COVID Information Commons (CIC) Research Lightning Talk

Transcript of a Presentation by Ioannis Paschalidis (Boston University), April 24, 2023

Title: Predictive Models of COVID-19 Severity and Patient Outcomes

Ioannis Paschalidis CIC Database Profile

NSF Award #: 2200052

YouTube Recording with Slides

Spring 2023 CIC Webinar Information

Transcript Editor: Lauren Close

Transcript

Slide 1

Ioannis Paschalidis:

Thank you so much and thank you for inviting me to be part of this exciting afternoon with a great set of speakers. I will talk about predictive models of COVID-19 severity and patient outcomes that we have developed at Boston University.

Slide 2

Just to give a little bit of a context - we've been working for quite some time before the pandemic on a variety of models that predict disease and important key events, for instance, hospitalizations, and also models that prescribe treatments. So when the pandemic started, we mobilized, like the rest of the scientific community, and we also mobilized a relatively large network of collaborators so we obtained access to a variety of different data sets.

You can sort of see in this slide the various data sets we had. Data sets that were local from Massachusetts - from two different hospital networks in Massachusetts. One was from Mass General Brigham network, from five different hospitals about 2,500 cases and then another relatively large cohort from the Boston Medical Center of about 7,000 cases. We also got some cases from Wuhan, China which was the origin, obviously, of the epidemic. And finally we were able to get access to some large national datasets - so a data set from Brazil and another data set from Mexico.

In this very short presentation I will focus more on the most recent work that considered the dataset - the series of patients from the Boston Medical Center. Boston Medical Center is the teaching hospital affiliated with BU Medical School and is also a safetynet hospital. As you will see, this has some interesting implications in the findings that we were able to get.

Slide 3

So we got access to the entire 2020 BMC cohort of about 7,000 patients. Those were patients who tested positive for COVID-19. Just some rough statistics - about 20 percent were admitted. From those admitted, about 23 or so were admitted to the ICU. From those admitted to the ICU, about 60 or 58.7 percent were intubated. From those intubated, about 70 percent unfortunately did not make it. We had lots of information about these patients, including demographics, their vitals throughout the hospital stay, radiology reports, their medical history, any symptoms, any lab results, any medications, and even information about depression status, zip code, and also we had information about the occupancy of the hospital at the time that each one of these patients were seen. We also had information on social determinants of Health. BMC runs a program called the Thrive Program that everyone who has an encounter with the hospital is given a survey where we asked them about needs in a variety of different areas, including housing, food, transportation, help with caretaking, medications, helping getting access to medications and paying for medications, education, and employment. A lot of the information was in tabular format that is more easily handled by AI machine learning methods, but also there was quite a bit of information, particularly radiology reports and other reports during the hospital stay, that were just narratives - reports by clinicians. We had to use quite a bit of natural language processing in order to extract appropriate findings from the text. What we did is we developed a set of robust interpretable models that can predict hospitalization, ICU admission, mechanical ventilation, and death. I'll show you some examples. I will not be exhaustive, again in the interest of time.

Slide 4

First, what we did is for every patient we constructed a timeline from the time that we were aware of the patient testing positive and having information about this patient up to the point of the event of interest whether that was, let's say, an ICU admission or mechanical ventilation. Here is the outcome of interest, and then looking back in time, we created these data buckets. We dropped any information that was available just before the time of the event of interest becuase we wanted models that could predict what will happen into the future. Perhaps even a first-year medical student can identify a patient heading to the ICU, so we wanted the models to make that prediction with earlier information. You will see that we obtained different versions of the models that had different cutoffs in terms of the information that the model used in order to make the prediction. Another reason we created these time buckets is we wanted to capture the dynamic evolution of the progress of the patient while in the hospital, for instance, the dynamic evolution of the vitals. We understood that this was rather important. So rather than just looking at the snapshot of vitals, let's say at some specific time and using that information for making a forward prediction, the exact values are important but trends are important as well. Physicians, when they look at patients, they sort of look at the trends of the vitals in the patient. So we developed a model that used some fairly sophisticated deep learning methodologies, including

LSTM type of networks and a transformer architecture that took as input the vitals - six vital signals at different points in time - and produced a score. That score captured the dynamic evolution of the vitals and that vital score was then used in an ensemble model that was attempting to make a prediction for the outcome of interest.

Slide 5

For instance, hospitalization predictions. You can see that these are fairly accurate. Some of the best models give you a 92 percent - this is in area under the curve. You can think of this as a measure of accuracy of the model. The best is 100 percent, a random guess will give you 50 percent. So 92 percent is quite good performance. You can see here that from some linear models that we developed we also found some of the factors that were important in making the hospitalization prediction. In blue, you see some of the variables that are associated with some earlier health conditions that, for instance, are highly correlated with hospitalization. You will also see that the occupancy of the hospital, if it was high, reduced the likelihood that the patient is going to be hospitalized. Also, you will see two social determinants of health: need for food and need for transportation. These were both contributing to a hospitalization decision. Patients with those needs were more likely to be hospitalized. I would like to emphasize the role of these social determinants of health. This is something that we also saw in other datasets, particularly in the Brazilian dataset, that was a national data set and we found that socio-demographic factors were impacting hospitalization decisions.

What we also found was that the model - the naive model that one is able to produce - is actually rather biased. You could see here how the model performs out of samples for Black individuals and white individuals. The false positive rate of the model for Black individuals was twice as much as the model for - the false positive rate for white individuals - even though we were controlling for race, we were controlling for socio-demographic factors, we were controlling for social determinants of health. Despite that, the model was much more eager to make the prediction that the Black individual was going to be hospitalized compared to a white individual. Correspondingly, it was most likely to make a false negative prediction for a white individual compared to a Black individual, which suggests that they are apparently hidden features in the data that are not visible to us, perhaps reflecting structural bias and other factors that make the model make that biased prediction. There are ways and we have addressed them in a paper we published on how one can correct for these factors and produce models that do not have this sort of bias.

Slide 6

These are some results on ICU prediction - predicting an ICU admission -roughly the median gap between an admission to the hospital and admission to the ICU at least in our data set was about four hours. We will use different cutoffs. If you use the latest information, then you get quite accurate models with AUCs on the order of 93-95 percent. If you start cutting off the information that you were going to use so 12 hours in advance the performance of the model drops to about 86 percent. 24 hours in advance, the performance of the model drops to about roughly 80 percent. What I would like to emphasize is that we compared these models that we developed to some standard models that predict ICU admissions, there are some well-known sepsis models, NEWS, qSOFA, they're called, and these are fairly inaccurate in this case, indicating that standard models for ICU prediction, at least in the COVID cases, fail to predict an ICU admission. This indicates a rather unique signature of the disease. Here, you find some of the variables that were again highly correlated with the outcome with an ICU admission. What was interesting to us was that this vital score that we produced that captured the dynamic evolution of the vitals pretty much tells the entire story. There are some other variables or some lab variables (LDH, CRP) that have been identified by other studies that also are contributing, but if one just takes the dynamic evolution of the vitals, that pretty much tells the ICU admission story.

Slide 7

Finally we produced a number of calculators that we made available on on the web. This, I understand, were used by our colleagues and collaborators at Mass General Hospital in the early stages of the epidemic. It was very easy to input some of the key variables and get a prediction, for instance, for an ICU admission, or for a mechanical intubation need of a patient, and we found that, you know, we were challenged to find cases - what is the model telling me that the very experienced clinician cannot potentially predict by just looking at the patient. Here are some cases, and there are many others, where patients were admitted, they were stable for a couple of days, nothing in their clinical outlook suggested that the condition of this patient was going to deteriorate, but then the model was able - upon admission - based on some special laboratory results, predict that the patient was going to need the ICU care. These are two cases with the details of those cases.

Slide 8

So that brings me to a conclusion. Of course, I'm just presenting. There have been many people that have contributed to this work and I would like to thank them, including the students in my group, but also our collaborators at Mass General Brigham Boston Medical Center and some of the other areas where we've been able to get data from. Thank you so much for your attention and looking forward to questions at the end of the session.