Modeling personalised treatments for intensive care patients using Dynamic Bayesian Networks

Federico Pirola[1,*]

[1]Department of Informatics, Systems and Communication, University of Milano-Bicocca

Abstract

I plan to extend dynamic Bayesian networks to assist physicians to make personalized and non myopic treatment decisions. The research activity will be oriented towards introducing the concept of personalized dynamic Bayesian network and to exploit counterfactuals for optimal sequential decision making. The main advantages and limitations of what proposed will be investigated and discussed. As for the domain of application I will concentrate on the healthcare sector with specific reference to two complex problems, i.e„ early detection of broncopulmonary dysplasia in premature babies and prognostic evaluation of functional outcomes of patients who suffered traumatic brain injuries.

Keywords

Dynamic Bayesian Networks, Causal Networks, Sequential Treatments, Personalized Treatments, Intensive Care

1. Introduction

In the last decades, significant advances have been made in the clinical practice, ranging from laboratory automation to real-time patients monitoring [1]. This progress enabled clinical researchers to discover key biomarkers, such as genes or proteins, involved in disease development. Still, predicting disease development over time and choosing the optimal treatment are complex tasks, making patient prognosis a challenging yet urgent problem to be tackled. When it comes to intensive care patients, physicians often provide treatments with immediate effects to stabilise the condition of critical patients: their goal is to avoid, or at least to reduce, the risk of severe outcomes such as coma or death. However, physicians lack precise and comprehensive data to prevent late-time undesired effects of the administered treatment [2]. Indeed, each patient has their own medical history, such as lifestyle, genetic profile, previous and ongoing therapies, making it a unique case on its own. This means that the same disease with different symptoms often evolves at different rates in different patients, which in turn may result into divergent outcomes depending on the specific patient. Nevertheless, physicians struggle to leverage the prior knowledge concerning patients' health, making personalised treatments difficult to provide timely [3]. Furthermore, these limitations make it hard to investigate the onset of side effects and comorbidities induced by administered treatments. Due to the recent achievements of Artificial Intelligence (AI) in many domains [4], the following question has emerged as an extremely relevant one: *Can AI assist clinicians in decision making?*

In this context, my research activity aims to investigate whether a specific instance of the class of probabilistic graphical models (PGMs), i.e., dynamic Bayesian networks (DBNs), can help physicians to make personalized and non-myopic treatment decisions in the case of intensive care patients. To tackle such a challenging problem I'm collaborating with physicians to collect and elicit domain knowledge according to the semantic of DBNs.

Doctoral Consortium at the 23rd International Conference of the Italian Association for Artificial Intelligence Bolzano, Italy, November 25-28, 2024.

*Corresponding author.

f.pirola17@campus.unimib.it (F. Pirola)

0009-0000-7090-4716 (F. Pirola)

2. Research Question

Randomized Controlled Trials (RCTs) represent the gold standard to study the effectiveness of novel treatments and provide personalised assistance to patients [5]. However, when it comes to model the temporal evolution of clinical conditions, RCTs are not viable. In fact, they are often unfeasible due to *i*) insufficient number of patients, *ii*) high costs and time constraints, *iii*) loss of patients during the follow-up. As a result, the only remaining alternative is offered by Observational Studies (OS). The key difference is that OS are known to be biased [6]. Compared to RCTs, OS are characterised by two major problems: *i*) patients are not randomly assigned to a certain treatment, thus they do not ensure similar characteristics between treated and untreated patients, *ii*) time-dependent data are not collected at fixed points in time. As a consequence, AI models deployed in the healthcare domain must be designed to deal with incomplete data which underlying mechanism is unknown. In practice they should represent a support to physicians in understanding why a certain value is missing and which are the factors involved.

For these reasons, in my research project I plan to develop explainable temporal models to understand the mechanisms underlying complex biological processes in the OS scenario. I also plan to exploit the models to make the treatments personalized based on the characteristics of severe patients. In greater detail, my plan is articulated accordingly to the three rungs of the ladder of causation [7]:

1. *Association Rung* - Development of a temporal model dealing with incomplete data to:
    - Predict the evolution of complex diseases or clinical conditions
    - Understand which patients are more likely to be in severe conditions
2. *Intervention Rung* - Integrate causal reasoning to estimate the effects of interventions to:
    - Provide personalised assistance based on patients' profiles
    - Understand which is the most appropriate time to treat
3. *Conterfactual Rung* - Estimate potential outcomes in case of alternative treatments to:
    - Enrich the knowledge with respect to intensive care patients

3. Related Works

The popularity of AI models in the healthcare domain has risen due to the excellent level of accuracy in predicting diseases [8]. Potentially, they offer a valid support to tackle complex decision making problems. Nevertheless, diseases are characterized by complex biological processes which evolve over time, but they are often modelled in a static fashion. Deploying atemporal models precludes the possibility to identify long-term effects of treatments. Additionally atemporal models make it difficult to provide insights about the best timing to treat a given patient. Moreover, when it comes to rare or severe clinical conditions, international healthcare committees are unwilling to allow the adoption of AI methods in the clinical practice. This reluctance is owed by the limitations of State-of-the-art (SoA) models, which are not transparent and miss interpretability [9]. In practice, computer scientists are unable to explain why a certain treatment is suggested by the AI model for a given patient. SoA techniques, such as Deep Learning algorithms, are also characterized by predetermined inputs and outputs. It means that the outcome prediction is precluded when the variables on which the models have been trained are not all observed. Making diagnosis and prognosis with these methods is therefore possible only in case of full observability of the reference set of covariates and when a considerable amount of data is available which is typically not the case in healthcare. Furthermore, these methods have no standard way of integrating prior clinical knowledge, thus they leave the physicians apart in the model design phase [8]. Building interpretable models, as PGMs, favours the interaction between different domain expertise, beyond making their underlying decision-making process transparent. These methods ensure to switch from a *competitive framework*, for which the models are deployed to outperform physicians' predictive accuracy, to a *collaborative framework*, incorporating prior clinical

knowledge while learning the structure and the parameters of the models. Allowing the physicians to participate in the definition of the models is likely to lower their reluctance to introduce AI to support decision-making in delicate situations. Thus, I think that employing DBNs, as discrete-time PGMs, is the optimal solution to study the temporal evolution of severe clinical conditions. Thanks to their superior explainability, they facilitate the communication between different domain experts and support physicians in providing personalised care to diverse groups of patients.

4. Methods

DBNs [10] are temporal PGMs expressed as a set of a prior network ($\mathcal{B}_0$), specifying distributions over the initial states and a transition network ($\mathcal{B}_{\mathcal{TR}}$), specifying the probabilities of transitioning from a previous to a current state. Additionally, DBNs are made of a graph $\mathcal{G}$ encoding the conditional independencies between variables and a set of conditional probability distributions with parameters Θ, which quantify the network [10]. Unlike other temporal models, they consider time in a discrete manner. Depending on the available data and the prior knowledge provided, a constant time window is defined. As a consequence, dependencies between nodes are set for each resulting time slice. By applying an unrolling operation, a DBN could be also viewed as a classic BN and all the standard methods for the static framework could be extended to temporal data too. In Figure 1 a representation of a sample DBN structure is reported both as the pair of networks $\mathcal{B}_0$ and $\mathcal{B}_{\mathcal{TR}}$ and in his unrolled form ($\mathcal{B}_{\mathcal{UN}}$).

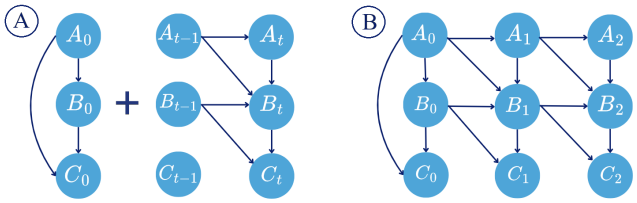

Figure 1: DBN Structure Representation. A): $\mathcal{B}_0$ + $\mathcal{B}_{\mathcal{TR}}$, B): $\mathcal{B}_{\mathcal{UN}}$

DBNs are generally employed to predict time-dependent outcomes by modeling multivariate time series data. Sometimes mixed cases considering both time-dependent (dynamic) and atemporal (static) variables could lead, though, to more accurate results. In [11] two different frameworks to encode dependencies between static and dynamic variables have been proposed. In the first case scenario the two categories of variables could be in a parent-children relationship only in $\mathcal{B}_0$. Instead, in the alternative setting arcs going from static to dynamic variables are allowed in $\mathcal{B}_{\mathcal{TR}}$ too. In Figure 2 two sample DBNs' structures are represented showing the differences between the two frameworks.

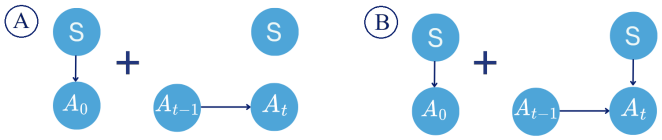

Figure 2: DBN's Static nodes representations. S: static node, A_j: dynamic node at time j

Missing data could also be managed during structure learning phase via Structural Expectation Maximization (SEM) method [12]. SEM has been demonstrated to outperform classical MICE imputation in structure recovery independently from the missing values rate, the missing mechanism underlying data and the data dimensionality. Indeed, as PGMs, methods of causal inference and causal discovery could be implemented too. It means that DBNs are learned coherently with potential confounders, namely set of factors which cause spurious associations between pairs of variables, to understand which treatments are truly beneficial to each subject [7]. Additionally, individual causal effects could be estimated through conterfactual queries to dispense personalised treatments according to patients' profiles.

5. Case Studies and Work Plan

5.1. Bronchopulmonary Dysplasia (BPD)

Bronchopulmonary dysplasia (BPD) is a chronic form of lung disease which affects preterm babies, resulting in both short- and long-term pulmonary defections and breathing difficulties. Novel treatments, such as surfactant cures or non-invasive ventilation techniques, have been developed in the last few decades [13]. However, effectively preventing BPD in extremely preterm newborns is still a difficult task. Previous works have already shown the higher propensity to BPD of extremely underweight infants and kids born at lower gestational age [13], but finding the best configuration of treatments over time is still a challenge due to the contraindications of many medications and therapies. In this study a cohort of approximately 450 preterm newborns from Mangiagalli clinic in Milan (Italy) is considered. The study details are reported in Figure 3.

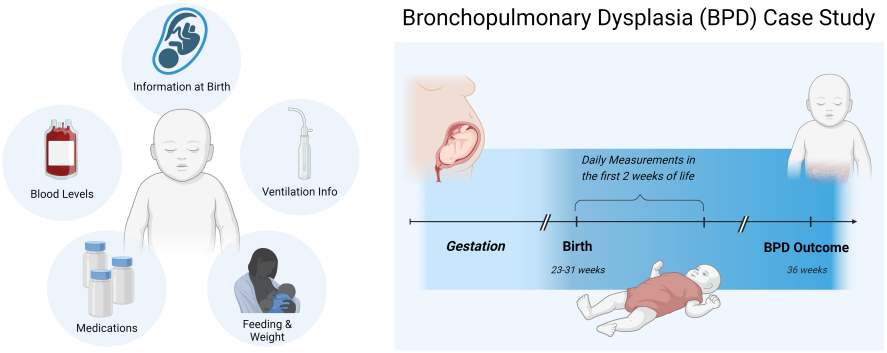

Figure 3: All extremely preterm (23-27 weeks) and very preterm infants (28-31 weeks) born between 2017 and 2023 at Mangiagalli Clinic are included in the study. For each patient both information at birth (static variables) and longitudinal data collected during the first 14 days of life (dynamic variables) are available. BPD ordinal outcome at 36 weeks at gestation is computed using Jensen scale [14], grouping newborns in four classes.

The primary objective of this study is to early diagnose BPD using DBNs, thus predicting BPD outcome in the first days of life. DBN's structure and parameters will be learnt functionally to the prediction of BPD ordinal outcome in preterm newborns. One of the main challenges of this study is to discretize both the time domain and the variable domain. Both the different granularities of each group of variables and between-subjects variability due to the observational nature of the study must be taken into account. Neonatologists from Mangiagalli clinic have provided standard intervals for blood levels, respiratory parameters and other clinical features. Discretization is fundamental in this setting because normal ranges for blood and respiratory parameters could vary in the first two weeks of life (i.e. PCO_2 in the first three days of life varies between 40 and 50 mmHg, whereas normal values during the second week ranges from 50 to 60 mmHg). Instead, statistics on medications' daily assumptions and respiratory supports are adopted for treatment variables discretization. Another trivial point of this case study is the presence of patient's incomplete daily profiles. In this setting, methods of structure learning and inference must be applied in absence of full observability. Missing data should be handled and, with some measurement uncertainty, explained functionally to observed or unobserved confounders. Supported by medical expertise from Mangiagalli clinic, those potential factors involved in the explanation of the missingness will be identified. Hence, missing data could be imputed coherently with the mechanism underlying the clinical data.

5.2. Traumatic Brain Injuries (TBI)

Traumatic Brain Injury (TBI) is defined as an alteration in brain function caused by an external force. It varies in severity from mild TBI to moderate and severe TBI. Despite the existence of many works investigating TBI impact on patients mental, physical and cognitive status, SoA models tend to model

TBI as a static event [15]. Moreover, if AI models are the gold standard in TBI diagnosis, they still struggle to predict long-term impact of the trauma or treatment effects. In Bhattacharyay et al [16] they proposed a novel approach which models dynamic information of TBI patients collected during Intensive Care Unit (ICU) stays as well as pre-ICU variables gathered immediately after the trauma. Although this work shown the potential of including dynamic information to predict neurological functional outcomes, they provided an interpretation of the dynamic treatments effects which is purely based on statistical associations. **The objective of this study is to create an interpretable model extending the findings of previous works to a causal framework.** In Figure 4 an overview of the Collaborative European NeuroTrauma Effectiveness Research in Traumatic Brain Injury (CENTER-TBI) [17] study characteristics is reported.

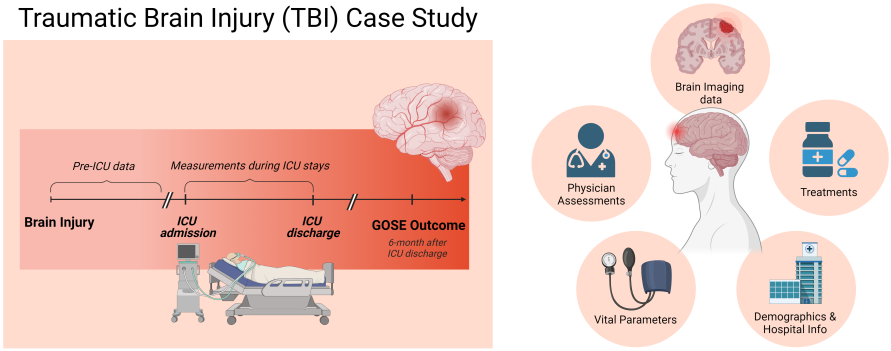

Figure 4: All adult patients from CENTER-TBI cohort with 24h+ ICU stays are included in the study. CENTER-TBI collection comprises data coming from 65 different hospitals around all Europe. The aim here is to predict 6-month GOSE outcomes [17] functionally to a set of 1.6k features collected before and during ICU stays.

The characteristics of CENTER-TBI database lead to some additional problem-specific challenges which I'm confident DBNs could undertake too. Firstly, compared to the cohort of patients from Mangiagalli clinic, in TBI study we are in a high-dimensional scenario. Constructing a model on the entire set of variables is, therefore, intractable. As a consequence, methods of variable selection must be adopted to reduce the dimensionality of the initial set. Moreover, CENTER-TBI data collection is multi-centric. As a result, it could not be hypothesized that all the patients are part of a sole population. Each country has his own protocols and standard procedures to be applied in case of TBIs and each facility may include patients with completely different characteristics. Therefore, in each hospital we could identify variable causal treatment effects. As a consequence, the information of the patients' provenance must be considered and multiple models may be constructed for distinct groups of patients.

6. Conclusions and Questions

My research project focuses on DBNs design and development in the healthcare domain, and is supported by two applications with similar characteristics in terms of the outcome, but facing different domain-specific issues. The first case study revolves around BPD early detection in preterm babies from Mangiagalli clinic. Instead, a second work in collaboration with CENTER-TBI, aims to model long-term effects of TBI in adult patients. I'll have tackle many challenges during the next years, ranging from methodological difficulties to task-specific issues. I'm aware the work plan is ambitious given the complexity of temporal and causal models I plan to build. However, my intention is to focus on one step a time. Firstly I'll work with incomplete observational data, with the aim of modeling the temporal evolution of severe clinical conditions. Then I'll integrate causal and counterfactual reasoning in a discrete-time setting using DBNs. I'm also working closely with neonatologist and neurologist, thus I'll have the support of expert physicians when it comes to interpreting clinical data and results.

I would like to benefit from this occasion to ask for suggestions concerning my research project. Here is a list of input questions to open the discussion:

- What are the main limitations and critical points of my work plan?
- Do you have any suggestions about managing missing data over time in an observational setting?
- Have you read any works revolving around personalised treatments and/or sequential treatments?
- Which conferences or journals could be interested to my work?

Acknowledgement

This work is funded by the National Plan for NRRP Complementary Investments (PNC, established with the decree-law 6 May 2021, n. 59, converted by law n. 101 of 2021) in the call for the funding of research initiatives for technologies and innovative trajectories in the health and care sectors (Directorial Decree n. 931 of 06-06-2022) - project n. PNC0000003 - AdvaNced Technologies for Human-centrEd Medicine (ANTHEM). This work reflects only the authors' views and opinions; neither the Ministry for University and Research nor the European Commission can be considered responsible for them.

References

[1] S. Wilson, et al., Innovative technological advancements in laboratory medicine: Predicting the lab of the future, Biotechnology & Biotechnological Equipment 36 (2022).
[2] S. Newsome, et al., Estimating long-term treatment effects in observational data: A comparison of the performance of different methods under real-world uncertainty, Statistics in Medicine (2017).
[3] L. Goetz, N. Schork, Personalized medicine: motivation, challenges, and progress, Fertility and sterility 109 (2018) 952–963.
[4] M. Khosravi, et al., Artificial intelligence and decision-making in healthcare: A thematic analysis of a systematic review of reviews, Health Services Research and Managerial Epidemiology 11 (2024) 1–15.
[5] E. Hariton, J. Locascio, Randomised controlled trials - the gold standard for effectiveness research: Study design: randomised controlled trials, BJOG: An International Journal of Obstetrics & Gynaecology 125 (2018).
[6] E. Hariton, J. Locascio, Confounding in observational studies evaluating the safety and effectiveness of medical treatments, Kidney360 2 (2021) 1156–1159.
[7] J. Pearl, Causality: Models, Reasoning and Inference, Cambridge University Press, New York, 2000.
[8] D. Umapathy, et al., Perspective of artificial intelligence in disease diagnosis: A review of current and future endeavours in the medical field, Cureus 15 (2023).
[9] R. Daneshjou, et al., Lack of transparency and potential bias in artificial intelligence data sets and algorithms: A scoping review, JAMA dermatology 157 (2021).
[10] K. Murphy, Dynamic bayesian networks: Representation, inference and learning (2002).
[11] R. Schafer, T. Weyrath, Dynamic bayesian networks: Representation, inference and learning, User Modelling: Proceedings of the Sixth International Conference (1997) 377–388.
[12] N. Friedman, The bayesian structural EM algorithm, CoRR abs/1301.7373 (2013).
[13] E. Bancalari, et al., Bronchopulmonary dysplasia: Clinical presentation, The Journal of Pediatrics 95 (1979) 819–823. Workshop on Bronchopulmonary Dysplasia.
[14] E. Jensen, et al., The diagnosis of bronchopulmonary dysplasia in very preterm infants: An evidence-based approach, American Journal of Respiratory and Critical Care Medicine 200 (2019).
[15] J. Bruns, W. Hauser, The epidemiology of traumatic brain injury, Epilepsia 44 (2003) 2–10.
[16] S. Bhattacharyay, et al., Mining the contribution of intensive care clinical course to outcome after traumatic brain injury, npj Digital Medicine 6 (2023).
[17] A. Maas, et al., Collaborative european neurotrauma effectiveness research in traumatic brain injury (center-tbi): A prospective longitudinal observational study, Neurosurgery 76 (2015) 67–80.