Tufts Approximate Bayesian Computation for an Explicit-Duration Hidden Markov Model of COVID-19 Hospital Trajectories **Tufts** Center

Gian Marco Visani^{*}, Alexandra Hope Lee^{*}, Cuong Nguyen^{*}, David M. Kent¹, John B. Wong², Joshua T. Cohen³, Michael C. Hughes^{*}

* Department of Computer Science, Tufts University, Medford, MA

2 Center for Evaluation of Value and Risk, Tufts Medical Center

1 Predictive Analytics and Comparative Effectiveness Center, Tufts Medical Center 3 Division of Clinical Decision Making, Tufts Medical Center

Goals

What: We wish to model demand for hospital resources (daily usage of general ward beds, ICU beds, and ventilators) during the COVID-19 pandemic in an **interpretable** and **broadly applicable** fashion.

Why:

- To inform decision-makers of future demand
- To assess the societal value of possible interventions.

How:

- We created a new model, the ACED-HMM, to model individual hospital trajectories with a compact parameterization → interpretability
- We used Approximate Bayesian Computation to fit a posterior distribution over model parameters given daily census count data from a specific site/region without patient-specific data. → broad applicability



Duration model. Sample duration Δ for stage k and health status h as:

$$\Delta|k,h \sim Cat(\pi_1^{k,h},\ldots,\pi_D^{k,h})$$

where vector $\pi^{k,h}$ defines a categorical probability mass function over 1, 2, ... D days.

We use a simple two-parameter formulation with mode $\lambda^{k,h}$ and temperature $\nu^{k,h}$ $\begin{bmatrix} \pi_1^{k,h}, \dots, \pi_p^{k,h} \end{bmatrix} \leftarrow \operatorname{softmax} \left(\frac{\log \operatorname{PoiPMF}(1 \mid \lambda^{k,h})}{\nu^{k,h}}, \dots, \frac{\log \operatorname{PoiPMF}(D \mid \lambda^{k,h})}{\nu^{k,h}} \right)$



Priors. All parameters (transition probabilities and duration modes/temperatures) are given informative priors to match statistics released by the US CDC in Sept. 2020.

Acknowledgments

The authors gratefully acknowledge funding support from a consortium of industry partners for this project – AstraZeneca, Bristol-Myers Squibb, Eli Lilly, Gilead Sciences, Janssen, Merck, Pfizer, and Regeneron – via an agreement with the Center for the Evaluation of Value and Risk in Health (CEVR) at Tufts Medical Center.

Approximate Bayesian Computation (ABC)

Purpose of ABC. Estimate a posterior over all ACED-HMM parameters given daily counts from a training period of a target site (a single hospital or all sites in a region).

Why ABC? Easier to specify distance function than a likelihood. Support extensions to model (simulator) without requiring new derivations/implementation.

Algorithm [1]: Let Θ be the set of parameters that specify the behavior of ACED-HMM. Let α be the annealing constant for ε. Initialize ε ← 1.0. Let d(y_{1:T}, y_{1:T}) ∈ [0, 1] be the distance function between the ground truth census counts (y_{1:T}), and the simulated census count (ỹ_{1:T}), over T timesteps. Initialize each θ in Θ as a sample from its own prior distribution π_θ(·).

Repeat until converged:

- S0. For each parameter θ in Θ :
 - S1. Propose a move to θ' according to a parameter-specific transition kernel $q_{\theta}(\theta \rightarrow \theta')$. Let $\Theta' \leftarrow \Theta[\theta \rightarrow \theta']$.
 - S2. Generate $\tilde{\mathbf{y}}_{1:T}$ using ACED-HMM with Θ' .
 - S3. If $d(\mathbf{y}_{1:T}, \tilde{\mathbf{y}}_{1:T}) \le \varepsilon$ go to S4, and otherwise stay at Θ and go to S6. S4. Calculate

$$a = h(\Theta, \Theta') = \min(1, \frac{\pi_{\theta}(\theta')q_{\theta}(\theta' \to \theta)}{\pi_{\theta}(\theta)q_{\theta}(\theta \to \theta')})$$

- S5. With probability h, accept $\Theta' (\Theta \leftarrow \Theta')$, and set $\mathbf{d}_{best} \leftarrow d(\mathbf{y}_{1:T}, \tilde{\mathbf{y}}_{1:T})$. Otherwise stay at Θ .
- S6. Update ε according to: $\varepsilon \leftarrow \max(\mathbf{d}_{best}, \alpha \varepsilon)$. Then go to S1.
- We define the distance as a normalized weighted mean-absolute-error:

$$d(\mathbf{y}_{1:T}, \tilde{\mathbf{y}}_{1:T}) = \frac{1}{K \cdot T} \sum_{t=1}^{T} \sum_{k=1}^{K} w_{tk} \frac{|y_t^k - \tilde{y}_t^k|}{\max(y_t^k, \tilde{y}_t^k)}$$

where w_{tk} is a scalar weight that determines the relative "importance" of matching the counts at time t and stage k.

Forecasts on MA

Training window: Nov. 11, 2020– January 11, 2021 Forecasting window: January 12 – February 11 2021 We assume we know the true hospital admissions in the forecasting window.



For ACED-HMM, we report the 2.5th and 97.5th percentiles of the MAE across multiple sets of samples from the posterior. For HME, as a distribution of forecasted samples is not available, we report the MAE computed using the mean estimate forecast, as well as the MAE computed using the provided lower and upper estimates count values. This second kind of interval (marked with *) should not be directly compared to the first, as they capture different aspects of uncertainty.

Posterior Distribution after training on MA



Retrospective What-If Scenarios on CA

In late 2020, hospitals in California experienced a surge that led to aggressive lockdowns.

We explore our model's potential to answer what-if questions about how different hypothetical pharmaceutical treatments rolled out on November 2nd (dashed line) might have helped avoid the surge in counts in CA hospitals that led to lockdown. (Note: our model is *not a causal model*).

Scenario 1 (top plot): Reduction in admitted patients by 25%, 50% and 87%, simulating the use of two drugs: bamlanivimab and etescwimab [2]. The therapy is assumed to be administered to 100% of people who test positive for COVID-19 and who are deemed at high risk of hospitalization, with a linear rame-up schedule of 30 days.



Scenario 2: 25% decrease in length-of-stay for recovering patients, simulating the use of Remdesivir [3].



References

[1] P. Marjoram et al. "Markov chain Monte Carlo without likelihoods." Proceedings of the National Academy of Sciences 100 (26), 2003.

(2) Eli Lilly and Company. Lilly's bamlanivimab and etesevimab together reduced hospitalizations and death in Phase 3 trial for early COVID-19. 2021. https://investor.lilly.com/news-releases/news-telease-deatals/lillys-bamlanivimab-and-etesevimab-together-reduced

[3] J. H. Beigel et al. "Remdesivir for the Treatment of Covid-19 — Final Report". New England Journal of Medicine, 383 (19), 2020.