Extended Abstract: Personalised Data-driven Healthcare*

Alexander Galozy[1] and Sławomir Nowaczyk[1]

I. INTRODUCTION

Medical advancement in the last century has led to an unprecedented increase in the populace's quality of life and lifespan. As a result, an ever-increasing number of people living with chronic health conditions require costly long-term treatment [1]. One promising direction to tackle the issue is the patient's active involvement in managing their care [2]. With the advent of intelligent devices, it is easier than ever to provide personalized digital interventions to patients and help them manage their treatment in their daily lives. If such new approaches are to succeed, scalability is necessary, and solutions are needed that can act autonomously without costly human intervention.

This Ph.D. was conducted as part of the project "improving Medication Adherence through Person-Centred Care and Adaptive Interventions (iMedA)" which aims to provide personalized adaptive interventions to hypertensive patients. The focus lies on inadequate Medication Adherence (MA), a pervasive issue where patients do not take their medication as prescribed. A schematic overview of the system architecture is shown in figure 1. The doctoral thesis is tackling the issue of providing personalized adaptive interventions in the domain of mHealth. Focus is placed on the problem of providing these interventions under domain-specific constraints.

II. RESEARCH QUESTIONS

R1: How can MA be measured? Individuals needing treatment support must be identified promptly. Timeliness and effectiveness of pharmacological treatment are tightly linked for many cardiovascular diseases, especially for secondary or tertiary prevention [3]. Current approaches aim to measure the level of MA throughout treatment and provide interventions as nonadherence is detected. The key challenges lie in the accurate and cost-effective measurement of adherence. Many measures of different types have been developed over the years. The most accurate way to measure MA is by measuring the metabolite concentration of the drug in body fluids, but it is expensive and burdensome for patients. Digital pills offer a less invasive alternative, but their wide-scale adoption and acceptance from care providers and patients are needed [4]. Indirect measures like medication refill adherence from electronic health records and pharmacy dispensation databases can approximate MA. Primary nonadherence can be readily observable from complete data sources, allowing for earlier interventions. The measurement

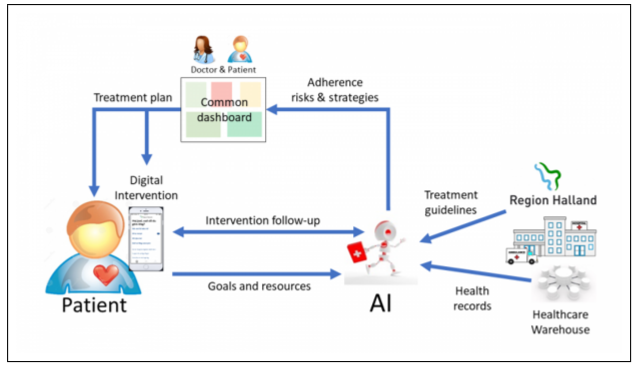

Fig. 1. Schematic view of iMedA

of refill adherence using EHRs and pharmacy records is prone to pitfalls that would affect the pool of candidates for intervention significantly [5]. Given that indirect adherence estimates may drive future studies, it is vital to investigate how data-related and measure-related issues affect MA approximation and how to alleviate them. Work on this research question has resulted in a published journal article [6]. We examined common pitfalls in refill adherence estimations using EHRs, showing a small yet statistically significant effect on population averages and a large effect in individual cases.

R2: What is the probability of MA? The prediction of MA allows the care provider to estimate the probability of nonadherence to medication. Training prediction models and analyzing adherence patterns in EHRs and pharmacy records would allow physicians to intervene early. Previous studies utilizing sociodemographic factors, clinical factors, or purchasing information had only marginal success in predicting long-term MA [7]. These works point out the importance of analyzing patients' refilling behaviors and patterns of healthcare utilization. We identified a research gap concerning the predictability of adherence using EHRs, focusing specifically on healthcare utilization factors and analyzing temporal patterns of refill MA. Work on this research question has resulted in a published journal article [8]. While model performance is relatively high, these models might not be suitable for selecting patients for intervention since performance is relatively low in interesting scenarios, such as predicting a sudden change in adherence, showing the apparent bias of models towards patients with high healthcare utilization. Furthermore, we discovered common patterns of refill adherence. Specific medication consumption patterns result in similar pickup patterns, obfuscating potential pathological patterns of medication-taking that can be relevant for

*Research funded by: Vinnova, grant number: 2017-04617

[1]A. Galozy and S. Nowaczyk are with CAISR, Halmstad University, Halmstad, Sweden, `alexander.galozy at hh.se`, `slawomir.nowaczyk at hh.se`

intervention.

R3: How can we provide adaptive and personalized intervention in mHealth? We argue that Reinforcement Learning (RL) is an excellent fit and promises to deliver state-of-the-art solutions for mHealth. The capability of RL methods for lifelong learning equips autonomous agents with the ability to adapt to changing circumstances of the patient quickly, promising higher patient engagement and better health outcomes through more relevant interventions.

While promising, several challenges in mHealth make the straightforward application of contemporary RL algorithms difficult. Some challenges include the requirement of a good initial policy avoiding too frequent or irrelevant interventions and robustness to failure of algorithmic assumptions [9]. We aim to develop algorithms and methods to address some of these challenges.

We investigate a novel variant of the contextual bandit problem with corrupted context motivated by mHealth applications [10]. We develop a meta-algorithm, which uses a "referee" that dynamically combines the policies of a contextual bandit and a multi-armed bandit. Our approach shows promising empirical performance on various simulated and real-world datasets.

Providing good intervention promptly is crucial. We further investigate the latent bandit problem, where a hidden state governs rewards. In real-world applications, agents often have access to imperfect contexts, which can help uncover the latent state. However, partial observability of the state may require occasional information-gathering, which can come at an additional cost. We propose a deliberate selection of actions to effectively uncover the latent state, allowing the agent to balance its knowledge about the state space and maximize cumulative reward. Our findings have been published in a recent journal article showing promising results on synthetic and real-world datasets [11].

In our most recent work, we are interested in important privacy aspects in latent bandits. In particular, we are interested in how differential privacy affects a user's de-anonymization level in data records, as measured through the score-board algorithm proposed by [12], and what impact on regret performance we can expect. Differential privacy is a concept in data privacy that focuses on limiting the disclosure of information from datasets, even when an adversary has access to most of the other data in the set [13]. This approach adds controlled noise to the data, making it difficult for an attacker to determine whether a particular user's data is in the set. This technique has been applied to various areas, including machine learning, where it preserves individual users' data privacy while allowing the system to learn from the overall dataset. We investigate the scenario where a new user (and potential adversary) joins a user pool and requests a transition matrix used in the latent-bandit algorithm to select actions. The reward distribution is known to the decision maker. The goal of each user in the pool is to protect their own transition matrix by adding noise through an arbitrary noise mechanism. We investigate different strategies for providing a transition matrix to a new user, such as nearest neighbor, cell-wise nearest neighbor, and cluster centers. In addition, we evaluate the ADS-GAN measure of privacy, a promising technique for data synthesis using generative adversarial networks, as proposed by [14]. We use this measure to assess the level of privacy in our experiments. Our experiments show that simply adding noise to existing records is insufficient for optimizing the regret-privacy trade-off. More sophisticated aggregation and privatization techniques are required to protect the transition matrix. We also demonstrate that using the ADS-GAN measure of privacy does not correlate well with the level of de-anonymization in data linkage attacks. It often overestimates the level of privacy, leaving users' private data vulnerable to de-anonymization. Our results have profound implications for mHealth where acceptance of new AI-based technologies requires strong privacy guarantees while maintaining good performance.

REFERENCES

[1] P. A. Grady and L. L. Gough. Self-management: a comprehensive approach to management of chronic conditions. *American Journal of Public Health*, 104(8):25–31, Aug 2014.

[2] C. Dowrick, M. Dixon-Woods, H. Holman, and J. Weinman. What is chronic illness? *Chronic Illness*, 1(1):1–6, Mar 2005.

[3] F D R Hobbs. Cardiovascular disease: different strategies for primary and secondary prevention? *Heart*, 90(10):1217–1223, 2004.

[4] T. Peters-Strickland, L. Pestreich, A. Hatch, S. Rohatagi, R. A. Baker, J. P. Docherty, L. Markovtsova, P. Raja, P. J. Weiden, and D. P. Walling. Usability of a novel digital medicine system in adults with schizophrenia treated with sensor-embedded tablets of aripiprazole. *Neuropsychiatric Disease and Treatment*, 12:2587–2594, 2016.

[5] Andrea M. Burden, J. Michael Paterson, Andrea Gruneir, and Suzanne M. Cadarette. Adherence to osteoporosis pharmacotherapy is underestimated using days supply values in electronic pharmacy claims data. *Pharmacoepidemiology and Drug Safety*, 24:67–74, 2014.

[6] Alexander Galozy, Slawomir Nowaczyk, Anita Sant'Anna, Mattias Ohlsson, and Markus Lingman. Pitfalls of medication adherence approximation through ehr and pharmacy records: Definitions, data and computation. *International Journal of Medical Informatics*, 136:104092, 2020.

[7] A.V. Prochazka J.F. Steiner. The assessment of refill compliance using pharmacy records: Methods, validity, and applications. *Journal of Clinical Epidemiology*, 50:105–116, 1997.

[8] Alexander Galozy and Slawomir Nowaczyk. Prediction and pattern analysis of medication refill adherence through electronic health records and dispensation data. *Journal of Biomedical Informatics: X*, 6-7:100075, 2020.

[9] Ambuj Tewari and Susan A. Murphy. *From Ads to Interventions: Contextual Bandits in Mobile Health*. Springer International Publishing, Cham, 2017.

[10] Alexander Galozy, Slawomir Nowaczyk, and Mattias Ohlsson. Corrupted contextual bandits with action order constraints. *Engineering Systems with Applications*, 2023, under peer-review.

[11] Alexander Galozy and Sławomir Nowaczyk. Information-gathering in latent bandits. *Knowledge-Based Systems*, 260:110099, 2023.

[12] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125, 2008.

[13] Cynthia Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation—TAMC*, volume 4978 of *Lecture Notes in Computer Science*, pages 1–19. Springer Verlag, April 2008.

[14] J. Yoon, L. N. Drumright, and M. van der Schaar. Anonymization Through Data Synthesis Using Generative Adversarial Networks (ADS-GAN). *IEEE J Biomed Health Inform*, 24(8):2378–2388, Aug 2020.