

A Dynamic Human-in-the-loop Recommender System for Evidence-based Clinical Staging of COVID-19

Yogatheesan Varatharajah
University of Illinois at Urbana-Champaign
varatha2@illinois.edu

Andrew Trotter
University of Illinois at Chicago
trottera@uic.edu

Haotian Chen
University of Illinois at Urbana-Champaign
hc19@illinois.edu

Ravishankar Iyer
University of Illinois at Urbana-Champaign
rkiyer@illinois.edu

ABSTRACT

In this position paper, we discuss the potential use of a reinforcement learning (RL)-based human-in-the-loop recommender system to support clinical management of COVID-19. COVID-19 is a disease of extraordinary complexity that even the most experienced clinicians are struggling to understand. There is an urgent need for an evidence-based model for predicting the severity of the COVID-19 disease and its complications that can guide individual clinical management decisions. Such a model will utilize a diverse set of information to determine a patient's disease severity and associated risk of complications. An immediate application would be a clinical protocol tailored for COVID-19 patient care; this is a critical need both today and for future studies of potential treatments.

CCS CONCEPTS

• **Human-centered computing** → *Human computer interaction (HCI)*; • **Applied computing** → *Health care information systems*.

KEYWORDS

COVID-19; reinforcement learning; human-in-the-loop; staging

ACM Reference Format:

Yogatheesan Varatharajah, Haotian Chen, Andrew Trotter, and Ravishankar Iyer. 2020. A Dynamic Human-in-the-loop Recommender System for Evidence-based Clinical Staging of COVID-19. In *5th International Workshop on Health Recommender Systems co-located with 14th ACM Conference on Recommender Systems (HealthRecSys'20)*, Online, Worldwide, September 26, 2020. , 2 pages.

1 INTRODUCTION

The emergence of the Severe Acute Respiratory Syndrome Coronavirus (SARS-CoV-2) poses significant challenges to the livelihood of the affected nations and, in the absence of directed treatment or a vaccine, requires drastic public health measures which have crippled national and international economies [2]. Preliminary data has shown that there is a spectrum of disease severity for which the disease mechanisms, patient characteristics, and risk factors are poorly understood. There is an urgent need for an evidence-based model to predict the severity of COVID-19 disease and its complications which can guide individual clinical management decisions

and anticipate system-level allocation of health resources [3]. Such a model will require consideration of diverse patient characteristics and clinical variables in order to determine the patient's disease severity and associated risk of complications including death. It also could be applied to estimate the demands placed on the medical and staff resources and the definition of a disease staging system would be a critical tool in future studies of potential treatments [1]. At present, patients are triaged predominantly using clinical assessments based on other respiratory illnesses and may not accurately reflect the trajectories they may follow under COVID-19. This disease is very new and there is a scarcity of research defining risk factors for severe disease or methods to predict patients at risk for rapid decline in their health. It is critical to develop dynamically evolving analytical tools that can make accurate recommendations using limited and readily available baseline data. These analytic tools should also adaptively incorporate new information prospectively from current encounters in order to clinically stage disease severity at baseline and throughout disease progression.

2 SYSTEM OVERVIEW

Our proposed system (shown in Figure 1) is based on a human-in-the-loop RL algorithm that leverages expert knowledge of clinical experts and data-driven analytics. Our system will operate as follows. Consider a situation in which the algorithm is challenged with a patient who presents to the emergency department with a defined set of symptoms, laboratory variables, clinical measurements and imaging results. The learning algorithm will be able to provide an estimate of the patient's probability of experiencing serious complications (such as requiring mechanical ventilation or death) using the patient's baseline characteristics. Based on this prognosis, a decision algorithm would recommend admission or discharge home and if admitted, the level of medical care required (e.g., general ward, step down unit, intensive care unit). However, the final decision regarding the level of hospital care and administration of supportive and directed treatments (such as anti-viral drugs) will be determined by a clinical expert, using both the algorithm's recommendation and his/her own clinical judgment (a human-in-the-loop ML model). The individual patient's clinical outcome (e.g., need for ventilator support, symptom severity, time spent in ICU, treatment response, recovery or death, side effects) will be used to reinforce the prognostic algorithm. The framework will adapt to continuously refine decisions based on new data and expert-clinician reinforcement.




Figure 1: A Dynamic ML-based Clinical Staging Scheme for COVID-19.

3 CHALLENGES

There are several challenges in developing a successful human-in-the-loop reinforcement-learning framework that generalizes across the entire disease severity spectrum.

Extracting actionable intelligence from heterogeneous and incomplete data: Owing to the complexity of COVID-19, the identification of distinct clinical stages of COVID-19 progression and patient trajectories requires the integration of multiple data sources. We believe that domain-guided models that integrate machine learning methods and clinical insights will be beneficial. Specifically, probabilistic graphical models, can represent the domain-driven relationships between different information sources, and can be transformed into discriminative models that can be trained using the available data. The goals are to improve outcomes, appropriately allocate healthcare resources, and reduce mortality rates while directed treatments and vaccines are being developed.

Quantifying the uncertainty in model predictions: Since data are limited in the beginning, uncertainty in the prognostics will be high in the early stages and will gradually decrease as the model is updated using new data. The ability to quantify such uncertainty is critical in order for clinicians to accurately gauge the importance of their own assessments relative to the model's predictions. We recommend the use of Bayesian methods and attribution-based approaches to quantify the uncertainty and interpret model predictions, respectively, both of which will inform the clinicians.

Time-frames for model development and reinforcement: It is unclear how many data points are required to train an accurate initial model. A potential measure that can guide this decision is the convergence of class probabilities to their respective population means. In addition, there have been multiple reports indicating varied lengths of hospitalizations. Such variability depends on multiple factors including disease severity, comorbidities, health provider policies, decisions regarding withdrawal of life support, cost of treatments, etc. Clearly, some of these factors are non-deterministic and data on such factors are typically unavailable for the machine learning model. Learning a sufficiently accurate and robust decision scheme with those difficult-to-measure elements and determining when to reinforce the learning algorithm remain challenging tasks.

Modeling the human-in-the-loop decision process: A typical RL approach relies on an effective balance between exploration and exploitation such that the algorithm is allowed sufficient exploration of the input space prior to basing predictions primarily on the space that it has already explored. However, that paradigm is not usable in this setting, because treatment decisions are a matter of life and death; we cannot take actions that would jeopardize medical ethics. Therefore, our approach requires the presence of a clinical expert who will make decisions after appraising the model's predictions in light of his/her own assessments.

4 CONCLUSION

In this paper, we described a novel domain-guided human-in-the-loop RL framework to assist physicians in clinical decision-making to stage COVID-19 patients across the disease severity spectrum. Going forward, the clinical stages as defined by this approach could form the basis for evaluating the efficacy of existing and new drugs related to the patients in different stages of disease progression. While the proposed model is specifically designed and trained for COVID-19, the underlying paradigm of our model, i.e., the human in the loop RL, affords the adaptivity to be applicable to other respiratory illnesses and other future pandemics, with re-calibration.

5 ACKNOWLEDGEMENTS

This project has been funded by the Jump ARCHES endowment through the Health Care Engineering Systems Center at the University of Illinois.

REFERENCES

- [1] Ezekiel J. Emanuel et al. 2020. Fair Allocation of Scarce Medical Resources in the Time of Covid-19. *New England Journal of Medicine* 382, 21 (2020), 2049–2055. <https://doi.org/10.1056/NEJMs2005114>
- [2] Center for Disease Control. 2020. Severe Outcomes Among Patients with Coronavirus Disease 2019 (COVID-19) – United States, February 12–March 16, 2020. *Morb Mortal Wkly Rep* 2020 69 (2020), 343–346. <https://doi.org/10.15585/mmwr.mm6912e2>
- [3] Hasan K Siddiqi and Mandeep R Mehra. 2020. COVID-19 illness in native and immunosuppressed states: A clinical–therapeutic staging proposal. *The Journal of Heart and Lung Transplantation* 39, 5 (2020), 405. <https://doi.org/10.1016/j.healun.2020.03.012>