

TOWARD WORLD MODELS FOR EPIDEMIOLOGY

Zeeshan Memon^{1,*}, **Yiqi Su**^{2,*}, **Christo Kurisummoottil Thomas**³, **Walid Saad**⁴,
Liang Zhao¹, **Naren Ramakrishnan**²

¹Department of Computer Science, Emory University, Atlanta, GA, USA

²Department of Computer Science, Virginia Tech, Alexandria, VA, USA

³Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, MA, USA

⁴Bradley Department of Electrical and Computer Engineering, Virginia Tech, Alexandria, VA, USA

{zeeshan.memon, liang.zhao}@emory.edu

{yiqisu, walids, naren}@vt.edu

cthomas2@wpi.edu

ABSTRACT

World models have emerged as a unifying paradigm for learning latent dynamics, simulating counterfactual futures, and supporting planning under uncertainty. In this paper, we argue that computational epidemiology is a natural and underdeveloped setting for world models. This is because epidemic decision-making requires reasoning about latent disease burden, imperfect and policy-dependent surveillance signals, and intervention effects are mediated by adaptive human behavior. We introduce a conceptual framework for epidemiological world models, formulating epidemics as controlled, partially observed dynamical systems in which (i) the true epidemic state is latent, (ii) observations are noisy and endogenous to policy, and (iii) interventions act as sequential actions whose effects propagate through behavioral and social feedback. We present three case studies that illustrate why explicit world modeling is necessary for policy-relevant reasoning: strategic misreporting in behavioral surveillance, systematic delays in time-lagged signals such as hospitalizations and deaths, and counterfactual intervention analysis where identical histories diverge under alternative action sequences.

1 INTRODUCTION

World models are growing in popularity across multiple domains in machine learning. Bridging latent dynamics modeling, imagination-based planning, and counterfactual rollouts, the core idea is to learn an internal representation of how the world evolves and then use that representation to support decision-making Ha & Schmidhuber (2018). Such learned simulators have been deployed for control, robotics, game playing, and embodied agents, including Dreamer-style agents and MuZero-like systems Hafner et al. (2020); Schrittwieser et al. (2020). At their core, world models provide the glue among perception, action, and reasoning: they separate latent state from noisy observations, thereby enabling prediction, simulation, and counterfactual reasoning at scale Silver et al. (2021).

A typical world model includes several ingredients. First, it is assumed to have a latent state representation that compresses the relevant aspects of the environment into a compact internal variable Kaelbling et al. (1998). Given the state representation, we also need transition dynamics that describe how this latent state changes over time, often as a function of actions or interventions Hızlı et al. (2025); Wang et al. (2024). World models also include an observation function that maps latent states back to the measurements we actually see, capturing noise, reporting artifacts, and partial observability Lillemark et al. (2026); Li et al. (2025). Expanding further, we aim to support a notion of interventions or actions, i.e., what can be changed, when, and how those changes propagate through the system. Finally, a world model is trained and evaluated against objectives or criteria that reflect the intended use, such as predictive accuracy, planning performance, or the ability to generate plausible counterfactual rollouts under alternative decisions Ha & Schmidhuber (2018); Ke et al. (2021).

*Equal contribution.

Why world models for epidemiology? In this paper, we investigate world models for computational epidemiology. Computational epidemiology has undergone a major transformation over the last two decades (even prior to COVID-19) Salathe et al. (2012); O’Shea (2017). Early work centered on compartmental, mass-action models with SEIR-style dynamics: analytically tractable, interpretable, and naturally aligned with policy questions Kong et al. (2022). Over time, the data landscape changed through new surveillance modalities, including search trends, digital traces, mobility data, social media, and other forms of syndromic surveillance, which in turn encouraged machine learning approaches layered on top of mechanistic epidemiological modeling.

The biggest beneficiary of this shift has been forecasting Tabataba et al. (2017). Short-term case prediction, hospital demand estimation, and ensemble forecasting have become central, with forecast accuracy often treated as the dominant evaluation metric. But the field has increasingly moved beyond forecasting toward scenario analysis Chen & Qiu (2020), wherein the goal is to support questions from policy makers such as: should we close schools? should we quarantine a sub-population? and what to expect under alternative interventions? This naturally invites counterfactual rollouts of policies and stress-testing of non-pharmaceutical interventions, vaccination strategies, and reopening plans Davies et al. (2020); Ferguson et al. (2020).

However, much of current scenario analysis research still assumes that the world is given, i.e., fixed dynamics, fixed behavior, and fixed observation processes Couzin & Heins (2023). Data are often treated as exogenous, and humans are treated as parameters rather than participants in the system being studied Krakowski (2025). From a world model perspective, these assumptions are limiting, because they conflate the underlying latent epidemic state with shifting observation mechanisms and adaptive human behavior, and they restrict counterfactual reasoning to interventions applied to an otherwise static simulator Ha & Schmidhuber (2018); Fenichel et al. (2011).

Our contributions are three-fold: (i) a formal conceptual framework for epidemiological world models, (ii) reframing of epidemic decision making into reasoning with world models, which in turn supports robustness of decisions under partial observability, endogenous measurement, and model misspecification, (iii) three empirical case studies from our research that demonstrate the kinds of policy reasoning world models make possible.

2 SOME EPIDEMIOLOGY FOR MACHINE LEARNERS

This section serves as a ‘101’ overview that provides the epidemiological background needed for future sections.

2.1 COMPARTMENTAL MODELS

Compartmental models have historically been the backbone of epidemiological modeling Choi & Shim (2020). In this approach, the entire population under study is expected to be distributed across several states and dynamical rules described as a system of ODEs govern how people move across states. Classic examples include SIR (Susceptible-Infected-Recovered) Prodanov (2022) and SEIR (Susceptible-Exposed-Infected-Recovered) Burke (2024). In their standard form, these models are expressed as systems of ordinary differential equations (ODEs) that describe how the fraction (or count) of the population in each compartment changes continuously over time Thiagarajan et al. (2022). Such models can be enlarged with additional states that distinguish age structure, spatial structure, and risk strata Wan et al. (2025). These models are often built around mass-action assumptions and rely on homogeneous mixing as a modeling convenience, even though real populations rarely mix homogeneously.

A primary virtue of compartmental models is their interpretability and policy relevance Gostic et al. (2020). Their parameters can often be treated as levers that correspond to actionable mechanisms, and summary quantities like R_0 (reproduction number) are commonly used to communicate epidemic status Dietz (1993); Smith & Schapira (2012). Informally, R_0 is the expected number of secondary infections caused by a single infected individual in a fully susceptible population, under a given set of behavioral and environmental conditions. When $R_0 = 1$, each infected person infects one other person on average, so incidence is expected to remain roughly steady rather than growing or shrinking. When $R_0 > 1$, infections grow because each case generates more than one new case on average, leading to expanding transmission. When $R_0 < 1$, infections decline because each case

generates fewer than one new case on average, so transmission tends to die out over time. In general, these reproduction numbers should be understood as summaries of system behavior rather than fixed constants Heesterbeek & Roberts (2007). Most importantly, interventions (by public health authorities) are imposed to bring R_0 to manageable levels.

2.2 AGENT-BASED MODELS

Agent-based models provide a more fine-grained representation of epidemic spread by simulating individuals and their interactions Tracy et al. (2018); Chopra et al. (2023). Here instead of ODEs, we model the population as a graph where nodes denote people and edges capture, e.g., contact structure. Once again nodes can be in various states (S/I/R or S/E/I/R, for instance). Time is typically discrete in these models and over the passage of time, people move from one state to another based on contact dynamics (e.g., if a susceptible node is exposed to a sufficient number of infected neighbors, the node moves to the exposed state in the case of SEIR). In an agent-based model we typically run simulations many times under the same model settings (often with different random seeds) to obtain a composite picture of disease spread across the network O’Gara et al. (2025).

2.3 OBSERVATIONS AS NOISY PROXIES

In epidemiology, true incidence is typically latent. This is because infections may be asymptomatic, reporting may be delayed, and testing is often partial or uneven. As a result, observed signals are noisy proxies for the underlying epidemic state: reported cases are not the same as infections, hospitalizations tend to lag transmission dynamics, and deaths reflect historical state rather than current conditions Dailey et al. (2007); Su et al. (2026a). For a machine learning audience, the key distinction is between latent state and observed data streams Memon et al. (2025); Gostic et al. (2020).

2.4 ADVERSARIAL ACTORS IN EPIDEMIOLOGY

Epidemiological systems can include adversarial or strategic behavior, especially with self-reported data from individuals. For example, a national study in Kenya found that while only 12% of individuals reported not wearing masks, direct observation revealed that nearly 90% were non-compliant, a discrepancy exceeding 75 percentage points that was consistent across demographic groups Jakubowski et al. (2021). Similar survey-based studies in countries like the USA document substantial social-desirability and recall biases in self-reported masking and vaccination behavior, compounded by spatial heterogeneity, policy variation, and limited auditability of daily behaviors Taube et al. (2023).

2.5 DATA REGIMES AND SURVEILLANCE EVOLUTION

Traditional surveillance systems have relied on lab-confirmed cases and sentinel networks. More recently, surveillance has expanded to include digital and opportunistic data streams such as mobility traces, search activity, social media signals, and wastewater measurements Zhang et al. (2024); Salathé (2018). In these modern regimes, data can be endogenous: policy affects behavior, behavior affects data, and data affects policy. As a result, the observation process itself can change under intervention, meaning that shifts in measured signals may reflect surveillance and behavioral feedback effects rather than true changes in underlying transmission.

2.6 LIMITS OF FORECASTING-CENTRIC THINKING

Forecasts are best understood as conditional extrapolations: they project observed trends forward under the assumptions encoded by the model and the data-generating process. As a result, forecast skill is not the same as policy usefulness. A model can produce accurate short-term predictions while still relying on wrong structural assumptions, and it can present misleading confidence when the observed signals are biased, delayed, or otherwise distorted Yuan et al. (2021). In particular, when we ask a high-performing model to adjust its forecasts under a change of scenario, (e.g., a new intervention), it must implicitly assume how that intervention would alter transmission, behavior, and surveillance assumptions that are often unstated, underidentified from the data, or simply absent from the model Eisenberg-Guyot et al. (2025). This motivates the need for true world models.

3 WORLD MODELS IN COMPUTATIONAL EPIDEMIOLOGY

3.1 EPIDEMICS AS DYNAMICAL SYSTEMS WITH INTERVENTION

We view epidemics as dynamical systems in which interventions are central drivers of system evolution rather than incidental covariates. In this framing, policies are actions applied over time, and humans are adaptive agents whose behavior responds to both disease conditions and policy choices Talukder et al. (2024). This creates a setting where decision-making must be counterfactual, i.e., reasoning about what would happen under alternative actions, not merely predictive. This perspective aligns with planning-as-inference views of decision-making, in which policies are evaluated by reasoning over future trajectories under alternative actions Wood et al. (2022). It also requires separating latent epidemic reality from the observed surveillance signals that imperfectly measure it, and it demands planning under uncertainty when both dynamics and observations can shift in response to interventions Parag et al. (2025).

3.2 INGREDIENTS OF AN EPIDEMIOLOGICAL WORLD MODEL

We model an epidemic as a controlled, partially observed dynamical system in which (i) the epidemic state is latent, (ii) observations are imperfect and policy-dependent proxies of that state, and (iii) interventions affect transmission primarily through behavioral and social feedback. This formulation extends classical state-space models for infectious-disease surveillance by explicitly treating interventions as actions in a sequential decision process rather than as fixed covariates. We adopt the standard formulation of controlled, partially observed dynamical systems from planning and reinforcement learning, and instantiate it for epidemiological policy reasoning Kaelbling et al. (1998).

World model as a controlled state-space system. An epidemiological world model specifies a generative process over latent states, actions, and observations:

$$a_t \sim \pi_\psi(\cdot | h_t), \quad h_t = \phi(o_{1:t-1}, a_{1:t-1}), \quad (1)$$

$$x_{t+1} \sim P_\theta(x_{t+1} | x_t, a_t), \quad (2)$$

$$o_t \sim \Omega_\theta(o_t | x_t, a_{t-1}), \quad (3)$$

where $x_t \in \mathcal{X}$ denotes the latent epidemic state, $o_t \in \mathcal{O}$ the observed surveillance signals, and $a_t \in \mathcal{A}$ the intervention applied at time t . The internal information state h_t summarizes past observations and interventions and provides the information available for action selection under partial observability.

The encoder $\phi(\cdot)$ corresponds to an explicit belief state obtained via filtering, but need not be an exact Bayesian posterior. More generally, $\phi(\cdot)$ may be implemented as a learned recurrent or memory-based representation that implicitly encodes uncertainty, delayed effects, and long-range policy context, without assuming identifiability of the latent state.

Latent state. The epidemic state x_t is entirely latent and need not be disentangled or semantically aligned with individual epidemiological variables. It can consist of structured components, continuous latent embeddings, or hybrid representations, provided it is sufficient to render the dynamics Markov under intervention. Conceptually, the state subsumes multiple interacting factors:

$$x_t = (x_t^{\text{epi}}, x_t^{\text{imm}}, x_t^{\text{net}}, x_t^{\text{beh}}, x_t^{\text{reg}}), \quad (4)$$

where x_t^{epi} captures true infections and susceptibility (including unobserved and asymptomatic cases), x_t^{imm} denotes immunity and waning, x_t^{net} captures the contact and mobility structure, x_t^{beh} represents the behavioral and social responses (e.g., compliance, risk perception), and x_t^{reg} encompasses slowly varying regime factors such as circulating variants, seasonality, or healthcare capacity. These quantities are unified as state variables because they are not directly observed yet jointly determine future epidemic evolution under intervention.

Dynamics. The transition model $P_\theta(x_{t+1} | x_t, a_t)$ governs how epidemic, behavioral, and regime components evolve in response to interventions. A defining feature of epidemiological dynamics is that actions affect transmission indirectly via behavior and contact patterns, inducing nonlinear and nonstationary effects:

$$P_\theta(x_{t+1} | x_t, a_t) \approx P_\theta(x_{t+1} | x_t, \eta(x_t^{\text{beh}}, a_t)), \quad (5)$$

where $\eta(\cdot)$ denotes the induced response of population behavior and mixing. Regime components evolve on longer timescales, enabling the model to represent gradual drift and abrupt shifts, such as variant emergence, within a unified controlled dynamical system.

Observations. Observations o_t correspond to quantities that are actually measured such as reported cases, hospitalizations, deaths, wastewater signals, or mobility traces, and are generated by a policy-dependent observation process:

$$o_t \sim \Omega_\theta(o_t \mid x_t, a_{t-1}). \tag{6}$$

This process subsumes all mechanisms by which latent epidemic states become observable, including but not limited to testing practices, reporting pipelines, delays, measurement noise, and behavioral responses induced by interventions. Because policies affect both disease dynamics and surveillance itself, the relationship between x_t and o_t is endogenous to past actions. This separation between latent reality and observation is essential for reasoning under partial observability, delayed feedback, and systematic or strategic measurement bias.

Policies and actions. A policy π_ψ maps the internal information state h_t to interventions a_t . The action space includes non-pharmaceutical interventions, vaccination and testing strategies, risk communication, and enforcement intensity. Epidemiologically, actions are typically coarse and temporally extended, and their effects are mediated through social response rather than direct biological control.

4 MOTIVATING CHALLENGES

We present three illustrative case studies that demonstrate why forecasting-centric approaches are insufficient for policy-relevant epidemiological reasoning. These studies are intentionally minimal and controlled. Rather than optimizing predictive accuracy, they are designed to expose structural failure modes that arise due to strategic behaviors, the nature of surveillance, and counterfactual intervention, settings that epidemiological world models are explicitly constructed to address. For each case study, we identify the problem, describe how a world-model perspective clarifies the failure mode, and enables more policy-relevant reasoning and representative results and findings.

4.1 EPIDEMIC MODELING UNDER USER DECEPTION

Epidemiological modeling increasingly relies on large-scale, crowdsourced behavioral data to track and forecast disease transmission dynamics in near real-time Banerjee & Lian (2022). Beyond clinical surveillance records, self-reported information on vaccination uptake, mask usage, and adherence to social distancing has become central to assessing the impact of non-pharmaceutical interventions (NPIs) Böff et al. (2025).

This misalignment between reported and actual behavior can systematically bias parameter estimates and weaken NPIs Wu et al. (2020b); Pullano et al. (2021); Milanesi & Nicolao (2023).

Unfortunately, there are substantial and persistent gaps between self-reported and directly observed behaviors Jakubowski et al. (2021); Taube et al. (2023). Even vaccination records, though more structured, exhibit persistent misreporting rates of 5–15% due to reporting delays, incomplete registries, and misclassification Hurley et al. (2026); Daley et al. (2024); Irving et al. (2009). Beyond measurement error, self-reported epidemiological data are often strategically distorted. Individuals may misreport or withhold information due to testing avoidance, stigma, policy incentives, or distrust in public health authorities Levy et al. (2022); Shiman et al. (2023); Su et al. (2026b).




Figure 1: Impact of strategic misreporting on epidemic control.

Figure 1 illustrates how strategic misreporting degrades epidemic control through biased observation of individual behavior under three regimes—no misreporting, mixed misreporting, and pure misre-

porting. Although all scenarios begin from the same initial transmission level, increasing degrees of misreporting systematically delay the reduction of the reproduction number below the threshold. This occurs because misreported vaccination and masking behavior leads the public health authority to overestimate compliance, resulting in weaker or delayed interventions despite persistent transmission. Consequently, misreporting acts as an endogenous distortion of the observation process that directly impairs control effectiveness rather than as random measurement noise.

However, most epidemiological models either ignore strategic misreporting Patel et al. (2005) or treat it as static parameter uncertainty (e.g., Tanner et al. (2008)), and are evaluated primarily by predictive accuracy on reported cases. World models address misreporting by explicitly modeling how biased, incentive-driven observations are generated from latent epidemic states, allowing policies to be evaluated and optimized on inferred reality rather than on distorted reports Ha & Schmidhuber (2018); Hafner et al. (2020).

4.2 EPIDEMIC MODELING UNDER A MOVING TARGET

A large body of work has shown that epidemiological surveillance signals are inherently delayed proxies of underlying transmission dynamics. Delays arise from incubation periods, care-seeking behavior, testing, hospital admissions, and reporting pipelines, causing observed trends in confirmed cases, hospitalizations, and deaths to lag true infections by days to weeks Gostic et al. (2020); Pullano et al. (2021); Wu et al. (2020a). These delays have been shown to bias real-time inference, degrade forecast performance near turning points, and delay apparent responses to interventions in both influenza and COVID-19 surveillance Reich et al. (2019); Davies et al. (2020).




Figure 2: Dynamics of stabilization of reported ILI case counts across various countries.

(2024); Fox et al. (2022). Incorporating these dynamics allows us to explicitly capture such observation dynamics into our world model.

Accurate infectious disease forecasting can support advance planning for vaccination clinics, targeted communication campaigns, and the distribution of antiviral medications, and may help guide mitigation measures such as school closures or cancellation of large gatherings during pandemics Reich et al. (2019); Mathis et al. (2024). However, because forecasting models are trained to extrapolate delayed surveillance signals, their predictions remain conditioned on the same lagged and potentially distorted observations Gostic et al. (2020); Pullano et al. (2021). As a result, forecasting alone cannot resolve the fundamental disconnect between observed data and the underlying epidemic state, motivating the need for world models that explicitly separate latent dynamics from the observation process and reason counterfactually about interventions.

In Fig 2 we track reported ILI (influenza like illness) case counts in multiple countries of Latin America against the final, ‘resting’, case counts after updates and corrections are incorporated over subsequent weeks. As can be seen, some countries are fast stabilizing (presumably, owing to the sophistication of their surveillance apparatus) and others are slower-stabilizing. The Centers for Disease Control and Prevention monitors influenza activity through a nationwide surveillance system in the United States that aggregates clinical, laboratory, and reporting data. Despite its importance for decision-making, this system inherently lags real-time influenza transmission due to delays in symptom onset, care-seeking, and reporting Mathis et al.

4.3 COUNTERFACTUAL ANALYSIS

We next illustrate the necessity of world models for counterfactual policy analysis. In practice, the true number of infected individuals is rarely observable, and decision makers must rely on other (delayed) signals such as hospitalizations. Accordingly, we treat hospitalizations as the available ground-truth outcome and relate them to latent infections through a fixed infection–hospitalization ratio Deng et al. (2025).

Figure 3 shows the resulting hospitalization trajectories from a calibrated baseline and a counterfactual scenario with earlier vaccine eligibility and stricter masking mandates. While the baseline scenario follows the realized course of the epidemic, the counterfactual intervention yields a lower and delayed hospitalization peak, corresponding to improved infection control. Crucially, this divergence arises despite identical observed histories prior to intervention, and cannot be recovered by forecasting approaches that condition only on the realized trajectory.




Figure 3: Hospital admission trajectories under baseline and counterfactual interventions.

This example highlights a fundamental limitation of forecasting-centric methods: predicting future observations does not answer counterfactual questions about alternative actions. By contrast, world models enable intervention-aware simulation, allowing policy evaluation in terms of latent epidemic control rather than passive extrapolation of observed trends Alles et al. (2025).

5 FRAMEWORK

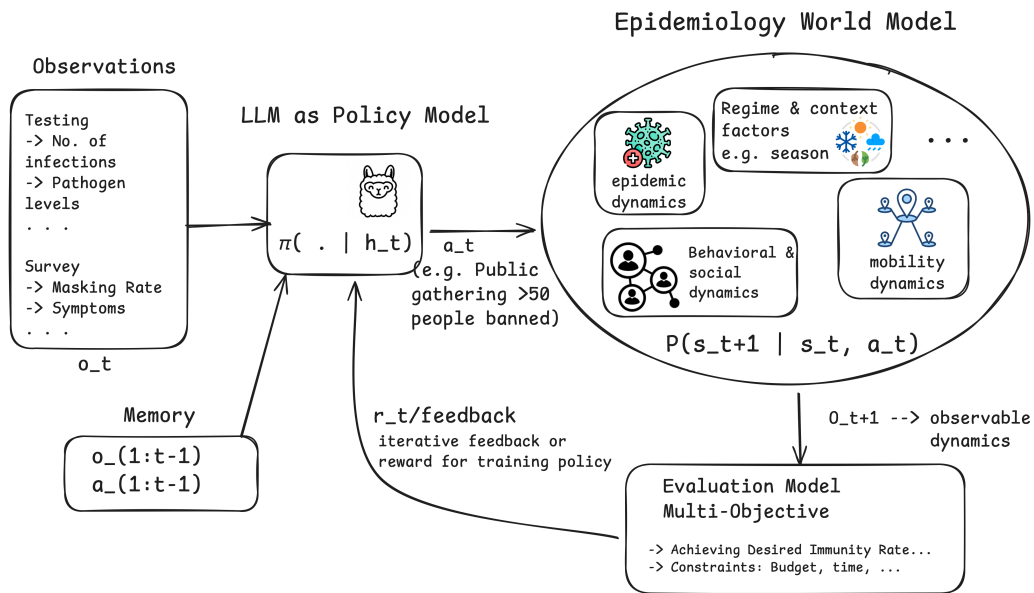


Figure 4: Epidemiological world model loop.

Having defined the components of an epidemiological world model, we now describe how these components are composed into a computational framework for policy reasoning. The framework is algorithm-agnostic: it does not prescribe how the world model is learned or how policies are opti-

mized, but specifies the interaction loop through which candidate interventions are proposed, rolled out, evaluated, and refined. Figure 4 illustrates this interaction loop, highlighting the separation between policy proposal, world-model simulation, observation generation, and feedback.

At each decision point, the framework maintains an internal information state h_t , as defined in Section 3.2. Given h_t , candidate policies propose interventions, which are evaluated by rolling out the epidemiological world model forward in time. The resulting latent trajectories and synthetic observations are scored using policy-relevant objectives. This loop can be executed once for scenario analysis or iterated to refine policies.

5.1 POLICY MODELS

Policies can be rule-based, learned, or hybrid. In particular, large language models (LLMs) can act as *policy models* that map the internal information state h_t into candidate interventions. In this role, LLMs provide a strong prior over plausible decision rules and intervention structures, but do not determine outcomes directly: all candidate actions are evaluated through world-model rollouts.

5.2 WORLD MODEL ROLLOUTS

Given a candidate intervention sequence $a_{t:t+H}$, the world model simulates forward trajectories of the latent epidemic state:

$$x_{t+1:t+H} \sim P_\theta(\cdot \mid x_t, a_{t:t+H}), \quad (7)$$

where H denotes the planning horizon. The use of a horizon reflects the fact that epidemiological interventions are temporally extended and exhibit delayed effects: changes in transmission, hospitalizations, and mortality typically manifest over weeks rather than instantaneously. Evaluating policies therefore requires reasoning over multi-step futures, including transient dynamics, delayed surveillance signals, and cumulative outcomes such as healthcare overload.

Rollouts propagate epidemiological, behavioral, and regime components, allowing interventions to affect transmission through contact structure and compliance. This enables the framework to capture both short-term responses (e.g., behavioral adaptation following an intervention) and longer-term consequences (e.g., immunity accumulation or variant-driven regime shifts).

In parallel, the observation model generates synthetic surveillance signals:

$$o_{t+1:t+H} \sim \Omega_\theta(\cdot \mid x_{t+1:t+H}, a_{t:t+H}), \quad (8)$$

making explicit how different policies shape both the underlying epidemic and the data that would be observed. This distinction is critical in epidemiology, where observed trends may reflect surveillance effects rather than true changes in transmission.

5.3 EVALUATION AND ITERATIVE REFINEMENT

Each rollout is evaluated using outcome measures aligned with epidemiological decision-making. These may include health outcomes (e.g., cumulative infections, deaths, hospital overload), social and economic costs, and equity considerations. Evaluation is typically multi-objective and can involve constraints rather than a single scalar score (e.g., avoiding ICU capacity violations). Importantly, evaluation operates on latent trajectories as well as synthetic observations, allowing policies to be assessed on their true epidemic impact rather than solely on predicted reported cases.

Evaluation outputs provide feedback to the policy model. In the simplest case, candidate interventions are compared and filtered based on dominance or constraint satisfaction. More generally, evaluation can induce scalar or vector-valued reward signals derived from multi-objective criteria, enabling reinforcement learning, bandit-style updates, or preference-based optimization. When policies are represented by language-based models, structured evaluation feedback can be incorporated through iterative prompting or refinement.

The framework is agnostic to the specific update mechanism, but makes explicit how intervention outcomes feed back into subsequent policy proposals. This iterative loop—policy proposal \rightarrow world model rollout \rightarrow evaluation \rightarrow refinement—constitutes the computational core by which epidemiological world models support counterfactual reasoning and planning under uncertainty, beyond forecasting-centric approaches.

5.4 LEARNING THE WORLD MODEL AND POLICY

We learn the world model from trajectories of observations and interventions, $(o_{1:T}, a_{1:T})$. Under partial observability, we maintain a belief state

$$h_t = f_\phi(h_{t-1}, o_t, a_{t-1}), \quad (9)$$

together with controlled latent dynamics and an observation model:

$$x_{t+1} \sim P_\theta(x_{t+1} | x_t, a_t), \quad (10)$$

$$o_t \sim \Omega_\theta(o_t | x_t, a_{t-1}), \quad (11)$$

where ϕ denotes the parameters of the belief-state encoder and θ denotes the jointly learned parameters of the transition and observation models. We train the world model by maximizing a sequential variational objective:

$$\mathcal{L}_{\text{WM}} = \sum_{t=1}^T \mathbb{E}_{q_\phi(x_t)} [\log \Omega_\theta(o_t | x_t)] - \beta \text{KL}(q_\phi(x_{1:T}) \| p_\theta(x_{1:T} | a_{1:T})), \quad (12)$$

which encourages latent states to explain observations while remaining consistent with controlled dynamics. After training, policies are optimized using model-based rollouts:

$$\pi^* = \arg \max_{\pi} \mathbb{E} \left[\sum_{\tau=t}^{t+H} r(x_\tau, a_\tau) \right], \quad (13)$$

where trajectories are simulated using the learned world model. This enables counterfactual evaluation of interventions beyond forecasting.

6 EXPERIMENTAL RESULTS

6.1 SETTINGS

We instantiate the framework using COVID-19 surveillance data from all 50 US states (January 2021–late 2022). State features include weekly epidemiological signals (hospitalizations per 100k, confirmed cases, vaccination rates), static demographic and healthcare-capacity attributes, and 13-dimensional policy actions from the Oxford COVID-19 Government Response Tracker (OxCGRT) Hale et al. (2021) encoded on a 0–4 ordinal scale. The world model is CovidLLM, a LLaMA-2-7B Touvron et al. (2023) base fine-tuned via SFT with LoRA Hu et al. (2022); epidemiological context is serialized into a structured prompt alongside the proposed policy action, with a GRU encoder providing temporal context over multi-week sequences. The model predicts next-week hospitalization trend as an ordinal classification over five classes. We evaluate two policy model families, GPT-4o-mini Achiam et al. (2024) and Qwen-7BBai et al. (2023), in closed-loop rollouts across all 50 states starting from September 5, 2022, with a planning horizon of $H=6$ weeks. For Qwen-7B, we additionally evaluate two optimization strategies: iterative feedback and GRPO. Policy alignment is measured as the percentage of proposed actions matching historically realized OxCGRT interventions; hospitalization reduction is the mean percentage decrease in hospitalizations per 100k over the horizon.

6.2 RESULTS

Policy evaluation. GPT-4o-mini achieves 44.8% alignment with historically realized interventions across all 50 states and 13 policy dimensions, correctly capturing the directionality and coarse structure of non-pharmaceutical interventions. Qwen-7B obtains 30.3% alignment, a notable gap that suggests smaller open-weight models are less reliably calibrated to the scale conventions of the OxCGRT action space, likely due to weaker instruction-following and reduced exposure to policy-structured data during pretraining or reasoning capabilities. GPT-4o-mini drives a mean hospitalization reduction of 46.5% over the 6-week horizon, while Qwen-7B achieves 39.1%, both exceeding the 37.1% reduction observed under historically realized interventions ($9.09 \rightarrow 5.67$ per 100k).

Policy optimization (Qwen-7B). The policy optimization conditions use world-model rollouts as a signal to refine the Qwen-7B policy model’s system prompt beyond zero-shot prompting. The

Table 1: Policy evaluation and optimization results across all 50 US states ($H=6$ weeks). Policy alignment is the percentage of proposed actions matching historically realized interventions. Hosp. reduction is the mean percentage decrease in hospitalizations per 100k over the rollout horizon.

Category	Method	Policy Alignment (%) \uparrow	Hosp. Reduction (%) \uparrow
Policy Eval.	GPT-4o-mini	44.8	-46.5
	Qwen-7B	30.3	-39.1
Policy Optim.	Iterative Feedback	49.5	-48.7
	GRPO	42.2	-46.1

iterative feedback method Agrawal et al. (2026) observes the full sequence of hospitalization outcomes across all rollout steps and uses these observations to rewrite the policy model’s system prompt, steering it toward lower-burden intervention sequences in subsequent iterations. This approach achieves 49.5% alignment and a hospitalization reduction of 48.7% (the highest across all conditions), indicating that prompt-level refinement guided by simulated epidemic outcomes can elicit more epidemiologically coherent intervention recommendations. The *GRPO* treats the negative end-of-horizon hospitalization count as a scalar reward and optimizes the policy model’s action distribution directly via group relative policy gradient updates over on-policy rollouts in the world model, achieving 42.2% alignment and a 46.1% hospitalization reduction.

7 DISCUSSION

While our framework provides a principled formulation of epidemiological world models, several challenges remain. First, latent epidemic states are inherently difficult to identify from noisy, delayed, and policy-dependent observations, leading to ambiguity between true infections, behavior, and reporting processes. This is further compounded by limited and heterogeneous data availability across regions, making multi-modal integration (e.g., cases, hospitalizations, mobility, surveys) both necessary and nontrivial. In addition, world models introduce significant computational overhead due to latent-state inference, multi-step rollouts, and policy optimization, particularly in multi-region or agent-based settings. Designing realistic and tractable policy representations also remains an open problem, as interventions are often high-level, constrained, and shaped by social and institutional factors that are difficult to encode.

Looking forward, several directions can strengthen this line of work. Incorporating mechanistic priors and causal structure may improve identifiability and interpretability of latent states, while advances in multi-modal learning and missing-data handling can make models more robust to real-world surveillance conditions. Scalable training and planning methods, including hierarchical and approximate approaches, will be essential for deployment at population scale. Additionally, developing richer evaluation protocols beyond forecasting accuracy, such as counterfactual validity and policy effectiveness, will be critical for assessing real-world impact. Ultimately, integrating world models with multi-agent, multi-region settings and adaptive behavioral dynamics offers a promising path toward more reliable and actionable epidemic decision support.

8 CONCLUSION

In closing, we advocate a shift from forecasting-centric epidemic modeling to an explicit world-model view that treats epidemics as controlled, partially observed dynamical systems with endogenous surveillance and adaptive human behavior. While our studies are minimal, they surface concrete failure modes that forecasting alone cannot resolve and illustrate the kinds of robustness-oriented analyses that epidemiological world models make possible. We believe that epidemiological world models can serve as reliable substrates for intervention design and stress-testing.

ACKNOWLEDGEMENTS

This work is supported in part by US National Science Foundation grants CCF-1918770, IIS-2509636, IIS-2312794, DBI-2412389, and CNS-2225511. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsor(s).

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint*, 2024. URL <https://arxiv.org/abs/2303.08774>.
- Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziemis, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. Gpa: Reflective prompt evolution can outperform reinforcement learning. *arXiv preprint*, 2026. URL <https://arxiv.org/abs/2507.19457>.
- Marvin Alles, Xingyuan Zhang, Patrick van der Smagt, and Philip Becker-Ehmck. Latent action world models for control with unlabeled trajectories. *arXiv preprint*, 2025. URL <https://arxiv.org/abs/2512.10016>.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint*, 2023. URL <https://arxiv.org/abs/2309.16609>.
- Sandipan Banerjee and Yongsheng Lian. Data driven covid-19 spread prediction based on mobility and mask mandate information. *Applied Intelligence*, 52(2):1969–1978, 2022. ISSN 0924-669X. doi: 10.1007/s10489-021-02381-8. URL <https://doi.org/10.1007/s10489-021-02381-8>.
- Lena Böff, Antonia Bartz, Manuela Harries, André Karch, Annette Aigner, Veronika K. Jaeger, Berit Lange, MuSPAD Consortium Group, COVIMOD Consortium Group, and RESPINOW Consortium Group. Dynamics of contact behaviour by self-reported COVID-19 vaccination and infection status during the COVID-19 pandemic in Germany: an analysis of two large population-based studies. *BMC Medicine*, 23(1):406, 2025. doi: 10.1186/s12916-025-04211-x. URL <https://doi.org/10.1186/s12916-025-04211-x>.
- Donald S. Burke. Origins of the problematic E in SEIR epidemic models. *Infectious Disease Modelling*, 9(3):673–679, 2024. ISSN 2468-0427. doi: <https://doi.org/10.1016/j.idm.2024.03.003>. URL <https://www.sciencedirect.com/science/article/pii/S2468042724000368>.
- Xiaohui Chen and Ziyi Qiu. Scenario analysis of non-pharmaceutical interventions on global covid-19 transmissions. *arXiv preprint*, 2020. URL <https://arxiv.org/abs/2004.04529>.
- Wongyeong Choi and Eunha Shim. Optimal strategies for vaccination and social distancing in a game-theoretic epidemiologic model. *Journal of theoretical biology*, 505, Nov. 2020. URL <https://doi.org/10.1016/j.jtbi.2020.110422>.
- Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Arnau Quera-Bofarull, Balaji Krishnamurthy, B. Aditya Prakash, and Ramesh Raskar. Differentiable agent-based epidemiology. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '23*, pp. 1848–1857, Richland, SC, 2023. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450394321. URL <https://dl.acm.org/doi/10.5555/3545946.3598851>.
- Iain D. Couzin and Conor Heins. Emerging technologies for behavioral research in changing environments. *Trends in Ecology & Evolution*, 38(4):346–354, 2023. ISSN 0169-5347. doi: <https://doi.org/10.1016/j.tree.2022.11.008>. URL <https://www.sciencedirect.com/science/article/pii/S0169534722002889>. Special issue: Animal behaviour in a changing world.
- Lynne Dailey, Rochelle E. Watkins, and Aileen J. Plant. Timeliness of data sources used for influenza surveillance. *Journal of the American Medical Informatics Association*, 14(5):626–631, 2007. doi: 10.1197/jamia.M2328. URL <https://doi.org/10.1197/jamia.M2328>.

- Matthew F. Daley, Liza M. Reifler, Jo Ann Shoup, Jason M. Glanz, Bruno J. Lewin, Nicola P. Klein, Elyse O. Kharbanda, Huong Q. McLean, Simon J. Hambidge, Jennifer C. Nelson, Allison L. Naleway, Eric S. Weintraub, Michael M. McNeil, Hilda Razzaghi, and James A. Singleton. Influenza vaccination accuracy among adults: Self-report compared with electronic health record data. *Vaccine*, 42(11):2740–2746, 2024. ISSN 0264-410X. doi: <https://doi.org/10.1016/j.vaccine.2024.03.052>. URL <https://www.sciencedirect.com/science/article/pii/S0264410X24003396>.
- Nicholas G. Davies, Adam J. Kucharski, Rosalind M. Eggo, Amy Gimma, W. John Edmunds, Thibaut Jombart, Kathleen O’Reilly, Akira Endo, Joel Hellewell, Emily S. Nightingale, Billy J. Quilty, Christopher I. Jarvis, Timothy W. Russell, Petra Klepac, Nikos I. Bosse, Sebastian Funk, Sam Abbott, Graham F. Medley, Hamish Gibbs, Carl A. B. Pearson, Stefan Flasche, Mark Jit, Samuel Clifford, Kiesha Prem, Charlie Diamond, Jon Emery, Arminder K. Deol, Simon R. Procter, Kevin van Zandvoort, Yueqian Fiona Sun, James D. Munday, Alicia Rosello, Megan Auzenbergs, Gwen Knight, Rein M. G. J. Houben, and Yang Liu. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *The Lancet Public Health*, 5(7):e375–e385, 2020. doi: [10.1016/S2468-2667\(20\)30133-X](https://doi.org/10.1016/S2468-2667(20)30133-X). URL [https://doi.org/10.1016/S2468-2667\(20\)30133-X](https://doi.org/10.1016/S2468-2667(20)30133-X).
- Yangyang Deng, Yun Kim, Anna Bratcher, Jefferson M. Jones, Muloongo Simuzingili, Adi V. Gundlapalli, Melissa Briggs Hagen, Ronaldo Iachan, and Kristie E. N. Clarke. Ratio of infections to COVID-19 cases and hospitalizations in the United States based on SARS-CoV-2 seroprevalence data, september 2021–february 2022. *Open Forum Infectious Diseases*, 12(1):ofae719, 2025. doi: [10.1093/ofid/ofae719](https://doi.org/10.1093/ofid/ofae719). URL <https://doi.org/10.1093/ofid/ofae719>.
- Klaus Dietz. The estimation of the basic reproduction number for infectious diseases. *Statistical Methods in Medical Research*, 2(1):23–41, 1993. doi: [10.1177/096228029300200103](https://doi.org/10.1177/096228029300200103). URL <https://doi.org/10.1177/096228029300200103>. PMID: 8261248.
- Jerzy Eisenberg-Guyot, Katrina L Kezios, Seth J Prins, and Sharon Schwartz. Inconsistent consistency: evaluating the well-defined intervention assumption in applied epidemiological research. *International Journal of Epidemiology*, 54(2):dyaf015, 03 2025. ISSN 1464-3685. doi: [10.1093/ije/dyaf015](https://doi.org/10.1093/ije/dyaf015). URL <https://doi.org/10.1093/ije/dyaf015>.
- Eli P. Fenichel, Carlos Castillo-Chavez, Michele G. Ccedia, Gerardo Chowell, P. A. Parra, Graham J. Hickling, George Holloway, Richard Horan, Bruce Morin, Charles Perrings, Michael Springborn, Luis Velázquez, and C. Villalobos. Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences of the United States of America*, 108(15):6306–6311, 2011. doi: [10.1073/pnas.1011250108](https://doi.org/10.1073/pnas.1011250108). URL <https://doi.org/10.1073/pnas.1011250108>.
- Neil M. Ferguson, Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina Cuomo-Dannenburg, Amy Dighe, Iliaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green, Arran Hamlet, Wes Hinsley, Lucy C. Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick G. T. Walker, Caroline Walters, Peter Winskill, Charles Whittaker, Christl A. Donnelly, Steven Riley, and Azra C. Ghani. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Technical Report 9, Imperial College London, 2020. URL <https://doi.org/10.25561/77482>.
- Spencer J. Fox, Michael Lachmann, Mauricio Tec, Remy Pasco, Spencer Woody, Zhanwei Du, Xutong Wang, Tanvi A. Ingle, Emily Javan, Maytal Dahan, Kelly Gaither, Mark E. Escott, Stephen I. Adler, S. Claiborne Johnston, James G. Scott, and Lauren Ancel Meyers. Real-time pandemic surveillance using hospital admissions and mobility data. *Proceedings of the National Academy of Sciences*, 119(7):e2111870119, 2022. doi: [10.1073/pnas.2111870119](https://doi.org/10.1073/pnas.2111870119). URL <https://www.pnas.org/doi/abs/10.1073/pnas.2111870119>.
- Katelyn M. Gostic, Lauren McGough, Edward B. Baskerville, Sam Abbott, Keya Joshi, Christine Tedijanto, Rebecca Kahn, Rene Niehus, James A. Hay, Pablo M. De Salazar, Joel Hellewell, Sophie Meakin, James D. Munday, Nikos I. Bosse, Katharine Sherratt, Robin N. Thompson, Laura F. White, Jana S. Huisman, Julien Scire, Sebastian Bonhoeffer, Tanja Stadler, Jacco

- Wallinga, Sebastian Funk, and Marc Lipsitch. Practical considerations for measuring the effective reproductive number, R_t . *PLoS Computational Biology*, 16(12):e1008409, 2020. URL <https://doi.org/10.1371/journal.pcbi.1008409>.
- David Ha and Jürgen Schmidhuber. World models. 2018. doi: 10.5281/ZENODO.1207631. URL <https://zenodo.org/record/1207631>.
- Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint*, 2020. URL <https://arxiv.org/abs/1912.01603>.
- Thomas Hale, Noam Angrist, Rafael Goldszmidt, Beatriz Kira, Anna Petherick, Toby Phillips, Samuel Webster, Emily Cameron-Blake, Laura Hallas, Saptarshi Majumdar, et al. A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature human behaviour*, 5(4):529–538, 2021. URL <https://doi.org/10.1038/s41562-021-01079-8>.
- J.A.P. Heesterbeek and M.G. Roberts. The type-reproduction number T in models for infectious disease control. *Mathematical Biosciences*, 206(1):3–10, 2007. ISSN 0025-5564. doi: <https://doi.org/10.1016/j.mbs.2004.10.013>. URL <https://www.sciencedirect.com/science/article/pii/S0025556405001549>. Alcalá Special Issue.
- Çağlar Hızlı, Çağatay Yıldız, Matthias Bethge, ST John, and Pekka Martinen. Identifying latent state transitions in non-linear dynamical systems. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=d16mJDyQN6>.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. URL <https://openreview.net/pdf?id=nZeVKeeFYf9>.
- Laura P. Hurley, Lori A. Crane, Matthew F. Daley, Brenda Beaty, John D. Rice, Kamonthip J. Homdayjanakul, Jason Lyons, Yingbo Lou, Laura Elliott, Jo Ann Shoup, Elise Robertson, Brian Manning, Carla L. Black, Peng-Jun Lu, and James A. Singleton. Accuracy of self-reported vaccination status using surveys in safety-net, integrated and rural health systems in colorado. *Vaccine*, 72:128101, 2026. ISSN 0264-410X. doi: <https://doi.org/10.1016/j.vaccine.2025.128101>. URL <https://www.sciencedirect.com/science/article/pii/S0264410X25013994>.
- Stephanie A. Irving, James G. Donahue, David K. Shay, Tina L. Ellis-Coyle, and Edward A. Belongia. Evaluation of self-reported and registry-based influenza vaccination status in a wisconsin cohort. *Vaccine*, 27(47):6546–6549, 2009. ISSN 0264-410X. doi: <https://doi.org/10.1016/j.vaccine.2009.08.050>. URL <https://www.sciencedirect.com/science/article/pii/S0264410X09012353>.
- Aleksandra Jakubowski, Dennis Egger, Carlyne Nekesa, Layna Lowe, Michael Walker, and Edward Miguel. Self-reported vs directly observed face mask use in kenya. *JAMA Network Open*, 4(7):e2118830–e2118830, 07 2021. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2021.18830. URL <https://doi.org/10.1001/jamanetworkopen.2021.18830>.
- Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998. URL [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X).
- Nan Rosemary Ke, Aniket Didolkar, Sarthak Mittal, Anirudh Goyal, Guillaume Lajoie, Stefan Bauer, Danilo Rezende, Yoshua Bengio, Michael Mozer, and Christopher Pal. Systematic evaluation of causal discovery in visual model based reinforcement learning. In *NeurIPS Workshop on Causal Inference & Machine Learning*, 2021. URL <https://openreview.net/forum?id=gp5Uzbl-9C->.
- Lingcai Kong, Mengwei Duan, Jin Shi, Jie Hong, Zhaorui Chang, and Zhijie Zhang. Compartmental structures used in modeling COVID-19: a scoping review. *Infectious Diseases of Poverty*, 11(1):72, 2022. doi: 10.1186/s40249-022-01001-y. URL <https://doi.org/10.1186/s40249-022-01001-y>.

- Sebastian Krakowski. Human-ai agency in the age of generative ai. *Information and Organization*, 35(1):100560, 2025. ISSN 1471-7727. doi: <https://doi.org/10.1016/j.infoandorg.2025.100560>. URL <https://www.sciencedirect.com/science/article/pii/S1471772725000065>.
- Andrea Gurmankin Levy, Alistair Thorpe, Laura D. Scherer, Aaron M. Scherer, Frank A. Drews, Jorie M. Butler, Nicole Burpo, Holly Shoemaker, Vanessa Stevens, and Angela Fagerlin. Misrepresentation and Nonadherence Regarding COVID-19 Public Health Measures. *JAMA Network Open*, 5(10):e2235837, 2022. doi: [10.1001/jamanetworkopen.2022.35837](https://doi.org/10.1001/jamanetworkopen.2022.35837).
- Xinqing Li, Xin He, Le Zhang, Min Wu, Xiaoli Li, and Yun Liu. A comprehensive survey on world models for embodied ai. *arXiv preprint*, 2025. URL <https://arxiv.org/abs/2510.16732>.
- Hansen Jin Lillemark, Benhao Huang, Fangneng Zhan, Yilun Du, and Thomas Anderson Keller. Flow equivariant world models: Memory for partially observed dynamic environments. *arXiv preprint*, 2026. URL <https://arxiv.org/abs/2601.01075>.
- Sarabeth M. Mathis, Alexander E. Webber, Tomás M. León, Erin L. Murray, Monica Sun, Lauren A. White, Logan C. Brooks, Alden Green, Addison J. Hu, and Roni Rosenfeld. Evaluation of FluSight influenza forecasting in the 2021–22 and 2022–23 seasons with a new target laboratory-confirmed influenza hospitalizations. *Nature Communications*, 15(1):6289, 2024. ISSN 2041-1723. doi: [10.1038/s41467-024-50601-9](https://doi.org/10.1038/s41467-024-50601-9). URL <https://doi.org/10.1038/s41467-024-50601-9>.
- Zeeshan Memon, Chen Ling, Ruochen Kong, Vishwanath Seshagiri, Andreas Zufle, and Liang Zhao. Deep identification of propagation trees. *arXiv preprint*, 2025. URL <https://arxiv.org/abs/2503.00646>.
- S. Milanesi and G. D. Nicolao. Correction of Italian under-reporting in the first COVID-19 wave via age-specific deconvolution of hospital admissions. *arXiv preprint*, 2023. URL <https://doi.org/10.48550/arxiv.2307.09978>.
- Jesse O’Shea. Digital disease detection: A systematic review of event-based internet biosurveillance systems. *International journal of medical informatics*, 101:15–22, 2017. URL <https://doi.org/10.1016/j.ijmedinf.2017.01.019>.
- David O’Gara, Matt Kasman, Laurent Hébert-Dufresne, and Ross A. Hammond. Adaptive behaviour during epidemics: a social risk appraisal approach to modelling dynamics. *Journal of The Royal Society Interface*, 22(222):20240363, 01 2025. ISSN 1742-5689. doi: [10.1098/rsif.2024.0363](https://doi.org/10.1098/rsif.2024.0363). URL <https://doi.org/10.1098/rsif.2024.0363>.
- Kris V. Parag, Ben Lambert, Christl A. Donnelly, and Sandor Beregi. Asymmetric limits on timely interventions from noisy epidemic data. *Communications Physics*, 8(1):450, 2025. doi: [10.1038/s42005-025-02358-w](https://doi.org/10.1038/s42005-025-02358-w). URL <https://doi.org/10.1038/s42005-025-02358-w>.
- Rajan Patel, Ira M. Longini Jr, and M. Elizabeth Halloran. Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *Journal of theoretical biology*, 234(2):201–212, 2005. URL <https://doi.org/10.1016/j.jtbi.2004.11.032>.
- Dimitar Prodanov. Chapter 10 - analytical solutions and parameter estimation of the sir epidemic model. In Praveen Agarwal, Juan J. Nieto, and Delfim F.M. Torres (eds.), *Mathematical Analysis of Infectious Diseases*, pp. 163–189. Academic Press, 2022. ISBN 978-0-323-90504-6. doi: <https://doi.org/10.1016/B978-0-323-90504-6.00015-2>. URL <https://www.sciencedirect.com/science/article/pii/B9780323905046000152>.
- Giulia Pullano, Laura Di Domenico, Chiara E. Sabbatini, Eugenio Valdano, Clément Turbelin, Maëlle Debin, Maroussia Guerisi, Chantal Kengne-Kuetche, Céline Souty, Thomas Hanslik, Thierry Blanchon, Julien Figoni, Séverine Vaux, Christine Campèse, and Vittoria Colizza. Underdetection of cases of COVID-19 in France threatens epidemic control. *Nature*, 590(7844): 134–139, 2021. URL <https://doi.org/10.1038/s41586-020-03095-6>.

- Nicholas G. Reich, Logan C. Brooks, Spencer J. Fox, Sasikiran Kandula, Chris J. McGowan, Evan Moore, Dave Osthus, Evan L. Ray, Abhinav Tushar, Teresa K. Yamana, Matthew Biggerstaff, Michael A. Johansson, Roni Rosenfeld, and Jeffrey Shaman. A collaborative multiyear, multi-model assessment of seasonal influenza forecasting in the United States. *Proceedings of the National Academy of Sciences*, 116(8):3146–3154, 2019. URL <https://doi.org/10.1073/pnas.1812594116>.
- Marcel Salathé. Digital epidemiology: what is it, and where is it going? *Life Sciences, Society and Policy*, 14(1):1–5, 2018. URL <https://doi.org/10.1186/s40504-017-0065-7>.
- Marcel Salathe, Linus Bengtsson, Todd J Bodnar, Devon D Brewer, John S Brownstein, Caroline Buckee, Ellsworth M Campbell, Ciro Cattuto, Shashank Khandelwal, Patricia L Mabry, et al. Digital epidemiology. *PLoS Computational Biology*, 8(7):e1002616, 2012. doi: 10.1371/journal.pcbi.1002616. URL <https://doi.org/10.1371/journal.pcbi.1002616>.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839): 604–609, 2020. doi: 10.1038/s41586-020-03051-4. URL <https://doi.org/10.1038/s41586-020-03051-4>.
- Lauren J. Shiman, Fatoumata Diallo, Christina I. Nieves, Brandon Brooks, Rachel Dannefer, Sheena Dorvil, Maria Lejano, and Jennifer Pierre. “Be honest and gain trust”: a population health study to understand the factors associated with building trust in local government related to COVID-19 and vaccination in three historically disinvested neighborhoods in New York City. *Frontiers in Public Health*, 11:1285152, 2023. doi: 10.3389/fpubh.2023.1285152. URL <https://doi.org/10.3389/fpubh.2023.1285152>.
- David Silver, Satinder Singh, Doina Precup, and Richard S Sutton. Reward is enough. *Artificial intelligence*, 299:103535, 2021. URL <https://doi.org/10.1016/j.artint.2021.103535>.
- Thomas Smith and Allan Schapira. Reproduction numbers in malaria and their implications. *Trends in Parasitology*, 28(1):3–8, 2012. ISSN 1471-4922. doi: <https://doi.org/10.1016/j.pt.2011.09.002>. URL <https://www.sciencedirect.com/science/article/pii/S1471492211001668>.
- Yiqi Su, Ray Lee, Jiaming Cui, and Naren Ramakrishnan. How (not) to hybridize neural and mechanistic models for epidemiological forecasting. *arXiv preprint*, 2026a. URL <https://arxiv.org/abs/2602.06323>.
- Yiqi Su, Christo Kurisummoottil Thomas, Walid Saad, Bud Mishra, and Naren Ramakrishnan. Modeling epidemiological dynamics under adversarial data and user deception. *arXiv preprint*, 2026b. URL <https://arxiv.org/abs/2602.20134>.
- Farzaneh Sadat Tabataba, Prithwish Chakraborty, Naren Ramakrishnan, Srinivasan Venkatramanan, Jiangzhuo Chen, Bryan Lewis, and Madhav Marathe. A framework for evaluating epidemic forecasts. *BMC infectious diseases*, 17(1):345, 2017. URL <https://doi.org/10.1186/s12879-017-2365-1>.
- Byomkesh Talukder, Jochen E. Schubert, Mohammadali Tofighi, Patrick J. Likongwe, Eunice Y. Choi, Gibson Y. Mphepo, Ali Asgary, Martin J. Bunch, Sosten S. Chiotha, Richard Matthew, Brett F. Sanders, Keith W. Hipel, Gary W. vanLoon, and James Orbinski. Complex adaptive systems-based framework for modeling the health impacts of climate change. *The Journal of Climate Change and Health*, 15:100292, 2024. ISSN 2667-2782. doi: <https://doi.org/10.1016/j.joclim.2023.100292>. URL <https://www.sciencedirect.com/science/article/pii/S2667278223000913>.
- Matthew W. Tanner, Lisa Sattenspiel, and Lewis Ntamo. Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical biosciences*, 215(2): 144–151, 2008. URL <https://doi.org/10.1016/j.mbs.2008.07.006>.

- Juliana C Taube, Zachary Susswein, and Shweta Bansal. Spatiotemporal trends in self-reported mask-wearing behavior in the united states: Analysis of a large cross-sectional survey. *JMIR Public Health Surveill*, 9:e42128, Mar 2023. ISSN 2369-2960. doi: 10.2196/42128. URL <https://publichealth.jmir.org/2023/1/e42128>.
- Jayaraman J. Thiagarajan, Rushil Anirudh, Peer-Timo Bremer, and et al. Machine learning-powered mitigation policy optimization in epidemiological models. In *Proceedings of the 1st Workshop on Healthcare AI and COVID-19, ICML 2022*, volume 184 of *Proceedings of Machine Learning Research*, pp. 63–72. PMLR, 2022. URL <https://proceedings.mlr.press/v184/thiagarajan22a.html>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint*, 2023. URL <https://arxiv.org/abs/2307.09288>.
- Melissa Tracy, Magdalena Cerdá, and Katherine M. Keyes. Agent-based modeling in public health: Current applications and future directions. *Annual Review of Public Health*, 39:77–94, 2018. doi: 10.1146/annurev-publhealth-040617-014317. URL <https://doi.org/10.1146/annurev-publhealth-040617-014317>.
- Guancheng Wan, Zewen Liu, Max S. Y. Lau, B. Aditya Prakash, and Wei Jin. EARTH: Epidemiology-Aware Neural ODE with Continuous Disease Transmission Graph. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=Cnfogmxymj>.
- Siyu Wang, Xiaocong Chen, and Lina Yao. On causally disentangled state representation learning for reinforcement learning based recommender systems. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, pp. 2390–2399, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704369. doi: 10.1145/3627673.3679674. URL <https://doi.org/10.1145/3627673.3679674>.
- Frank Wood, Andrew Warrington, Saeid Naderiparizi, Christian Weilbach, Vaden Masrani, William Harvey, Adam Ścibior, Boyan Beronov, John Grefenstette, Duncan Campbell, et al. Planning as inference in epidemiological dynamics models. *Frontiers in Artificial Intelligence*, 4:550603, 2022. doi: 10.3389/frai.2021.550603.
- Senlin Wu, Amanda N. Mertens, Yulin S. Crider, Anh Nguyen, Nattinee N. Pokpongkiat, Samuel Djajadi, Ankur Seth, Solomon M. Hsiang, John M. Colford, Arthur Reingold, Benjamin F. Arnold, Alan Hubbard, and Jade Benjamin-Chung. Substantial underestimation of SARS-CoV-2 infection in the United States. *Nature Communications*, 11:4507, 2020a.
- Shu Liang Wu, Andrew N. Mertens, Yulin S. Crider, An Nguyen, Natcha N. Pokpongkiat, Sinta Djajadi, Aaron Seth, Matthew S. Hsiang, John M. Colford, Arthur Reingold, Benjamin F. Arnold, Alan Hubbard, and Joshua Benjamin-Chung. Substantial underestimation of SARS-CoV-2 infection in the united states. *Nature Communications*, 11(1):4507, 2020b. doi: 10.1038/s41467-020-18272-4. URL <https://doi.org/10.1038/s41467-020-18272-4>.
- Wei Yuan, Brett K. Beaulieu-Jones, Kun-Hsing Yu, et al. Temporal bias in case-control design: preventing reliable predictions of the future. *Nature Communications*, 12:1107, 2021. doi: 10.1038/s41467-021-21390-2. URL <https://doi.org/10.1038/s41467-021-21390-2>.
- Li Zhang, Wenqiang Guo, and Chenrui Lv. Modern technologies and solutions to enhance surveillance and response systems for emerging zoonotic diseases. *Science in One Health*, 3: 100061, 2024. ISSN 2949-7043. doi: <https://doi.org/10.1016/j.soh.2023.100061>. URL <https://www.sciencedirect.com/science/article/pii/S2949704323000550>.