
MOSS - Michigan Open Source
Support

Inventor

Marisa Eisenberg

Further information

Ashwathi lyer ashwathi@umich.edu

View online page



Mantis fundamentally changes the paradigm by learning not from historical outbreaks, but from over 400 million days of realistic, simulated epidemic scenarios. By training solely on mechanistic simulations spanning diverse diseases, interventions, and surveillance conditions, Mantis can produce accurate, probabilistic forecasts for any disease or location—immediately and without model retraining. Critically, each prediction comes with an epidemiological rationale, thanks to a "back-to-simulation" approach that identifies the driving mechanisms behind trends, such as shifts in transmission or reporting. This enables both high accuracy and actionable transparency, making Mantis both scalable and trustworthy—addressing gaps left by traditional and Al-driven models alike.

ADDITIONAL INFORMATION

PROJECT LINKS: GitHub

DEPARTMENT/LAB: Marisa Eisenberg (Department of Epidemiology); Carson Dudley

LICENSE: Polyform-Noncommercial