Early Recognition of Sepsis with Gaussian Process Temporal Convolutional Networks and Dynamic Time Warping

Michael Moor[1,2], Max Horn[1,2], Bastian Rieck[1,2], Damian Roqueiro[1,2], Karsten Borgwardt[1,2]
[1]Department of Biosystems Science and Engineering, Basel, ETH Zurich; [2]SIB Swiss Institute of Bioinformatics

1 Introduction

Sepsis is a life-threatening host response to infection associated with high mortality, morbidity, and health costs [1]. Its management is highly time-sensitive since each hour of delayed treatment increases mortality due to irreversible organ damage.
Meanwhile, despite decades of clinical research, robust biomarkers for sepsis are missing. Detecting sepsis *early* by utilizing the affluence of high-resolution intensive care records has become a challenging machine learning problem.

2 Dataset & Filtering

Dataset: MIMIC-III critical care database [2]

Exclusion criteria: Age < 15 years old, relevant data missing (e.g. admission time), sepsis onset earlier than 7 hours after ICU admission. Not logged via Metavision.

Filtering: We apply case-control matching and extract input data of 44 vital and laboratory variables up to 48 hours preceding their (matched) sepsis onset and after ICU admission.

Variable	Sepsis Cases	Controls
n	570 (9.2%)	5,618 (90.8%)
Female	236 (41.4%)	2,548 (45.4%)
Male	334 (58.6%)	3,070 (54.6%)
Mean time to sepsis onset in ICU (median)	16.7 h (11.8 h)	—
Age ($\mu \pm \sigma$)	67.2 ± 15.3	64.2 ± 17.3
Admission type		
Emergency	504 (88.4%)	4,689 (83.5%)
Elective	60 (10.5%)	872 (15.5%)
Urgent	6 (1.1%)	57 (1.0%)

3 Sepsis Definition

To define sepsis, we make use of the most recent definition, Sepsis-3 [1], which requires a co-occurrence of suspected infection (SI) and organ dysfunction (as measured by the SOFA score). We determine Sepsis-3 on a hourly basis to approximate a temporally resolved label.

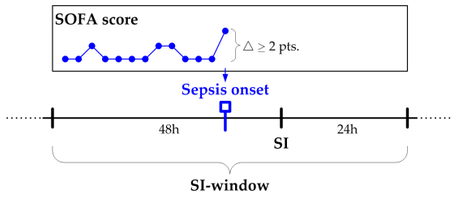

4 Methods

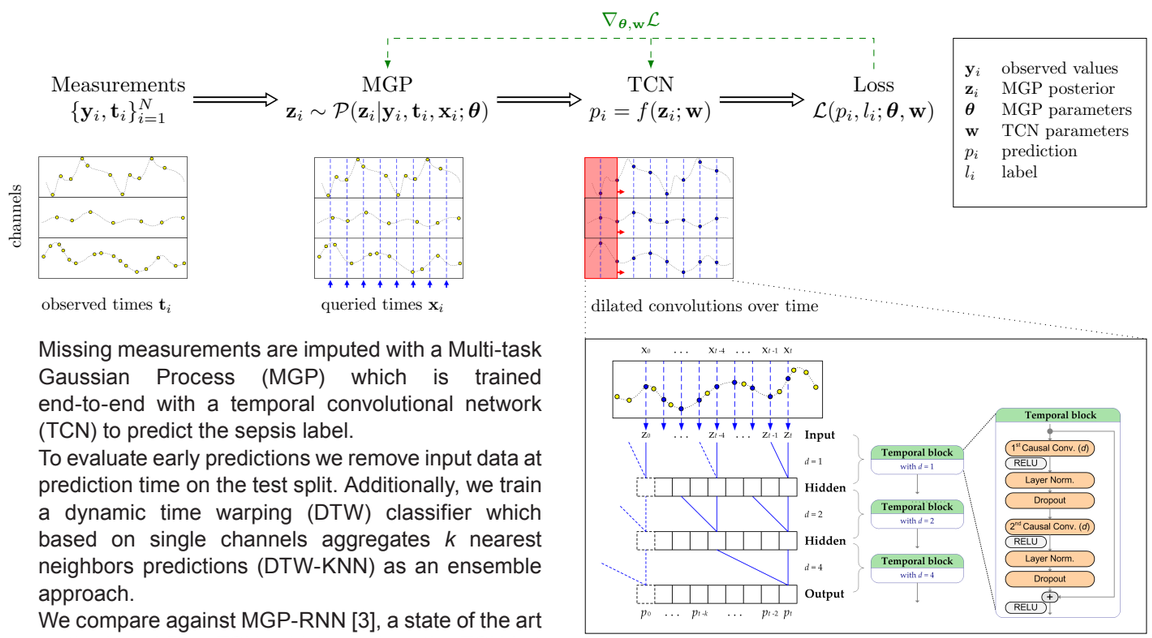

Missing measurements are imputed with a Multi-task Gaussian Process (MGP) which is trained end-to-end with a temporal convolutional network (TCN) to predict the sepsis label.
To evaluate early predictions we remove input data at prediction time on the test split. Additionally, we train a dynamic time warping (DTW) classifier which based on single channels aggregates *k* nearest neighbors predictions (DTW-KNN) as an ensemble approach.
We compare against MGP-RNN [3], a state of the art method, and Raw-TCN, which employs the TCN but with simple carry-forward imputation.

5 Results

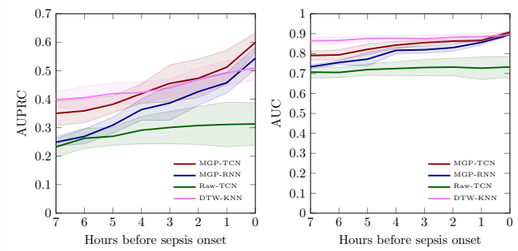

Our proposed sepsis prediction methods MGP-TCN and DTW-KNN exhibit favorable performances over all prediction horizons and consistently outperform the baselines [3].

To our knowledge, MGP-TCN is the first method to leverage temporal causal dilated convolutions on irregularly sampled time series. DTW-KNN ensemble performs surprisingly well in earlier horizons while MGP-TCN dominates in the 4 hours before sepsis onset.

6 References

[1] Singer, Mervyn, et al. "The third international consensus definitions for sepsis and septic shock (Sepsis-3)." Jama 315.8 (2016): 801-810.
[2] Johnson, Alistair EW, et al. "MIMIC-III, a freely accessible critical care database." Scientific data 3 (2016): 160035.
[3] Futoma, Joseph, Sanjay Hariharan, and Katherine Heller. "Learning to detect sepsis with a multitask Gaussian process RNN classifier." Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org, 2017.

Code & Data: