6 months research engineer: Policy learning for personalized medicine. Finding the optimal dose of hormone for ovarian stimulation

Julie Josse (Inria) & Mathieu Dellenbach (Elixir)

Summary

Infertility affects 1 in 5 couples of childbearing age.The most common solution is to resort to In Vitro Fertilization. However, one of the main challenges is to determine the initial dose and duration of gonadotropin hormone administration to maximize the number of oocytes obtained at the end of stimulation, under the constraint that estradiol levels must not be too high to avoid hyperstimulation. To tackle this challenge, we will leverage rich observational multi-centric data as well as techniques of causal inference. More precisely, we will consider methods for learning optimal treatment policies and in particular for establishing the appropriate dose and duration of treatment for each patient. One of the challenges will be to propose methods to manage missing data in this framework. Finally, we will consider techniques of dynamic treatment regimes to enrich the analysis with monitoring data.

Keywords: In Vitro Fertilization, Causal inference, policy learning, missing values, dynamic treatment regime

1 Statistical context: Policy learning and missing values

Patient management protocols require adjustments to the individual patient and the organizational context to be effective. Yet, most of the available techniques to estimate the effect of treatment on patients are dedicated to measuring the average treatment effect (ATE), i.e. the effect of the treatment for the population [3]. The graal, the conditional average treatment effect (CATE) function, [1] is well defined from a mathematical point of view and allows us to estimate how treatment effects vary depending on the patient's characteristics. Exploiting the treatment effect heterogeneity allows us to personalize the treatments. The benefits are high but this is still a very hard problem, from a theoretical and applied point of views. This is especially true when you do not target the CATE but the optimal policy [2]: the aim is to learn optimal treatment assignment rules (who should be treated and who should not?) to apply it in the future. Targeting CATE or the best policy are related, but policy learning makes it easier to include constraints, such as fairness ones (you can not include some features), budget constraints (at most 15% of people get treated), deployment constraints and functional form constraints. For instance, in medical domain, decision trees are often used and it will be better accepted by the community to improve the decisions rather than modifying the entire process. We can construct a class of acceptable policies that respects these features. The machine learning community has proposed a variety of practical, high performance algorithms [12] but causal inference can bring theory, i.e., tightened bounds on regret (gap between optimal policy and estimated policy) [15] which provides guidance for algorithm choice.

The primarily aim of the work consists in **learning a decision rule** which is a mapping from the covariate space to the treatment space that maximizes the expected potential outcome for the population in a counterfactual world had this rule been implemented. However, this task is not straightforward as there are many missing values.

The problematic of **missing values** is ubiquitous in data analysis practices and it is exacerbated when aggregating data of different sources. Naive approaches such as complete-case analysis which can lead to important bias, cannot be applied in **high-dimensional settings** when almost all data can be deleted (with only 300 features, and a probability to be missing for each individual and feature of 0.01%, complete case analysis would result in keeping around 5% of the rows). There exists an abundant literature on the topic [7, 14, 5, 9] and many methods available either to estimate some parameters (EM, multiple imputation) or to do supervised learning with missing values [4]. However, in the context of causal inference the literature is scarce, [8, 11, 6, 10] and as far as we know, no solution have been suggested in the framework of policy learning.

A final step in modeling that is closer to the complexity of the data is to consider **dynamic treatment regime** (DTR) [13] and time varying covariates. It consists of a sequence of decision rules, that indicate what to do at every step: how to individualize treatments based on previous treatment choices and covariate history. The benefits of such approaches are immense in improving disease management and in balancing treatment efficacy and risk. Nevertheless, they require strong assumptions of sequential ignorability that are difficult to meet and methods are not meant to handle missing values in this longitudinal case. From the algorithm point of view, there has tight connection with particular case of Reinforcement learning (in a non markovian set up, a batch mode framework, the system dynamics are unknown, etc) which consists of evaluating and learning dynamic policies.

2 Application context and objective: Ovarian Stimulation at Elixir

Patient management protocols is an acute problematic in the field of fertility, specifically for In Vitro Fertilization (IVF) procedures where different patients profile have long been established, but not the optimal treatment. IVF requires to collect multiple oocytes (at least 3) from the woman ovaries, to produce these multiple oocytes patients must undergo an hormonal stimulation phase. Choosing the right hormone, at the right dose, and for the right duration, defines the success of this hormonal stimulation. Physicians have been trying to define optimal personalized protocols and specifically the right dose with diverse statistical tools that are not used in practice. We propose to develop a decision support system based on 2 algorithms: The first algorithm takes as input the patient's clinical and biological characteristics, as well as the initial dose of gonadotropin administered, to estimate as output the number of oocytes obtained. The aim is to establish a dose-response curve based on the patient's characteristics. A second algorithm taking as input the initial clinical and biological characteristics and the results of each monitoring: biological results and follicular diameters (follicles are cellular structures that house oocytes and are observed in the ovaries by ultrasound), to estimate as output the evolution of these same monitoring parameters. This should enable clinicians to anticipate the ideal day for ovulation induction.

Note that the aim is to start by developing a solution for this specific application. However the methods are generic and can be applied for many other questions whether in the medical domain or in other domains such as economy, etc. In addition, even if the project is motivated by practical questions, the project requires strong methodological and theoretical contributions. There is a possibility to pursued in a PhD Cifre at the end of the 6 months contrat.

3 Laboratory - company

The candidate will be supervised by both Julie Josse (expert in Missing Values and Causal Inference), Elixir Chief Technical Officer Xavier Hurst and their Chief Medical Officer Mathieu Dellenbach. Julie Josse has many international connections in causal inference (she was invited to the semester on causality in Berkeley, to the Rousseeuw prize in Belgium, etc.) and often sends her PhD students to do research internships abroad, in particular with the Department of Statistics at Stanford University with whom she has many connections. Xavier Hurst is a DevOps engineer with experience in

Premedical Team - Inria Montpellier The Premedical (Precision Medicine by Data Integration and Causal Learning) team[1], is a recent Inria-Inserm team located in Montpellier. It is an interdisciplinary team composed of statisticians,

[1] https://team.inria.fr/premedical/

biostatisticians, machine learners, and clinicians. Premedical develops methods for optimal treatment policy (drug efficacy, who gets treated and when, etc.) from heterogeneous data (clinical trials, observational data) that come with methodological challenges. In particular, Premedical develops methods for causal inference, statistical learning, management of missing data, federated learning, etc. Premedical holds the missing data and causality research group[2] and has created a taskview on causal inference methods.

Elixir Elixir Health is a french deeptech start-up building a decision support system for fertility physicians. The core of the decision support system will assist them during the most critical steps of the In Vitro Fertilization care pathway, starting with ovarian stimulation. Elixir is located in Strasbourg supported by the Quest for Health Incubator and Paris supported by the Hotel Dieu Incubator from the Public Parisian Hospital (Assistance Publique–Hôpitaux de Paris: APHP). This work is supported and subsidized by the French Public Investment Bank (BPI) through the i-Lab funding.

4 Contact

We are looking for excellent candidates, highly motivated, with background knowledge in mathematics, statistics /machine learning and interested by interdisciplinary research and collaboration. We will focus on both the theoretical and practical aspects including implementation.

We are offering a research engineer position at INRIA (Montpellier) or at Elixir for a duration of six months, ideally starting in September 2023 (duration and starting dates can be discussed). This can be pursued by a PhD thesis, depending on the overall quality of the internship.

Qualifications:

- Master in Statistics, Machine Learning, Biostatistics, Data-science or related fields
- Experience with machine learning and high-dimensional statistics
- Strong statistical computing skill
- Excellent writing and communication skills
- Background on causal inference is a plus
- Strong interest for the medical domain

Required application materials:

- CV

[2] https://misscausal.gitlabpages.inria.fr/misscausal.gitlab.io/

- Academic performance in recent years
- Complete contact information for two references (that will be contacted if needed).

Interested candidates should apply as early as possible since the positions will be filled when suitable candidates are found.

Email your application to julie.josse@inria.fr and ...

References

[1] Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.

[2] Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1):133–161, 2021.

[3] Miguel A Hernán and James M Robins. *Causal Inference*. Chapman & Hall/CRC, 2019.

[4] Julie Josse, Nicolas Prost, Erwan Scornet, and Gaël Varoquaux. On the consistency of supervised learning with missing values. *arXiv preprint*, 2019.

[5] Julie Josse and Jerome P. Reiter. Introduction to the special section on missing data. *Statist. Sci.*, 33(2):139–141, 05 2018.

[6] N. Kallus, X. Mao, and M. Udell. Causal inference with noisy and missing covariates via matrix factorization. In *Advances in Neural Information Processing Systems*, pages 6921–6932, 2018.

[7] R. J. A. Little and D. B. Rubin. *Statistical Analysis with Missing Data*. Wiley, 2002.

[8] Alessandra Mattei and Fabrizia Mealli. Estimating and using propensity score in presence of missing background data: an application to assess the impact of childbearing on wellbeing. *Statistical Methods and Applications*, 18(2):257–273, 2009.

[9] Imke Mayer, Aude Sportisse, Julie Josse, Nicholas Tierney, and Nathalie Vialaneix. R-miss-tastic: a unified platform for missing values methods and workflows, 2021.

[10] Imke Mayer, Erik Sverdrup, Tobias Gauss, Jean-Denis Moyer, Stefan Wager, and Julie Josse. Doubly robust treatment effect estimation with missing attributes. *Ann. Appl. Statist.*, 14(3):1409–1431, 2020.

[11] Shaun Seaman and Ian White. Inverse probability weighting with missing predictors of treatment assignment or missingness. *Communications in Statistics-Theory and Methods*, 43(16):3499–3515, 2014.

[12] A. Swaminathan and T. Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning (ICML)*, pages 814–823, 2015.

[13] Anastasios A Tsiatis. *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. CRC press, 2019.

[14] S. van Buuren. *Flexible Imputation of Missing Data*. Chapman and Hall/CRC, Boca Raton, FL, 2018.

[15] Zhengyuan Zhou, Susan Athey, and Stefan Wager. Offline multi-action policy learning: Generalization and optimization. *arXiv preprint arXiv:1810.04778*, 2018.