

Sihan ZHU, Jian PU, 2021. A self-supervised method for treatment recommendation in sepsis. *Frontiers of Information Technology & Electronic Engineering*, 22(7):926-939. <https://doi.org/10.1631/FITEE.2000127>

A self-supervised method for treatment recommendation in sepsis

Key words: Treatment recommendation; Sepsis; Self-supervised learning; Reinforcement learning; Electronic health records

Corresponding author: Jian PU

E-mail: jianpu@fudan.edu.cn

 ORCID: <https://orcid.org/0000-0002-0892-1213>

Motivation

1. Sepsis treatment is a highly challenging effort to reduce mortality in hospital intensive care units since the treatment response may vary for each patient.
2. Tailored treatment recommendations are desired to assist doctors in making decisions efficiently and accurately.
3. Reinforcement learning methods can learn from patients' responses and give personalized recommendations.

Main idea

1. An uncertainty evaluation method is proposed to separate patient samples into two domains according to their responses to treatments and the state value of the chosen policy.
2. Examples of two domains are then reconstructed with an auxiliary transfer learning task.
3. A distillation method of privilege learning is tied to a variational auto-encoder framework for the transfer learning task between the low- and high-quality domains.

Method

1. A reinforcement learning (RL) based uncertainty evaluation method is proposed by comparing states under different policies.
2. An auxiliary transfer learning task is designed using variational auto-encoder (VAE) to reconstruct samples from different domains.
3. Samples from the high-quality domain are seen as “teachers” to samples from the low-quality domain, which results in a privilege learning task to help train the model.

Major results

- Main framework




Fig. 1 General architecture of our work

Major results

- Performance comparison of our model and related methods

Table 3 Performance comparison of different methods

Method	Expected return	Estimated mortality	Major mortality	Jaccard coefficient
BLSTM	–	22.1%	–	0.376
RLSTM	–	21.3%	–	0.378
DDQN	14.3	14.6%	12.4%	0.289
Our method	15.1	12.3%	9.5%	0.357

Major results

- Performance under different settings




Fig. 5 Expected Q-value distribution under different reconstruction parameter l settings (a) and estimated mortality in different Q-value intervals under the setting of reconstruction for the transfer module (b)

Table 4 Performance with different reconstruction parameter l settings used in the transfer module $M(\cdot)$

l	Estimated mortality (%)	Expected return
0.1	12.5	15.00
0.2	12.6	15.01
0.5	12.3	14.98
0.8	14.3	15.28
0.9	14.7	15.41

Best results are in bold




Fig. 6 Estimated mortality distribution in DDQN prediction and our model with a conservative setting $l = 0.3$ and an efficient setting $l = 0.7$

Conclusions

1. We have proposed a self-supervised RL method with privileged learning to tackle the treatment recommendation problem.
2. The proposed method provides flexibility on the trade-off between efficiency and accuracy.
3. The conservative policy provides predictions which are more likely to be regarded as guiding suggestions, and the efficient policy provides exact treatment prescriptions on easy-handling cases while sacrificing accuracy on ambiguous ones, which then can be left to doctors.