Centre of Innovation for Data tech and Artificial Intelligence

Report Proof of Concept
CIDAI-POC-2021-01

Predicting the evolution of COVID-19 mortality risk:
A recurrent neural network approach

Report developed by:
Marta Villegas
Aitor Gonzalez-Aguirre
Asier Gutiérrez-Fandiño
Casimiro Pio Carrino
Jordi Armengol-Estapé
Alfonso Valencia
Marina Valls Soler

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

CIDAI-POC-2021-01

Some rights reserved. This work is available under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0).

Under the terms of this licence, you may copy, redistribute and adapt the work for non-commercial purposes, provided the work is appropriately cited, as indicated below.
In any use of this work, there should be no suggestion that CIDAI endorses any specific organization, products or services. The use of the CIDAI logo is not permitted.
If you adapt the work, then you must license your work under the same or equivalent Creative Commons licence.

If you create a translation of this work, you should add the following disclaimer along with the suggested citation: “This translation was not created by the Centre of Innovation for Data tech and Artificial Intelligence (CIDAI). CIDAI is not responsible for the content or accuracy of this translation. The original English edition shall be the binding and authentic edition”.
Any mediation relating to disputes arising under the licence shall be conducted in accordance with the mediation rules of the World Intellectual Property Organization.

Suggested citation. CIDAI-POC-2021-01 // Predicting the evolution of COVID-19 mortality risk: A recurrent neural network approach by Marta Villegas, Aitor Gonzalez-Aguirre, Asier Gutiérrez-Fandiño, Casimiro Pio Carrino, Jordi Armengol-Estapé, Alfonso Valencia, Marina Valls Soler, CIDAI, 2021. Licence: CC BY-NC-SA 4.0

Index

1. Abstract

1.1. Objectives

The propagation of COVID-19 in Spain prompted the declaration of the state of alarm on March 14, 2020. On 2 December 2020, the infection had been confirmed in 1,665,775 patients and caused 45,784 deaths. This unprecedented health crisis challenged the ingenuity of all professionals involved. Decision support systems in clinical care and health services management were identified as crucial in the fight against the pandemic.

Our main goal is to develop a model based on recurrent neural networks (RNNs) that responds to the reality of hospitals during the pandemic. The goal for the model to generate daily predictions in a hospital setting that can be used as an early warning system and allow for proper management of existing resources.

1.2. Summary of the problem

Decision support systems for clinical care and health service management are crucial for responding effectively to the challenge of the COVID-19 pandemic and optimizing the hospitalization process for this emerging disease. Artificial intelligence and Deep Learning tools offer a range of possibilities for obtaining models that, trained with historical data, can anticipate future scenarios.

1.3. Solution

The model developed predicts the risk of mortality daily for each patient admitted taking into account all available data. The model has obtained encouraging sensitivity scores, highly relevant metrics in the clinical domain. This model can support decision making in healthcare systems.

2. Description of the problem

The problem that the Recurrent Neural Network model developed in the BSC-CNS seeks to solve is to support decision-making in healthcare systems with limited resources and under additional pressure by the magnitude of the spread of COVID-19.

The consensus is generally to focus the allocation of life support measures to those most likely to be in a critical situation and lose their life project. These decisions are usually made taking into account different parameters that include various dimensions: clinical situation, cognitive situation, presence of irreversible functional impairment, quality of life, probability of significant response to treatments, quality of life expectancy, future care scenarios, etc. These criteria, however, are merely indicative as action must be taken on a case-by-case basis, also taking into account the patient’s values and preferences.

Medical care, therefore, require reflective analysis that must be done efficiently, particularly in an epidemic scenario. The Recurrent Neural Network is presented as an additional resource that can be added to the criteria for prioritizing, allocating, and withdrawing existing care resources that allow professionals to make these challenging decisions.

3. Implementation of the solution

3.1. Architecture, technology and data used.

3.1.1 Architecture.

As depicted in Figure 1, we designed an Artificial Neural Network with four modules, namely, the embedding module (grey boxes), the recurrent module (dark green boxes), the classifier module (light green boxes) and, optionally, the attention modules (blue boxes).

The embedding module only acts on the static vector through a fully-connected layer that encodes the high-dimensional sparse feature vector into a lower dimensional dense vector.

The recurrent module accounts for both the encoding and the memorization of the temporal information provided by the dynamic vectors (red boxes). It consists of an unidirectional RNN, built with either an LSTM or a Gated Recurrent Unit (GRU). Each day, the RNN cells process the relevant information of the patient dynamic data and produce an output vector that stores the relevant clinical information until that day to perform the predictions.

The attention module of dynamic field finds the correlations of all previous RNN's outputs and merges all the global relevant information of the sequence until a given day.

Finally, the classifier module consists of two fully-connected layers followed by a sigmoid activation function to produce the binary mortality predictions. We placed the classifier module on top of the RNN cells and the static embeddings to predict the mortality risk each day.

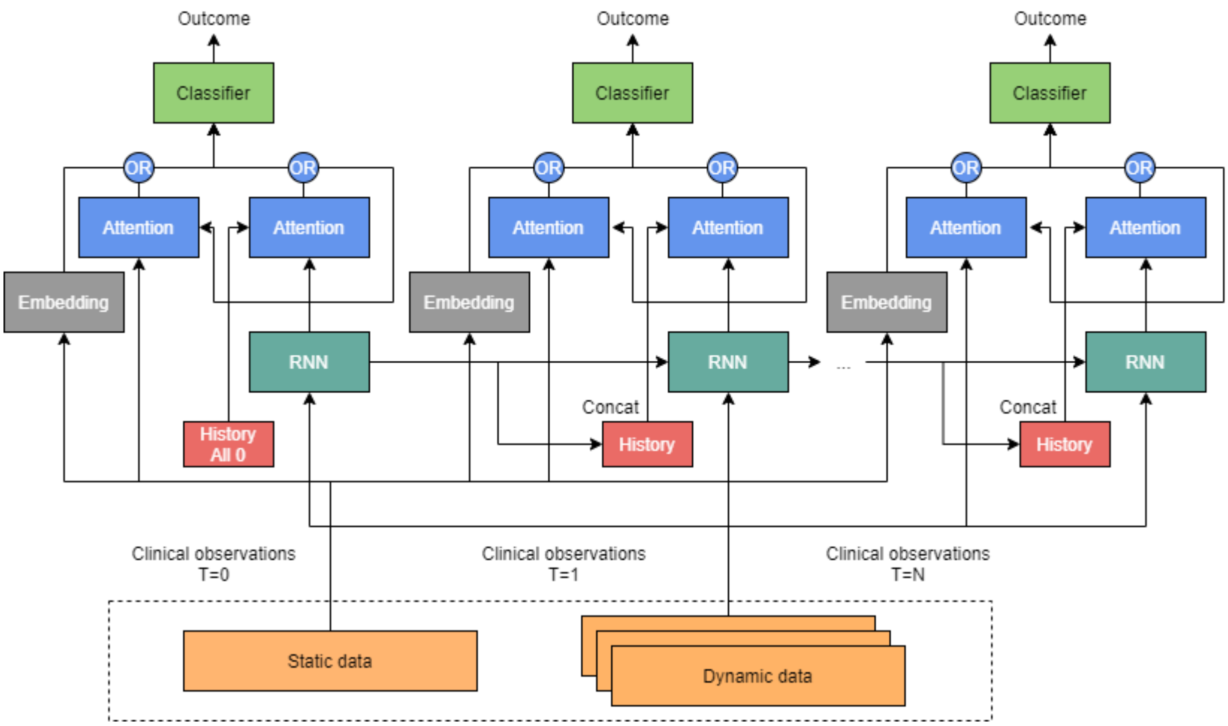

Figura 1 – RNN architecture as described in section 1.1

3.1.2 Data used.

In this work, we apply Deep Learning techniques for predicting the clinical outcome of patients with COVID-19. The result is a model that predicts mortality risk to support decision making. This study leverages two datasets with clinical information on 2,307 and 3,870 patients from HM Hospitales (HM) and the Hospital Universitario 12 de Octubre (H12O) in Spain.

The two datasets used to form and validate the predictive models include all patients admitted with COVID-19 in HM and H12O. In the HM dataset, the cohort includes all patients with COVID from the beginning of the pandemic in Spain until April 20, 2020. In the H12O dataset, the cohort extends until October 14, 2020. In both cases, the data contain two different types of variables: static variables that are constant throughout admission, such as sex or age, and dynamic variables that are measured at different times during the hospital stay, such as medications, lab tests and vital signs.

Having two different datasets allowed us to validate the results and check whether the architecture and design developed are equally suitable for both cases. Although both datasets contain information from patients with COVID, they have different characteristics. This corroborates the robustness and applicability of the proposal.

3.1.3 Implementation.

We designed a model to monitor the mortality risk of the patient by producing a daily prediction during the patient's hospital stay. Therefore, unlike standard training schemes for recurrent models, we designed an ad-hoc training scheme to produce daily predictions. Specifically, we do not feed the entire sequence of dynamic data into the RNN to output a single prediction. But, we feed each daily records of temporal and static data day-by-day and output a prediction for each day.

In addition, we experimented with other model's hyperparameter configurations. All the layers are trained end-to-end with stochastic gradient descent.

At inference, the model is fed with the static vector and then, at each time step, we input the corresponding dynamic vector and the model outputs the probability of patient mortality. To obtain the mortality predictions, the probability is discretized into a binary label using a threshold, as in logistic regression. We set the threshold to the default value of 0.5. Notice that, after training, other values of the model's threshold can be used, depending on how much we want to be confident in the prediction.

In addition, we designed specific evaluation metrics shown in Figure 2 to better test the model and reflect our intended use case providing daily predictions. In this case, we get the prediction vectors already returned by the system, and use them to calculate the daily performances as follows:

- **Daily performances from the admission:** computes the performance every day from the admission day. Thus, for instance, the accuracy at day 3 is the average accuracy of all 3rd day predictions, and the accuracy of day 4 from admission is the average accuracy of all 4th day predictions, and so to the end.
- **Daily performances from the outcome:** computes the performance every day from the outcome day

Day from admission															
Patient	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 10	Day 11	Day 12	Day 13	Day 14	...
1	0	0	0	0	0	0	0	0	1	1	1	1	1	1	
2	0	0	0	1	0	0	0	0	0	0					
3	0	0	1	1	0	1	1	1							
4	0	0	0	1	1	1	1								
5	0	0	0	1	1	1									
...	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score

Days before outcome															
Patient	...	Day 14	Day 13	Day 12	Day 11	Day 10	Day 9	Day 8	Day 7	Day 6	Day 5	Day 4	Day 3	Day 2	Day 1
1		0	0	0	0	0	0	0	0	1	1	1	1	1	1
2						0	0	0	1	0	0	0	0	0	0
3								0	0	1	1	0	1	1	1
4									0	0	0	1	1	1	1
5										0	0	0	1	1	1
...	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score	Score

Figure 2 – Daily performances comparison.

3.2. Challenges solved and results obtained.

Regarding the global results, for HM, the best ensemble RNN model is significantly better than the other models when considering sensitivity, reaching an important result of 0.84 in sensitivity with no penalty in F1 and a small impact in accuracy. Similarly, for H12O, the RNN ensemble model outperforms all models in sensitivity. In particular, the sensitivity oriented ensemble model outperforms RF in 18 percentage points and SVC in 16 percentage points reaching an 0.80 in sensitivity.

When evaluating day-by-day performance from the outcome, the HM RNN model shows a higher and flatter curve in sensitivity compared to the baselines, starting at 0.97, ending at 0.55 and maintaining a performance of 0.79 or better until day 14. This shows that the system is capable of making early predictions. Similarly, the H12O RNN model shows a better performance in sensitivity, maintaining a score of 0.64 or better until 7 days before the outcome.

Even in our data constrained scenario, the models outperform the strong baselines of tuned RF and SVC in terms of sensitivity. We have demonstrated the viability of our approach to predict the clinical outcome of patients infected with SARS-CoV-2. The result is a time series model that can support decision making in healthcare systems and aims at interpretability. Despite the low resource scenario, the results are promising and suggest that with more data the performance of the model will increase.

3.3. Current limitations of the Technology.

For implementing the system into a new facility or database, it is necessary to have a sufficient amount of patient data to retrain the model.

4. Potential impact of the solution.

This system can be implemented in every hospital that stores sufficient amounts of previous patient data with COVID-19.

5. Cited works

-Villegas, Marta & Gonzalez-Agirre, Aitor & Gutiérrez-Fandiño, Asier & Armengol-Estapé, Jordi & Carrino, Casimiro & Fernández, David & Soares, Felipe & Serrano-Balazote, Pablo & Pedrera Jiménez, Miguel & Garcia Barrio, Noelia & Valencia, Alfonso. (2020). Predicting the Evolution of COVID-19 Mortality Risk: a Recurrent Neural Network Approach. 10.1101/2020.12.22.20244061.

6. Bibliography

-Villegas, Marta & Gonzalez-Agirre, Aitor & Gutiérrez-Fandiño, Asier & Armengol-Estapé, Jordi & Carrino, Casimiro & Fernández, David & Soares, Felipe & Serrano-Balazote, Pablo & Pedrera Jiménez, Miguel & Garcia Barrio, Noelia & Valencia, Alfonso. (2020). Predicting the Evolution of COVID-19 Mortality Risk: a Recurrent Neural Network Approach. 10.1101/2020.12.22.20244061.

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

www.cidai.eu