Modelling patient trajectories using multimodal information

João F Silva[1], Sérgio Matos[1]

Electronic Health Records (EHRs) offer a multi-dimensional longitudinal view of the patient medical history and thus contain a trajectory that is representative of the evolution of the patient health status through time. This context can be used by physicians to monitor patient health and make more accurate prognoses or diagnoses, but the volume, variety and variability of this information hinders its effective use. Patient trajectory modelling can assist by exploring existing information in a scalable manner and can contribute to health care quality by guiding management decisions and fostering preventive medicine practices. Such models can be applied in various problems, such as predicting disease progression, recommending therapies, or predicting when a patient will be readmitted to the health care system.

This work describes a flexible neural network architecture for patient trajectory modelling that combines different types of information and considers the temporal aspect of clinical data. The model encodes each patient admission as a continuous vector representation derived from different data modalities. These patient admission representations are then combined through a recurrent network to model the patient trajectory. In this work we used data from the public MIMIC-III clinical database and considered the prediction of two different clinical outcomes: patient readmission within a period of 30 days after the patient is discharged, and disease progression, by predicting the expected diagnoses on a future admission. We encoded clinical text using ClinicalBERT contextual embeddings and explored the inclusion of timing and coding information. The proposed model can be easily expanded to support more data modalities as well as the prediction of other clinical outcomes and is publicly available at https://github.com/bioinformatics-ua/PatientTM.

FIGURE 1
The trajectory modelling architecture generates vector representations for each patient admission at the Feature Concatenation layer. The model considers sliding windows with k admissions for each patient, scanning from the oldest to the most recent admission.
The resulting vector representations are forwarded through a recurrent model and a fully connected layer. Finally, N sigmoid activations are used to calculate class probabilities, where N corresponds to the number of classes to predict. In diagnose prediction this corresponds to the diagnoses considered, while in readmission prediction the output layer is simplified to consider a single output.