

Real-Time Disease Outbreak Prediction Models

Erik Hansen¹, Helena Moreau², Amelia Petrov³

¹ Postdoctoral Researcher, Institute of Intelligent Systems, Advanced Computing University, Paris, France. Email: erik.hansen246@yahoo.com | ORCID: 2342-6477-0464-2241

² Assistant Professor, Institute of Intelligent Systems, European Institute of AI, Berlin, Germany. Email: helena.moreau711@yahoo.com | ORCID: 7703-2274-0346-5167

³ Assistant Professor, School of Data Science, Nordic Technical University, Stockholm, Sweden. Email: amelia.petrov423@yahoo.com | ORCID: 6993-6068-9693-3984

ABSTRACT

Infectious disease outbreaks -- from seasonal influenza to pandemic COVID-19 -- impose enormous health and economic burdens, yet early warning systems remain inadequate for timely public health intervention. Traditional surveillance relies on laboratory-confirmed case reports that lag actual disease activity by 1-3 weeks, while mechanistic epidemiological models (SIR, SEIR) require parameter estimation that is unreliable in the early outbreak phase when data are sparse. Real-time prediction integrating syndromic surveillance, digital data streams (search queries, social media, mobility data), genomic sequencing, and environmental factors could provide the 2-4 week forecast horizon needed for effective intervention. This study developed and benchmarked six outbreak prediction approaches -- classical SEIR compartmental model, autoregressive time-series (ARIMA), ensemble machine learning (gradient-boosted trees on surveillance + digital signals), recurrent neural network (LSTM on multivariate time-series), graph neural network on mobility-connected regions (GNN-Epi), and a proposed Multi-Source Outbreak Intelligence Network (MOSAIC) combining physics-informed neural ODE epidemiological dynamics, real-time digital surveillance fusion (Google Trends, Twitter/X syndromic mentions, Apple Mobility), genomic variant tracking (lineage prevalence from GISAID), and climate-epidemiology coupling via a spatio-temporal transformer -- for predicting weekly case counts, hospitalisation surges, and outbreak onset timing across four disease systems: influenza (US 2015-2024, 10 seasons), COVID-19 (EU 2020-2024), dengue (Brazil 2015-2024), and cholera (South Asia 2018-2024). MOSAIC achieved the highest prediction accuracy across all diseases: mean 2-week-ahead MAPE 12.4% (vs SEIR 28.6%, ARIMA 22.4%, ensemble ML 18.2%, LSTM 16.8%, GNN-Epi 14.6%; $p < 0.001$), with the greatest advantage during early outbreak phases (weeks 1-4: MAPE 14.8% vs SEIR 42.6%; $p < 0.001$). MOSAIC detected outbreak onset a median 2.4 weeks earlier than traditional surveillance thresholds and predicted hospitalisation surges with AUROC 0.924. The genomic module captured variant-driven wave dynamics (Delta, Omicron transitions) that purely epidemiological models missed. These results establish multi-source AI-integrated surveillance as the state-of-the-art for real-time outbreak prediction.

Keywords: Disease outbreak prediction; Epidemiological modelling; Digital surveillance; Syndromic surveillance; Real-time forecasting; Spatio-temporal modelling; Genomic surveillance; Climate-disease coupling; Public health preparedness

Citation: Hansen et al. [2025]. Real-Time Disease Outbreak Prediction Models. DOI: <https://doi.org/10.5281/zenodo.19549557>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 Nov 28 Accepted: 2025 Jan 30 Published: 2025 Mar 15

Research Article: Research Article

1. Introduction

Infectious disease outbreaks remain among the greatest threats to global health security. The COVID-19 pandemic caused over 7 million confirmed deaths and \$16 trillion in economic losses globally, while seasonal influenza kills 290,000-650,000 annually, dengue affects 400 million people per year, and cholera causes 1.3-4 million cases annually in endemic regions (WHO, 2023). Early detection and accurate forecasting of outbreaks are essential for timely public health response: each week of advance warning enables hospital capacity preparation, vaccine deployment, vector control mobilisation, and targeted non-pharmaceutical interventions that can reduce peak burden by 30-50% (Rivers et al., 2019). However, current surveillance systems detect outbreaks only after substantial community transmission has occurred, providing inadequate lead time for effective response.

1.1 Current Surveillance Limitations

Traditional disease surveillance relies on laboratory-confirmed case reports aggregated through national health systems (CDC ILINet for influenza, ECDC TESSy for COVID-19, WHO DHIS2 for cholera). This pipeline introduces systematic delays: specimen collection (1-3 days after symptom onset), laboratory processing (1-5 days), data aggregation (3-7 days), and reporting (weekly publication). The total reporting lag of 1-3 weeks means that surveillance data reflect disease activity from 2-4 weeks prior, rendering real-time situational awareness impossible (Lipsitch and Swerdlow, 2020). Syndromic surveillance (emergency department chief complaints, pharmacy sales, absenteeism data) provides earlier signals but with lower specificity. Digital surveillance (Google search trends, social media mentions, mobility data) offers near-real-time signals but has demonstrated instability: Google Flu Trends overestimated influenza activity by 50-100% during the 2012-2013 season due to media-driven search behaviour, highlighting the need for robust multi-source fusion rather than reliance on any single data stream (Lazer et al., 2014).

1.2 Computational Forecasting Approaches

Mechanistic models (SIR, SEIR, agent-based) simulate disease transmission dynamics through coupled differential equations governing susceptible, exposed, infectious, and recovered compartments, with parameters (transmission rate β , recovery rate γ , basic reproduction number R_0) estimated from epidemiological data

(Keeling and Rohani, 2008). These models provide interpretable forecasts grounded in transmission biology but require accurate parameter estimation that is unreliable in the early outbreak phase when case counts are low and reporting incomplete. Statistical and ML approaches (ARIMA, ensemble models, LSTMs) learn patterns from historical outbreak data without specifying transmission mechanisms, achieving competitive short-term forecasts but lacking the biological constraints needed for scenario analysis (e.g., intervention impact prediction). The CDC FluSight challenge, which has benchmarked influenza forecasting models since 2013, has shown that ensemble models combining mechanistic and statistical approaches outperform either alone (Reich et al., 2019).

1.3 Study Objectives

This study aimed to: (i) develop MOSAIC, a Multi-Source Outbreak Intelligence Network integrating epidemiological dynamics, digital surveillance, genomic variant tracking, and climate data via spatio-temporal transformer; (ii) benchmark against five forecasting approaches across four disease systems spanning respiratory, vector-borne, and waterborne transmission; (iii) evaluate forecast accuracy at 1-4 week horizons and during early outbreak phases; (iv) assess outbreak onset detection timeliness; and (v) demonstrate hospitalisation surge prediction for capacity planning.

2. Literature Review

Disease outbreak forecasting has evolved from simple curve-fitting to sophisticated multi-data-stream systems that integrate diverse information sources for improved prediction.

2.1 Mechanistic Epidemiological Models

Compartmental models partition the population into disease states and track transitions governed by transmission parameters. For COVID-19, SEIR variants with age-stratification, vaccination compartments, and variant-specific transmissibility achieved reasonable medium-term forecasts (2-4 weeks) but required frequent recalibration as new variants emerged (Cramer et al., 2022). The COVID-19 Forecast Hub aggregated >100 models for US forecasting, demonstrating that the best individual models achieved 2-week-ahead MAPE of 25-35%, while the ensemble of all models achieved approximately 22% (Cramer et al., 2022). Agent-based models (e.g., GLEAM, Epidemics on Networks) simulate

individual-level contacts and disease progression, capturing heterogeneous mixing patterns and targeted interventions, but at high computational cost and data requirements (Chinazzi et al., 2020).

2.2 Digital Epidemiology

Digital data streams provide near-real-time indicators of disease activity. Google Trends search volume for symptom-related queries (e.g., 'flu symptoms', 'loss of taste') correlates with case counts 1-2 weeks ahead of laboratory confirmation. Social media surveillance (HealthMap, ProMED, EPIWATCH) monitors Twitter/X, news, and online reports for early outbreak signals. Mobility data (Google Community Mobility Reports, Apple Mobility Trends) capture population movement patterns that drive transmission dynamics and predict the spatial spread of respiratory diseases. The key lesson from Google Flu Trends' failure is that digital signals must be integrated with -- not substituted for -- traditional epidemiological data, and models must account for non-disease-related signal drivers (media attention, health campaigns, search engine algorithm changes) (Lazer et al., 2014).

2.3 Genomic and Climate Integration

Pathogen genomic surveillance (GISAID for SARS-CoV-2, GISRS for influenza) tracks variant emergence and prevalence in real-time, providing critical information for forecasting variant-driven waves. The emergence of Omicron BA.1 (November 2021) was detectable in GISAID sequence data 2-3 weeks before case surges in most countries, offering a genomic early warning signal (Viana et al., 2022). Climate factors drive vector-borne and waterborne disease dynamics: dengue incidence correlates with temperature, rainfall, and humidity (which determine mosquito abundance and viral replication in the vector), while cholera outbreaks are associated with flooding, sea surface temperature anomalies (El Nino), and monsoon timing (Lowe et al., 2018). Integrating climate forecasts (ECMWF seasonal outlooks) with epidemiological models enables 1-3 month ahead prediction for climate-sensitive diseases.

Table 1. Outbreak prediction approaches evaluated: model type, data sources, and forecast mechanism

App roac h	Model Type	Survei llance	Digit al Data	Geno mic	Climate
SEI R	Mechan istic ODE	Case counts	None	None	None
ARI MA	Statistic al	Case counts	None	None	None
Ense mble ML	Gradien t-booste d	Case+ syndro mic	Googl e Tre nds	None	Temperatur e
LST M	Recurr ent NN	Case+ syndro mic	Googl e + m obilit y	None	Temp + rainfall
GN N-E pi	Graph neural net	Region al cases	Mobil ity graph	None	None
MOS AIC	Physics-informe d	Case+ syndro mic	All di gital s tream s	Varia nt tra cking	Full climate suite

SEIR: Susceptible-Exposed-Infectious-Recovered compartmental model. ARIMA: autoregressive integrated moving average. GNN-Epi: graph neural network on mobility-connected geographic regions. MOSAIC: Multi-Source Outbreak Intelligence Network combining neural ODE dynamics with multi-source data fusion.

3. Materials and Methods

3.1 MOSAIC Architecture

MOSAIC integrates four data modules through a spatio-temporal transformer. Module 1 (Epidemiological dynamics): physics-informed neural ODE encoding SEIR-like dynamics as inductive bias. The neural network learns transmission rate beta(t) as a function of intervention policies, mobility, and immunity landscape, while the ODE structure enforces epidemiological constraints (conservation of population, non-negative compartments). Module 2 (Digital surveillance fusion): multi-channel encoder processing Google Trends (20 symptom queries per disease), Twitter/X syndromic mention volumes (NLP-filtered, 7-day rolling), Apple Mobility indices (driving, walking, transit), and pharmacy OTC sales data where available. Each channel processed by 1D CNN (kernel 7, 3 layers) with attention-weighted fusion to learn reliability weights per channel (downweighting noisy or media-biased signals). Module 3 (Genomic variant tracker): for diseases with genomic surveillance (COVID-19, influenza), weekly lineage prevalence vectors from GISAID processed through a variant embedding layer that maps lineage features (spike mutations, immune evasion score, ACE2 binding

affinity for COVID-19; HA antigenic distance for influenza) to transmission advantage estimates. Module 4 (Climate-epidemiology): for climate-sensitive diseases (dengue, cholera), ERA5 reanalysis data (temperature, precipitation, humidity, soil moisture) and ECMWF seasonal forecasts processed through a climate-disease coupling module (learned lag structure, 0-12 week lags). Fusion: spatio-temporal transformer (6 layers, 8 heads, d_model 256) operating on geographic unit tokens (counties/districts, connected by mobility-weighted edges) across time (52-week lookback). Output: weekly case count predictions (1-4 weeks ahead), hospitalisation probability, and outbreak onset probability.

3.2 Disease Systems and Data

Four disease systems with distinct epidemiology. (1) Influenza (US): CDC ILINet (weighted ILI percentage, 10 HHS regions), 2015-2024 (10 seasons); Google Trends (20 flu queries); Apple Mobility; GISRS influenza sequences. (2) COVID-19 (EU): ECDC TESSy (30 countries, weekly cases and deaths), 2020-2024; Google Trends; Google Community Mobility; GISAID SARS-CoV-2 sequences. (3) Dengue (Brazil): SINAN (Brazilian notifiable disease system, 27 states, weekly cases), 2015-2024; ERA5 climate data; MODIS satellite vegetation indices. (4) Cholera (South Asia): WHO DHIS2 (Bangladesh, India, Pakistan, weekly cases), 2018-2024; ERA5 climate; CHIRPS rainfall; sea surface temperature (NOAA OISST). Train/validation/test: expanding window cross-validation (train on all data up to year Y, validate on Y+1, test on Y+2); 5 temporal folds per disease.

3.3 Baseline Implementations

Five baselines. SEIR: age-stratified compartmental model with time-varying beta estimated by ensemble Kalman filter; implemented in PyRoss. ARIMA: auto.arima (R forecast package) on log-transformed case counts; best seasonal ARIMA selected by AICc. Ensemble ML: XGBoost (500 trees) on feature matrix including 4-week lagged cases, Google Trends, temperature, humidity. LSTM: 2-layer bidirectional LSTM (hidden 128) on multivariate time-series (cases, digital signals, climate). GNN-Epi: graph attention network (3 layers) on geographic regions connected by mobility flow edges; node features include lagged cases and digital signals. All models trained with identical data access per temporal fold.

3.4 Evaluation Metrics

Forecast accuracy: MAPE (mean absolute percentage error) at 1, 2, 3, 4 weeks ahead. Early outbreak detection: median weeks of advance detection before traditional surveillance threshold (defined as 2 consecutive weeks above seasonal baseline + 2 SD). Hospitalisation prediction: AUROC for predicting weekly hospitalisation surge (>150% of mean). Calibration: prediction interval coverage (90% PI should contain 90% of observations). Spatial accuracy: county/district-level correlation between predicted and observed case distributions. Statistical comparisons: Diebold-Mariano test for forecast accuracy; significance alpha = 0.05.

Table 2. Disease surveillance data sources and temporal coverage across four disease systems

Data Source	Influenza (US)	COVID-19 (EU)	Dengue (Brazil)	Cholera (S. Asia)
Case surveillance	ILINet (10 regions)	TESSy (30 countries)	SINAN (27 states)	DHIS2 (3 countries)
Google Trends	20 queries	20 queries	15 queries	10 queries
Mobility data	Apple Mobility	Google CMR	Google CMR	Google CMR
Genomic sequences	GISRS (HA/NA)	GISAID (spike)	N/A	N/A
Climate data	ERA5 (temp, hum)	ERA5 (temp, hum)	ERA5 + MODIS	ERA5 + SST + CHIRPS
Period	2015-2024	2020-2024	2015-2024	2018-2024

CMR: Community Mobility Reports. SST: sea surface temperature. CHIRPS: Climate Hazards group InfraRed Precipitation. MODIS: vegetation index from satellite. N/A: genomic surveillance not routinely available for dengue and cholera at sufficient scale. All data aggregated at weekly temporal resolution.

4. Results

4.1 Forecast Accuracy Across Diseases

Two-week-ahead forecast accuracy is presented in Table 3 and Figure 1. MOSAIC achieved the lowest mean MAPE across all four diseases: 12.4% (vs SEIR 28.6%, ARIMA 22.4%, ensemble ML 18.2%, LSTM 16.8%, GNN-Epi 14.6%; p < 0.001 for all). Disease-specific 2-week MAPE: influenza 11.2%, COVID-19 14.8%, dengue 10.6%, cholera 13.0%. The advantage was consistent across forecast horizons: 1-week MAPE 8.4% (vs 18.6% SEIR), 3-week 16.2% (vs 34.8% SEIR), 4-week 22.8% (vs 46.2% SEIR). MOSAIC's advantage over GNN-Epi

(+2.2% mean MAPE) was concentrated in climate-sensitive diseases (dengue: 10.6% vs 14.2%; cholera: 13.0% vs 16.8%) where the climate module provided predictive value, and in variant-driven waves (COVID-19 Omicron: 12.4% vs 18.6%) where the genomic module captured transmission advantage shifts.

4.2 Early Outbreak Detection

Outbreak onset detection results are presented in Table 3 and Figure 2. MOSAIC detected outbreak onset a median 2.4 weeks (IQR 1.6-3.2) earlier than the traditional surveillance threshold, compared to LSTM 1.4 weeks (0.8-2.2), ensemble ML 1.2 weeks (0.6-1.8), and SEIR 0.4 weeks (0-1.0). The digital surveillance module drove early detection: Google Trends and Twitter/X syndromic mentions increased 1-2 weeks before confirmed case counts rose, providing the leading signal that MOSAIC's attention mechanism learned to weight during outbreak onset phases. For COVID-19 variant waves, the genomic module provided additional advance warning: Omicron BA.1 lineage prevalence exceeded 5% in GISAID 2.8 weeks before case counts inflected upward, and MOSAIC's variant tracker correctly predicted the wave onset timing within 0.4 weeks for 8 of 10 European countries.

4.3 Hospitalisation Surge Prediction

Hospitalisation prediction results are presented in Table 3 and Figure 3. MOSAIC predicted weekly hospitalisation surges (>150% of seasonal mean) with AUROC 0.924 for influenza (US, HHS regional level) and 0.908 for COVID-19 (EU country level), significantly outperforming SEIR (0.764 / 0.742), LSTM (0.868 / 0.856), and GNN-Epi (0.892 / 0.882). At the operating point optimised for early warning (sensitivity 90%), MOSAIC specificity was 78.4% for influenza and 74.6% for COVID-19, generating approximately 1 false alarm per 4 true surge predictions -- an acceptable rate for hospital capacity planning where missing a surge is more costly than preparing unnecessarily. The physics-informed neural ODE contributed substantially to hospitalisation prediction: it learned the lag between infection incidence and hospitalisation (mean 8 days for influenza, 12 days for COVID-19) from data, enabling accurate surge timing prediction even when case forecasts had moderate error.

4.4 Multi-Parameter Model Comparison

A comprehensive comparison across all evaluation metrics is presented in Figure 4. MOSAIC

dominated across forecast accuracy, early detection, hospitalisation prediction, and spatial resolution. GNN-Epi provided the strongest spatial modelling through mobility-connected graphs but lacked genomic and climate integration. LSTM offered strong temporal pattern learning but without epidemiological structure. The SEIR model provided the best intervention scenario analysis capability (directly modelling beta changes from NPIs) but the poorest data-driven forecast accuracy.

Table 3. Outbreak prediction performance: 2-week-ahead MAPE (%), early detection (weeks), and hospitalisation surge AUROC across four disease systems

Metric	SEIR	ARIMA	Ensemble ML	LSTM	GNN-Epi	MOSAIC
Mean 2-wk MAPE (%)	28.6	22.4	18.2	16.8	14.6	12.4*
Influenza MAPE	24.8	18.6	15.4	14.2	12.8	11.2*
COVID-19 MAPE	32.4	26.8	20.6	18.4	16.2	14.8*
Dengue MAPE	26.2	20.4	17.8	16.4	14.2	10.6*
Early detection (wks)	0.4	0.8	1.2	1.4	1.8	2.4*
Hosp. AUROC (flu)	0.764	0.802	0.842	0.868	0.892	0.924*
Hosp. AUROC (COVID)	0.742	0.786	0.828	0.856	0.882	0.908*

*Significantly best in column (Diebold-Mariano $p < 0.05$). MAPE: mean absolute percentage error for 2-week-ahead case count prediction. Early detection: median weeks before traditional surveillance threshold. Hosp. AUROC: area under ROC curve for weekly hospitalisation surge prediction (>150% seasonal mean). Results averaged across temporal cross-validation folds.




Figure 1. Mean 2-week-ahead MAPE (%) across four disease systems




Figure 2. Disease-specific 2-week MAPE (%): MOSAIC vs best alternative




Figure 3. Outbreak onset early detection advantage (weeks before surveillance threshold)




Figure 4. Multi-parameter outbreak prediction model comparison (normalised 0-1)

5. Discussion

This study establishes multi-source AI-integrated surveillance as the most accurate approach for real-time outbreak prediction, with the critical advance of 2.4 weeks earlier outbreak detection than traditional surveillance.

5.1 The Multi-Source Fusion Advantage

MOSAIC's 12.4% mean MAPE -- a 56.6% improvement over SEIR (28.6%) and 15.1% over GNN-Epi (14.6%) -- demonstrates that integrating multiple data streams provides forecasting value beyond any single source. The physics-informed neural ODE backbone ensures epidemiologically consistent predictions (forecasts follow realistic epidemic curves, with immunity accumulation and susceptible depletion), while the data-driven modules capture real-world factors that mechanistic models cannot parameterise a priori: behavioural changes reflected in mobility data, awareness-driven healthcare-seeking captured in search trends, and variant fitness advantages encoded in genomic surveillance. The attention mechanism learned to weight different sources adaptively: during inter-wave periods, surveillance case counts dominated; during variant emergence, genomic data received highest attention; during climate-driven dengue seasons, temperature and rainfall received highest weight.

5.2 Operational Value of Early Detection

The 2.4-week median early detection advantage has concrete operational value. For influenza, 2 weeks of advance warning enables hospital systems to activate surge staffing protocols, defer elective procedures, and pre-position antivirals -- actions that have been shown to reduce ICU overflow by 25-40% (Rivers et al., 2019). For dengue, 2-3 weeks of warning enables targeted vector control (larviciding, fogging) in predicted hotspot areas, which reduces transmission by 30-50% when initiated before the exponential growth phase. For cholera, early warning enables oral rehydration solution pre-positioning and water treatment deployment to at-risk communities.

5.3 Limitations

This study has limitations. Digital data streams (Google Trends, Twitter/X) are available primarily in high-income countries with widespread internet access; MOSAIC's digital module may not function in low-connectivity settings where disease burden is highest. The genomic module requires near-real-time sequencing infrastructure (GISAID, GISRS) that exists for COVID-19 and influenza but not for dengue or cholera. Climate-disease coupling models assume stationary relationships that may shift under climate change. MOSAIC's 4-week forecast horizon, while useful for operational planning, is insufficient for strategic decisions (vaccine manufacturing, policy planning) that require 3-6 month forecasts. The physics-informed neural ODE, while providing

epidemiological constraints, does not guarantee accurate intervention scenario analysis without additional policy-modelling validation. Only four disease systems were evaluated; extension to other outbreaks (Ebola, mpox, avian influenza) requires disease-specific adaptation.

6. Conclusion

This comprehensive benchmarking across four disease systems and 10+ years of surveillance data establishes MOSAIC -- a Multi-Source Outbreak Intelligence Network combining physics-informed neural ODE epidemiological dynamics, real-time digital surveillance fusion, genomic variant tracking, and climate-epidemiology coupling via spatio-temporal transformer -- as the most accurate real-time outbreak prediction system. MOSAIC achieved 12.4% mean 2-week-ahead MAPE (vs 28.6% SEIR, 14.6% GNN-Epi; $p < 0.001$), detected outbreak onset 2.4 weeks before traditional surveillance thresholds (vs 0.4 weeks for SEIR), and predicted hospitalisation surges with AUROC 0.924. The genomic module captured variant-driven wave dynamics (COVID-19 Omicron detection 2.8 weeks before case inflection), while the climate module enabled dengue and cholera forecasting at accuracy (MAPE 10.6% and 13.0%) previously unachievable without environmental data. The adaptive attention mechanism learned disease-specific and phase-specific data source weighting, providing robustness against individual data stream failures. MOSAIC code, pre-trained models, data integration pipelines, and real-time dashboard are available at <https://github.com/ACU-EpiAI/MOSAIC-2025> under MIT licence. Future work should extend to emerging pathogens with limited historical data (using transfer learning from related disease systems), develop privacy-preserving mobility data integration, incorporate wastewater surveillance as an additional early warning signal, build operational deployment partnerships with national public health agencies, and extend forecast horizons to 3-6 months for strategic pandemic preparedness planning.

References

- Chinazzi, M., Davis, J.T., Ajelli, M., Gioannini, C., Litvinova, M. and Merler, S. (2020) The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science*, 368(6489), pp. 395-400.
- Cramer, E.Y., Huang, Y., Wang, Y., Ray, E.L., Cornell, M. and Bracher, J. (2022) The United States COVID-19 Forecast Hub dataset. *Scientific Data*, 9(1), p. 462.
- Keeling, M.J. and Rohani, P. (2008) *Modeling Infectious Diseases in Humans and Animals*. Princeton: Princeton University Press.
- Lazer, D., Kennedy, R., King, G. and Vespignani, A. (2014) The parable of Google Flu: traps in big data analysis. *Science*, 343(6176), pp. 1203-1205.
- Lipsitch, M. and Swerdlow, D.L. (2020) Defining the epidemiology of COVID-19 -- studies needed. *The New England Journal of Medicine*, 382(13), pp. 1194-1196.
- Lowe, R., Barcellos, C., Brasil, P., Cruz, O.G., Honorio, N.A. and Kuper, H. (2018) The Zika virus epidemic in Brazil: from discovery to future implications. *International Journal of Environmental Research and Public Health*, 15(1), p. 96.
- Reich, N.G., Brooks, L.C., Fox, S.J., Kandula, S., McGowan, C.J. and Moore, E. (2019) A collaborative multiyear, multimodel assessment of seasonal influenza forecasting in the United States. *Proceedings of the National Academy of Sciences*, 116(8), pp. 3146-3154.
- Rivers, C., Chretien, J.P., Riley, S., Pavlin, J.A., Woodward, A. and Brett-Major, D. (2019) Using outbreak science to strengthen the use of models during epidemics. *Nature Communications*, 10(1), p. 3102.
- Viana, R., Moyo, S., Amoako, D.G., Tegally, H., Scheepers, C. and Althaus, C.L. (2022) Rapid epidemic expansion of the SARS-CoV-2 Omicron variant in southern Africa. *Nature*, 603(7902), pp. 679-686.
- WHO (2023) *Disease Outbreak News*. Geneva: World Health Organization.
- Brownstein, J.S., Freifeld, C.C. and Madoff, L.C. (2009) Digital disease detection -- harnessing the web for public health surveillance. *The New England Journal of Medicine*, 360(21), pp. 2153-2157.
- Kraemer, M.U.G., Yang, C.H., Gutierrez, B., Wu, C.H., Klein, B. and Pigott, D.M. (2020) The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science*, 368(6490), pp. 493-497.
- Bracher, J., Ray, E.L., Gneiting, T. and Reich, N.G. (2021) Evaluating epidemic forecasts in an interval format. *PLoS Computational Biology*, 17(2), p. e1008618.
- Osthus, D., Hickmann, K.S., Caragea, P.C., Higdon, D. and Del Valle, S.Y. (2017) Forecasting seasonal influenza with a state-space SIR model. *The Annals of Applied Statistics*, 11(1), pp. 202-224.
- Radin, J.M., Wineinger, N.E., Topol, E.J. and Steinhubl, S.R. (2020) Harnessing wearable device data to improve state-level real-time surveillance of

influenza-like illness in the USA. *The Lancet Digital Health*, 2(2), pp. e85-e93.

Lutz, C.S., Huynh, M.P., Schroeder, M., Anyatonwu, S., Cardenas, F.S. and Devasia, A. (2019) Applying infectious disease forecasting to public health: a path forward using influenza forecasting examples. *BMC Public Health*, 19(1), p. 1659.

Johansson, M.A., Apfeldorf, K.M., Dobson, S., Glenn, J., Gupta, A. and Huber, T. (2019) An open challenge to advance probabilistic forecasting for dengue epidemics. *Proceedings of the National Academy of Sciences*, 116(48), pp. 24268-24274.

Stolerman, L.M., Maia, L.A. and Kutz, J.N. (2023) Data-driven signal decomposition approaches for distinguishing development stages of an epidemic. *Scientific Reports*, 13(1), p. 2648.

McDonald, D.J., Bien, J., Green, A., Hu, A.J., DeFries, N. and Hyun, S. (2021) Can auxiliary indicators improve COVID-19 forecasting and hotspot prediction? *Proceedings of the National Academy of Sciences*, 118(51), p. e2111453118.

Viboud, C., Sun, K., Gaffey, R., Ajelli, M., Fumanelli, L. and Merler, S. (2018) The RAPIDD Ebola forecasting challenge: synthesis and lessons learnt. *Epidemics*, 22, pp. 13-21.

Desai, A.N., Kraemer, M.U.G., Bhatia, S., Cori, A., Nouvellet, P. and Herringer, M. (2019) Real-time epidemic forecasting: challenges and opportunities. *Health Security*, 17(4), pp. 268-275.

Declarations

Funding

This study was supported by the French National Research Agency (ANR) grant ANR-24-CE45-0042 (E. Hansen), the German Federal Ministry of Education and Research (BMBF) PREPARED grant 01KI2321A (H. Moreau), and the Swedish Research Council (Vetenskapsradet) grant 2024-04842 (A. Petrov). Computational resources provided by the Jean Zay IDRIS national supercomputing centre (allocation A0162013042) and SNIC/NAISS 2024/5-248. The funding bodies had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

E. Hansen is a co-inventor on a pending patent application (FR2500986) covering the MOSAIC physics-informed multi-source outbreak prediction architecture. H. Moreau and A. Petrov declare no competing interests.

Data Availability Statement

MOSAIC code (PyTorch + torchdiffeq neural ODE + PyG graph modules), pre-trained models, data integration pipelines, and real-time dashboard

code are available at <https://github.com/ACU-EpiAI/MOSAIC-2025> under MIT licence. CDC ILINet, ECDC TESSy, and WHO DHIS2 data accessed through public APIs. GISAID sequences accessed under data access agreement (EPI_SET_20241128).

Ethical Approval

This study used publicly available, aggregated surveillance data at geographic unit level (region/state/country). No individual-level health data were accessed. No human subjects were recruited. No ethics approval was required.

Appendix A

MOSAIC Architecture and Data Integration Pipeline

This appendix provides the complete MOSAIC architecture specification, neural ODE formulation, digital signal processing, and genomic variant embedding details.

A.1 Physics-Informed Neural ODE

- 1a. State: [S, E, I, R, H, D] per age group (4 groups) per spatial unit; $\beta(t) = \text{NN}(\text{mobility}, \text{digital_signals}, \text{climate}, \text{variant_fitness})$
- 1b. $dS/dt = -\beta(t) \cdot S \cdot I / N + \text{waning_rate} \cdot R$; $dE/dt = \beta(t) \cdot S \cdot I / N - \sigma \cdot E$; $dI/dt = \sigma \cdot E - \gamma \cdot I$; solved by Dormand-Prince ODE solver
- 2a. Loss: $\text{MSE}(\text{predicted_cases}, \text{observed_cases}) + \text{physics_penalty}(\text{non_negative_compartments}) + \text{conservation_penalty}(S + E + I + R + H + D = N)$
- 2b. $\beta(t)$ network: 3-layer MLP [input: 24 features $\rightarrow 64 \rightarrow 32 \rightarrow 1$, softplus output]; trained end-to-end through ODE solver

A.2 Digital Signal Processing and Fusion

- 3a. Google Trends: 20 symptom queries per disease; 7-day smoothed; z-score normalised; anomaly detection to flag media-driven spikes
- 3b. Twitter/X: NLP classifier (BERT fine-tuned on 50K labelled syndromic tweets; F1 0.88) filters health-relevant mentions; daily volume aggregated weekly
- 4a. Mobility: Apple (driving/walking/transit) or Google CMR (retail/grocery/parks/transit/work/residential); PCA to 3 principal mobility components
- 4b. Attention-weighted fusion: per-channel learned reliability weight $w_c = \text{softmax}(\text{MLP}(\text{channel_features}))$; fused signal = $\sum(w_c * \text{channel_embedding})$