Training and validating a treatment recommender with partial verification evidence.

Authors: Vishnu Unnikrishnan, Clara Puga, Miro Schleicher, Uli Niemann, Berthold Langguth, Stefan Schoisswohl, Birgit Mazurek, Rilana Cima, Jose Antonio Lopez-Escamez, Dimitris Kikidis, Eleftheria Vellidou, Rüdiger Pryss, Winfried Schlee, Myra Spiliopoulou

Publisher: Elsevier B.V.

In: Artificial Intelligence in Medicine, 2025, volume 160, article 103062, published online: February 2025

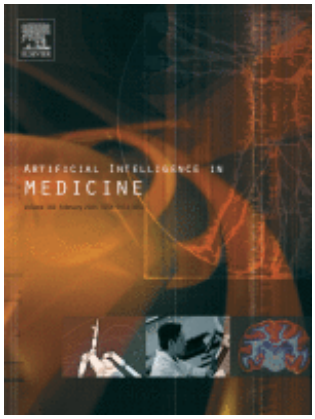

Abstract:

Background:
Current clinical decision support systems (DSS) are trained and validated on observational data from the clinic in which the DSS is going to be applied. This is problematic for treatments that have already been validated in a randomized clinical trial (RCT), but have not yet been introduced in any clinic. In this work, we report on a method for training and validating the DSS core before introduction to a clinic, using the RCT data themselves. The key challenges we address are of missingness, foremost: missing rationale when assigning a treatment to a patient (the assignment is at random), and missing verification evidence, since the effectiveness of a treatment for a patient can only be verified (ground truth) if the treatment was indeed assigned to the patient — but then the assignment was at random.

Materials:
We use the data of a multi-armed clinical trial that investigated the effectiveness of single treatments and combination treatments for 240+ tinnitus patients recruited and treated in 5 clinical centres.

Methods:
To deal with the 'missing rationale for treatment assignment' challenge, we re-model the target variable that measures the outcome of interest, in order to suppress the effect of the individual treatment, which was at random, and control on the effect of treatment in general. To deal with missing features for many patients, we use a learning core that is robust to missing features. Further, we build ensembles that parsimoniously

exploit the small patient numbers we have for learning. To deal with the 'missing verification evidence' challenge, we introduce counterfactual treatment verification, a verification scheme that juxtaposes the effectiveness of the recommendations of our approach to the effectiveness of the RCT assignments in the cases of agreement/disagreement between the two.

Results and limitations:
We demonstrate that our approach leverages the RCT data for learning and verification, by showing that the DSS suggests treatments that improve the outcome. The results are limited through the small number of patients per treatment; while our ensemble is designed to mitigate this effect, the predictive performance of the methods is affected by the smallness of the data.

Outlook:
We provide a basis for the establishment of decision supporting routines on treatments that have been tested in RCTs but have not yet been deployed clinically. Practitioners can use our approach to train and validate a DSS on new treatments by simply using the RCT data available to them. More work is needed to strengthen the robustness of the predictors. Since there are no further data available to this purpose, but those already used, the potential of synthetic data generation seems an appropriate alternative.

Related links:

- Online publication: https://doi.org/10.1016/j.artmed.2024.103062
- Download full article: https://www.sciencedirect.com/science/article/pii/S093336572400304X/pdfft?md5=6d73a6e6b96a827446865cc66443afcf&pid=1-s2.0-S093336572400304X-main.pdf