Considerations, Good Practices, Risks and Pitfalls
in Developing AI Solutions Against COVID-19

Alexandra Luccioni[1,2] , **Joseph Bullock**[3,4] , **Katherine Hoffmann Pham**[3,5] ,
Cynthia Sin Nga Lam[6] and **Miguel Luengo-Oroz**[3]

[1] Université de Montréal [2] Mila Québec Artificial Intelligence Institute
[3] United Nations Global Pulse [4] Institute for Data Science, Durham University
[5] NYU Stern School of Business [6] Global Coordination Mechanism on NCDs, WHO
sasha.luccioni@mila.quebec, {joseph,katherine,miguel}@unglobalpulse.org, lams@who.int

Abstract

The COVID-19 pandemic has been a major challenge to humanity, with 12.7 million confirmed cases as of July 13th, 2020 [1]. In previous work, we described how Artificial Intelligence can be used to tackle the pandemic with applications at the molecular, clinical, and societal scales [2]. In the present follow-up article, we review these three research directions, and assess the level of maturity and feasibility of the approaches used, as well as their potential for operationalization. We also summarize some commonly encountered risks and practical pitfalls, as well as guidelines and best practices for formulating and deploying AI applications at different scales.

1 Introduction

In recent months, an already remarkable body of research on potential applications of Artificial Intelligence (AI) and Machine Learning (ML) to address the health and societal challenges created by the COVID-19 pandemic has been produced. In a recent survey of the literature [2], we identified over 300 relevant journal articles and preprints with applications ranging from tracking the spread of misinformation to molecular modeling of the SARS-CoV-2 virus. These applications are emerging in a wide array of settings, increasingly driven by the priorities of responders and multidisciplinary collaborations.

In this article, we present a more in-depth analysis of AI applications at the molecular, clinical, and societal scales. We build on our previous work by sharing some high-level observations on the broader landscape of AI for COVID-19 and providing a critical take on best practices and application strategies. While the speed and scale of the AI community's research response to the pandemic is unprecedented, we feel nonetheless that there are a number of common, structural challenges faced by many of these projects. We extract shared design considerations and promising directions, as well as recurring problems, risks, and areas of improvement that could help researchers and funding bodies when developing further applications and research agendas.

2 Molecular Scale

Promising Directions AI has been applied in a wide spectrum of molecular research – ranging from better understanding the structure of the SARS-CoV-2 virus to assisting drug development and improving molecular diagnosis.

Some of the earliest responses to the pandemic were produced by AI researchers leveraging existing protein structure prediction algorithms [3] and drug discovery pipelines [4], or exploring pre-built knowledge graphs [5]. Having the infrastructure for this research already in place facilitated a rapid response to the new challenge of COVID-19. Open science, based on previously published peer-reviewed papers detailing the methodologies involved, has helped to accelerate the evaluation of these approaches.

A number of projects studied drug re-purposing, in the hope of discovering a therapy for which the engineering of novel compounds is not needed, clinical trials have been authorized, or the compound's use has already been approved (e.g., [6]). Others used ML approaches to reduce the computational burden of docking simulations by narrowing down the set of candidate compounds which needed to be docked (e.g., [7]). Another notable trend in molecular applications involves identifying desirable or undesirable properties of a candidate compound – e.g., novelty, drug likeness, or toxicity – and then training drug discovery models to suggest compounds that meet these desiderata (e.g., [8]). These applications allowed creative exercises from the modeling perspective, in which researchers experimented with different representations of the same data to design different AI pipelines.

Many papers relied on a common set of open datasets such as PDBbind, ChemBL, and DrugBank, which makes their methodologies more accessible and replicable for other researchers. Several even posted models or results on Github or other web pages, providing open access to all (e.g., [9]).

Risks and Pitfalls Many ML applications in molecular science remain at the research level. It is understandable that molecular research is part of a multidisciplinary and multistep process; nonetheless, few of the efforts we observed had advanced through the drug development pipeline to more formal evaluation. One challenge is that synthesizing and testing compounds is costly and time-consuming, so researchers who are not embedded within a larger infrastructure are often unable to execute these later steps required to translate research
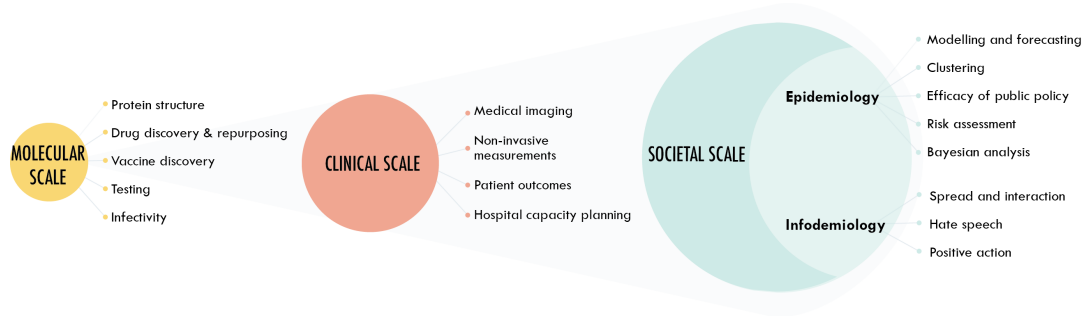
Figure 1: AI applications for the COVID-19 response organized at three levels: the molecular scale, the clinical scale, and the societal scale.

to practice.

Another barrier in this research-to-practice process could be limitations in data and model sharing. To our knowledge, two candidate vaccines that reported the use of ML in their development have been approved for clinical evaluation [10]. However, both vaccines came from corporations which published limited information on their approach and the extent to which ML was utilized in vaccine development. More generally, we found that some of the research produced by private-sector entities relied on datasets such as knowledge graphs which are not publicly accessible, describing proprietary models only in vague or high-level terms. Hence, we would like to reiterate the importance of open science in ensuring accessible vaccines and treatments for vulnerable communities.

3 Clinical Scale

Promising Directions From a clinical perspective, AI and ML have already been used to assist and improve patient-level assessment of COVID-19. One major application for ML in this field has been the analysis of medical imagery such as CT and X-Ray scans with the help of common neural network architectures, either to provide a supplementary data point to corroborate COVID-19 diagnosis [11], or to assess the severity and progression of the disease [12]. Some of these approaches are now fully operational and have received institutional approval for deployment in hospitals as triage tools [13] or as human-in-the-loop systems for radiologists [14].

Approaches using complementary data sources have also been proposed, ranging from wearable devices and mobile phones for detecting symptoms [15], to electronic health records to improve diagnosis and outcome prediction. We find one of the most promising applications of AI at a clinical level to be the prediction of patient outcomes and the proposal of triage approaches based on features extracted from medical data. Such approaches are transparent and practicable, pinpointing key measurable features which enable hospitals to plan the use of resources such as ventilators and ICU beds.

Furthermore, there are promising hybrid studies that leverage both medical imagery and clinical features to predict patient-level characteristics such as the severity of COVID-19 [16]. These approaches leverage different complementary sources of data in order to make more precise and more generalizable predictions of patients' prognosis. In fact, the clinical

features identified by studies using solely clinical data were also corroborated by hybrid studies; e.g. both sets of studies found that high levels of substances such as lactic dehydrogenase (LDH) and high-sensitivity C-reactive protein (CRP) were correlated with mortality risk and longer hospitalization.

Risks and Pitfalls While many studies are being carried out in situ on COVID-19 patients, there is still much that we do not know about the virus itself and the factors that can put patients at risk for hospitalization, developing acute respiratory distress syndrome (ARDS), and eventual death from respiratory failure. For instance, the extent to which medical imagery alone can be used for the diagnosis of COVID-19 is still debated by the medical community [17]. Above and beyond the feasibility of diagnosis, many of the medical imaging papers we reviewed had methodological issues, relying on small and poorly-balanced datasets that mix data from several populations, coupled with flawed evaluation procedures [18]. Most also presented no plan for inclusion in clinical workflows and no attempt to provide a transparent explanation for the diagnosis, which is especially important in patient-level applications of AI [19]. Finally, while approaches that leverage ML to analyze non-invasive measurements are potentially promising given the ubiquity and accessibility of sensor technologies, we found that these approaches are not sufficiently mature to evaluate their performance. We would advocate more extensive testing and clinical investigations to validate their performance in deployment.

4 Societal Scale

Promising Directions From a societal perspective, AI has been applied to the field of epidemiological modelling, as well as to understanding and combating the "infodemic" spread of misinformation [20]. At the epidemiological level, many studies have sought to produce forecasting models for national and regional level statistics. While a vast body of literature on modeling already exists, AI based models could augment classical models in situations where analytic transmission equations are not well known, such as when modelling the effects of public policy measures such as social distancing and self-quarantine [21]. In addition, AI methods can also be used to incorporate new data sources, such as social media and search information [22]. Other works use AI to identify similarities and differences in the evolution of the pandemic between regions. These approaches have leveraged

both supervised and unsupervised techniques, and may help inform policy makers at a high level and highlight areas for more detailed exploration.

AI has also been applied to investigate the scale and spread of the infodemic in order to address the propagation of misinformation and disinformation including the emergence of hate speech. Given the vast amount of information now being disseminated and shared, there is a need for tools to help identify and promote reliable information sources, and understand the spread of misinformation. Promising work has analyzed patterns in the transmission of such information and developed infodemic risk scoring algorithms (e.g., [23]). Moreover, there has been an increasing focus on assessing the emergence of hate speech, particularly using network analysis techniques, which could help inform preventative efforts or contribute to the development of continuous monitoring platforms [24].

Risks and Pitfalls Policy decisions must be based on justifiable models which stand up to public scrutiny. Since much of the data collected for COVID-19 epidemiological modeling tasks is extremely limited, the choice of models and datasets can have significant effects on overall performance and models may lack generalizability. A significant limitation of many articles in this category is the heterogeneous data collection in different countries due to multiple factors including variations in testing, case tracking, and reporting quality and standards. Moreover, applying models trained in one context to another raises concerns surrounding the model's ability to capture aspects such as different cultural norms which may impact the spread and effect of the virus. In such examples, a transferred model will have to be tailored for local contexts given that there may be different demographic characteristics and behaviors. Indeed, developing proper model transferability procedures and guidelines is especially important for data poor regions. While synthetic data approaches have been proposed, these should be applied with caution.

In order to understand and tackle the infodemic, it is important to capture and analyze information from a diverse range of sources. While much of the work on this topic uses data from online sources such as social media and Google searches, information propagated through alternative channels such as radio is important for capturing wider trends. Moreover, many of the existing approaches rely on language modelling techniques developed for English or other widely spoken languages, but relying on such models might leave many populations behind, including some of those most vulnerable. Finally, we note that while numerous methods have been proposed for identifying hate speech, further research is needed to identify the targeted groups and to learn how to use the insights gained from these techniques in an effective way (e.g. see Section II [25]).

5 Discussion

We believe that there are several considerations to keep in mind when applying AI to a global problem such as the COVID-19 pandemic. These include:

1. *Application relevance and context*: Does the application make sense from both an application and a method-ological perspective? Were domain experts consulted to properly assess the needs and define the problem to be solved? Does the solution serve its target audience?

2. *Data availability and quality*: Has the quantity of data being used to train and evaluate the model been assessed and deemed to be of a reasonable size and diversity to justify the claims made? Has bias in the data been considered and documented? Were privacy measures taken?

3. *AI methodology and complexity*: Is the approach proposed justified? If an AI based model, has it been benchmarked against more traditional approaches? Has the approach been validated by other researchers and, if possible, peer-reviewed?

4. *Transparency and explainability*: Have efforts been made to render the results and the approach understandable by humans? Is it possible to identify the key features being used by the algorithm?

5. *Dissemination of knowledge*: Are the data, code and models being shared in any form? Are there reporting guidelines and standards that should be followed?

6. *Operationalization and performance*: Can the approach be integrated into decision-making workflows? What is the implementation plan? Are there regulatory frameworks that have to be taken into account? What are the potential risks? Will the models incorporate user feedback, and if so, how?

The COVID-19 pandemic is a global emergency that has overstretched health care networks and posed significant health, economic and social challenges to humanity. AI can play an important role in alleviating this pressure, but we would like to reiterate that in order for any technological solution to make an impact, it must be deployed contextually and appropriately. We advocate for aspiring research initiatives to be carried out in partnership with stakeholders who have the necessary domain knowledge. It is also important to investigate how proven solutions can be adapted to local contexts to address unmet needs, particularly in areas of the world with fewer resources. This requires developing appropriate model and data sharing solutions, and specific measures to address data scarcity (see e.g. [26] for more details).

Finally, given the rapidly changing nature of human understanding regarding the pandemic, and therefore the inability to fully validate many approaches, we suggest that models should not be designed to process data in an end-to-end fashion at this stage, but rather to augment human decision making. With careful attention to the implementation context and operational needs, we believe that AI solutions can be a valuable asset in the fight against the pandemic.

Acknowledgements

References

[1] WHO. Coronavirus disease (COVID-19) outbreak situation; 2020. https://www.who.int/emergencies/diseases/novel-coronavirus-2019.

[2] Bullock J, Luccioni A, Pham KH, Lam CSN, Luengo-Oroz M, et al. Mapping the landscape of artificial intelligence applications against COVID-19. arXiv preprint arXiv:200311336. 2020;.

[3] Jumper J, Tunyasuvunakool K, Kohli P, Hassabis D, AlphaFold Team. Computational predictions of protein structures associated with COVID-19; 2020.

[4] Zhavoronkov A, Aladinskiy V, Zhebrak A, Zagribelnyy B, Terentiev V, Bezrukov DS, et al. Potential COVID-2019 3C-like Protease Inhibitors Designed Using Generative Deep Learning Approaches. Insilico Medicine Hong Kong Ltd A. 2020;307:E1.

[5] Richardson P, Griffin I, Tucker C, Smith D, Oechsle O, Phelan A, et al. Baricitinib as Potential Treatment for 2019-nCoV Acute Respiratory Disease. The Lancet. 2020;395(10223):e30–e31.

[6] Nguyen D, Gao K, Chen J, Wang R, Wei G. Potentially highly potent drugs for 2019-nCoV. bioRxiv preprint bioRxiv:20200205936013v1. 2020;.

[7] Ton AT, Gentile F, Hsing M, Ban F, Cherkasov A. Rapid Identification of Potential Inhibitors of SARS-CoV-2 Main Protease by Deep Docking of 1.3 Billion Compounds. Molecular Informatics. 2020;.

[8] Tang B, He F, Liu D, Fang M, Wu Z, Xu D. AI-Aided Design of Novel Targeted Covalent Inhibitors against SARS-CoV-2. bioRxiv preprint bioRxiv:20200303972133. 2020;.

[9] Lopez-Rincon A, Tonda A, Mendoza-Maldonado L, Claassen E, Garssen J, Kraneveld AD. Accurate Identification of Sars-Cov-2 from Viral Genome Sequences Using Deep Learning. bioRxiv preprint bioRxiv:20200313990242v1. 2020;.

[10] WHO. Draft landscape of COVID 19 candidate vaccines; 2020. https://www.who.int/who-documents-detail/draft-landscape-of-covid-19-candidate-vaccines.

[11] Li L, Qin L, Xu Z, Yin Y, Wang X, Kong B, et al. Artificial intelligence distinguishes covid-19 from community acquired pneumonia on chest ct. Radiology. 2020;.

[12] Tang Z, Zhao W, Xie X, Zhong Z, Shi F, Liu J, et al. Severity Assessment of Coronavirus Disease 2019 (COVID-19) Using Quantitative Features from Chest CT Images. arXiv preprint arXiv:200311988. 2020;.

[13] Alaa AM, van der Schaar M. Autoprognosis: Automated clinical prognostic modeling via bayesian optimization with structured kernel learning. arXiv preprint arXiv:180207207. 2018;.

[14] Shan F, Gao Y, Wang J, Shi W, Shi N, Han M, et al. Lung Infection Quantification of COVID-19 in CT Images with Deep Learning. arXiv preprint arXiv:200304655. 2020;.

[15] Radin JM, Wineinger NE, Topol EJ, Steinhubl SR. Harnessing wearable device data to improve state-level real-time surveillance of influenza-like illness in the USA: a population-based study. The Lancet Digital Health. 2020;.

[16] Shi W, Peng X, Liu T, Cheng Z, Lu H, Yang S, et al. Deep Learning-Based Quantitative Computed Tomography Model in Predicting the Severity of COVID-19: A Retrospective Study in 196 Patients. 2020;.

[17] Weinstock MB RJea Echenique A. Chest x-ray findings in 636 ambulatory patients with COVID-19 presenting to an urgent care center: a normal chest x-ray is no guarantee. The Journal of Urgent Care Medicine. 2020; p. 13–18.

[18] Wynants L, Van Calster B, Bonten MM, Collins GS, Debray TP, De Vos M, et al. Prediction models for diagnosis and prognosis of COVID-19 infection: Systematic review and critical appraisal. bmj. 2020;369.

[19] Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. Nature medicine. 2019;25(9):1337–1340.

[20] WHO. Infodemic management - Infodemiology; 2020. https://www.who.int/teams/risk-communication/infodemic-management.

[21] Dandekar R, Barbastathis G. Neural Network aided quarantine control model estimation of COVID spread in Wuhan, China. arXiv preprint arXiv:200309403. 2020;.

[22] Lampos V, Moura S, Yom-Tov E, Edelstein M, Majumder M, Hamada Y, et al. Tracking COVID-19 using online search. arXiv preprint arXiv:200308086. 2020;.

[23] Gallotti R, Valle F, Castaldo N, Sacco P, Domenico MD. Assessing the risks of "infodemics" in response to COVID-19 epidemics. arXiv preprint arXiv:200403997. 2020;.

[24] Velásquez N, Leahy R, Restrepo NJ, Lupu Y, Sear R, Gabriel N, et al. Hate multiverse spreads malicious COVID-19 content online beyond individual platform control. arXiv preprint arXiv:200400673. 2020;.

[25] United Nations. United Nations Guidance Note on Addressing and Countering COVID-19 related Hate Speech; 2020. https://www.un.org/en/genocideprevention/documents/Guidance%20on%20COVID-19%20related%20Hate%20Speech.pdf.

[26] Wright J, Verity A. Artificial Intelligence Principles For Vulnerable Populations in Humanitarian Contexts. Digital Humanitarian Network. 2020;.