



THE 2ND WORKSHOP ON WORLD MODELS

Toward World Models for Epidemiology

Zeeshan Memon^{*,1}, Yiqi Su^{*,2}, Christo Thomas³, Walid Saad², Liang Zhao¹, Naren Ramakrishnan²

- Equal contribution ¹Emory University ² Virginia Tech ³ Worcester Polytechnic Institute

April 2026

Why World Models for Epidemiology?

- Epidemiology has evolved from compartmental models to ML-augmented forecasting, yet **forecast accuracy \neq policy usefulness**
 - Humans are **adaptive agents**; behavior shifts under interventions
 - Observations are **endogenous**: policy \rightarrow behavior \rightarrow data \rightarrow policy
 - Policy-making requires **counterfactual reasoning** (“what if?”)
- Forecasting alone cannot answer these questions
- **World models (WMs) provide the solution** by separating latent state from noisy observations, enabling simulation for counterfactual analysis

WMs in Computational Epidemiology

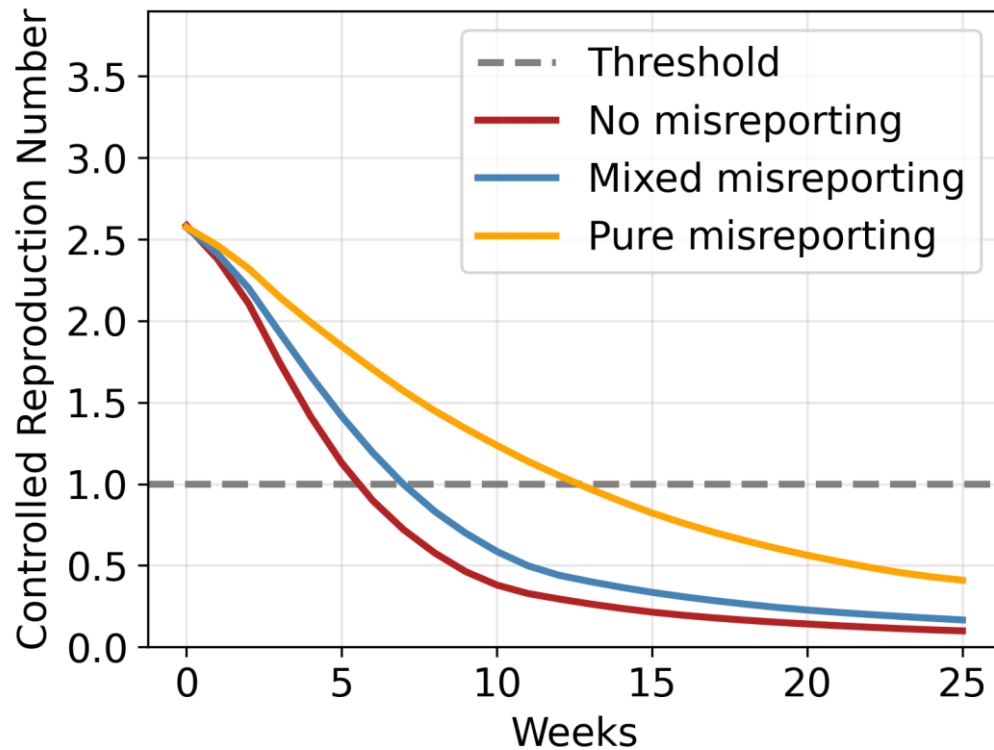
Core idea: **Epidemics as controlled partially observed dynamical systems**

Ingredient	Equation	Interpretation
Latent state	$x_t = (x^{epi}, x^{imm}, x^{net}, x^{beh}, x^{reg})$	Infection, immunity, contact, behavior & regime
Transition	$x_{t+1} \sim P_\theta(x_{t+1} x_t, a_t)$	Latent epidemic state under interventions
Observation	$o_t \sim \Omega_\theta(o_t x_t, a_{t-1})$	Policy-dependent surveillance signals
Belief state	$h_t = \phi(o_{1:t-1}, a_{1:t-1})$	Belief state from past observation & actions
Action	$a_t \sim \pi_\psi(\cdot h_t)$	Proposed actions on current belief state
Dynamics	$P_\theta(x_{t+1} x_t, \eta(x^{beh}, a_t))$	Indirect effects of actions through behaviors

- **Dynamics are action-mediated:** interventions affect transmission via behaviors
- **Observations are policy-dependent:** the same latent state produces different signals under different testing/reporting regimes

Motivating Challenges

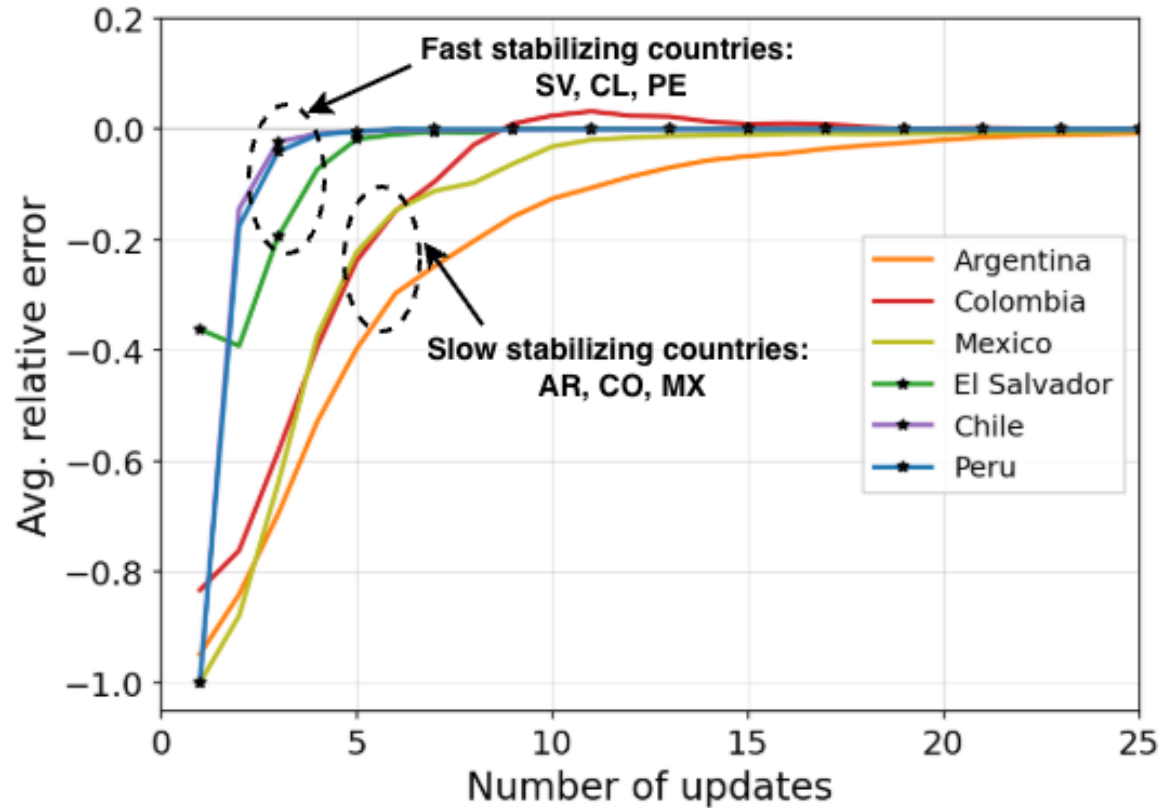
(a) Strategic misreporting in behavioral data



- Self-reported compliance is strategically distorted
- Misreporting delays suppression by biasing policy decisions
- ✓ WM separates latent behavior from distorted observations

Motivating Challenges

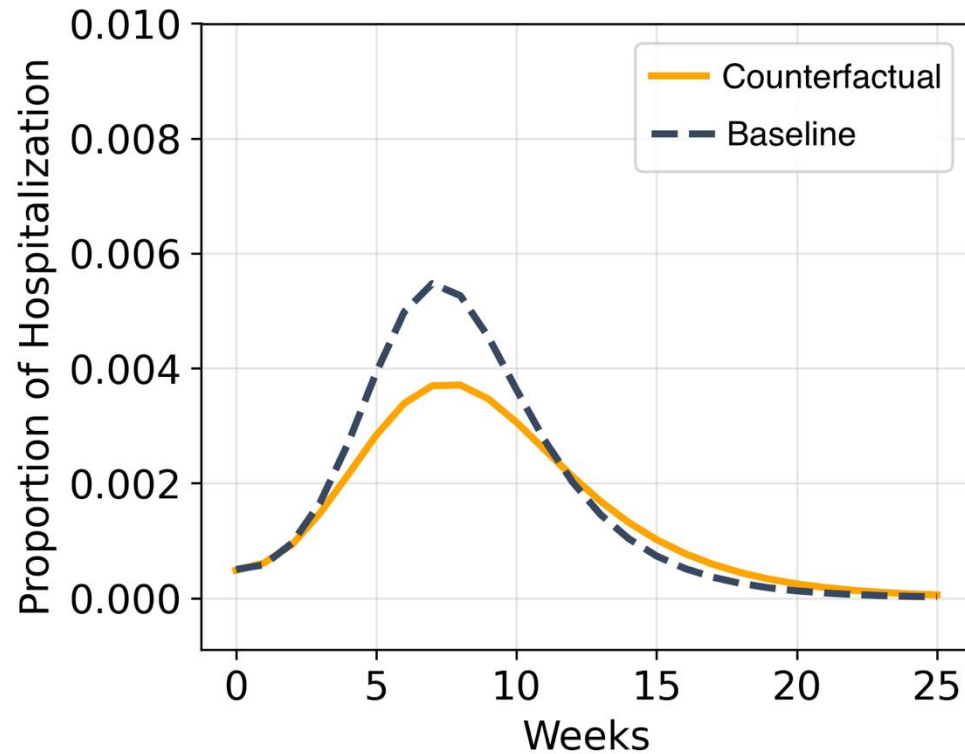
(b) Time-lagged surveillance signals



- Surveillance pipelines introduce systematic lag
- Forecasts conditioned on outdated signals
- ✓ WM infers latent state from delayed observations

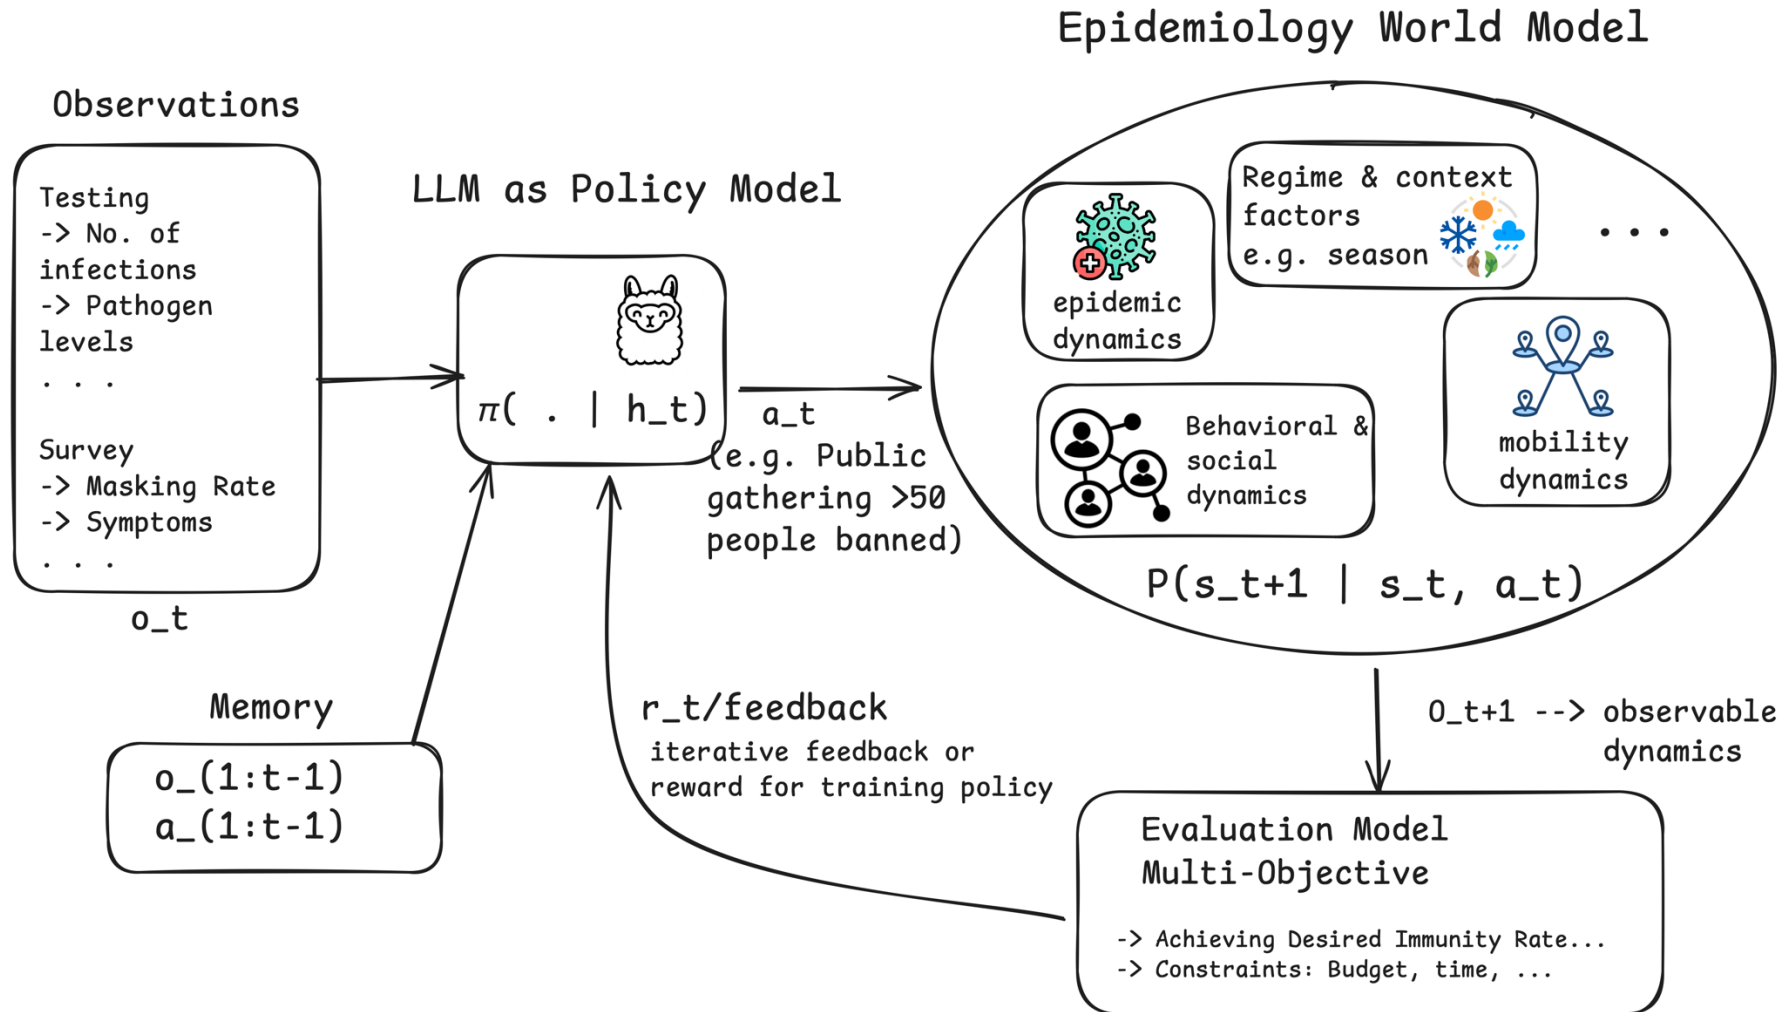
Motivating Challenges

(c) Counterfactual intervention analysis



- Identical histories diverge under alternative policy actions
- Forecasting cannot answer "what if we acted differently?"
- ✓ WM enables intervention-aware counterfactual simulation

Epidemic WM Framework



Epidemic WM Framework

- **Learning & optimization**

- **World model objective** (variational):

$$L_{WM} = \sum_{t=1}^T q_{\phi}(x_t) [\log \Omega_{\theta}(o_t | x_t)] - \beta KL(q_{\phi}(x_{1:T}) | p_{\theta}(x_{1:T} | a_{1:T}))$$

- **Policy optimization**

$$\pi^* = \operatorname{argmax}_{\pi} E \left[\sum_{\tau=t}^{t+H} r(x_{\tau}, a_{\tau}) \right]$$

- **Rollouts** simulate H-week forward trajectories

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

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

Epidemic WM Framework

- **Two optimization strategies**

- **Iterative Feedback:**

- WM outcomes rewrite the LLM policy prompt (no weight updates)

- **GRPO:**

- policy gradient using hospitalization reduction as reward signal

Experimental Settings

Component	Implementation
Data	Weekly COVID-19, 50 US states, 2021–2022
Policy (Actions)	13-dim OxCGRT NPIs
Horizon (H)	6 weeks
World Model	CovidLLM (LLaMA-2-7B + LoRA + GRU)
Policy Model	GPT-4o-mini / Qwen-7B

Experimental Results

Category	Method	Alignment (%)	Hosp. Reduction (%)
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

- Both LLMs exceed historical policy performance (37.1% reduction)
- Iterative feedback: best overall prompt refinement via WM rollouts
- GRPO: gradient-based optimization achieves comparable reduction

Takeaways

- Epidemics are partially observed control systems, well-suited for world models
- Forecasting alone cannot support policy reasoning; it fails under misreporting, delays, and counterfactual needs
- World models disentangle latent state from noisy, policy-driven observations
- Epidemiological world models enable intervention-aware simulation and reliable decision support

Thank You!

