Investigating Gender Fairness in Machine Learning-driven Personalized Care for Chronic Pain

Pratik Gajane
Department of Mathematics and Computer Science,
Eindhoven University of Technology,
p.gajane@tue.nl

Sean Newman
School of Public Health,
University of Michigan,
scnewman@umich.edu

John D. Piette
Department of Internal Medicine, School of Medicine, and
School of Public Health,
University of Michigan,
jpiette@umich.edu

February 29, 2024

Abstract

This study investigates gender fairness in personalized pain care recommendations using machine learning algorithms. Leveraging a contextual bandits framework, personalized recommendations are formulated and evaluated using LinUCB algorithm on a dataset comprising interactions with 164 patients across 10 sessions each. Results indicate that while adjustments to algorithm parameters influence the quality of pain care recommendations, this impact remains consistent across genders. However, when certain patient information, such as self-reported pain measurements, is absent, the quality of pain care recommendations for women is notably inferior to that for men.

1 Introduction

Pain is an unpleasant sensation and emotional experience that leads to poor quality of life for millions of people worldwide. Chronic pain syndrome is a prevalent and increasingly common problem in many countries [Zimmer et al., 2022]. Chronic pain, as indicated by a World Health Organization study, is associated with a four-fold increase in the occurrence of depression or anxiety [Slack, 2022]. Chronic pain frequently disrupts individuals' capacity to focus, eat, and sleep. Considering the significant impact of chronic pain on individuals and society, a variety of machine learning solutions have been proposed to provide effective pain care. See Matsangidou et al. [2021] for a comprehensive survey. While there have been advancements in comprehending and treating chronic pain, providing effective pain care continues to be a major problem. A significant challenge in effectively dealing with chronic pain lies in the considerable variability among patients, both in their response to treatments and their susceptibility to adverse effects [Ablin and Buskila, 2012]. This variability hampers the efficacy of generic interventions, emphasizing the necessity for more personalized approaches to address the intricate nature of chronic pain.

Chronic pain care necessitates treatments over a long period of time taking the feedback from patients into consideration. Thus reinforcement learning algorithms, especially suited to make effective sequential decisions using feedback, provide a suitable solution for machine learning-driven pain care. For example, Komorowski et al. [2018] used reinforcement learning solutions to recommend personalized treatments that are on average more reliable than human clinicians. Saria [2018] proposed reinforcement learning methods for individualized treatment strategies to correct hypotension in Sepsis. Roggeveen et al. [2021] propose similar solutions for optimizing hemodynamic treatment for critically ill patients with Sepsis. As evidenced by these articles, the ability of reinforcement learning algorithms to learn effectively from user feedback makes them particularly suitable for making personalized decisions.

However, a drawback of the ability to make personalized decisions is the potential for introducing biases related to protected attributes such as gender. It has been seen that the design of the majority of machine learning algorithms ignores the sex and gender dimension and its contribution to health differences among individuals [Cirillo et al., 2020]. In particular, machine learning tools for healthcare that use unisex treatments may disadvantage female patients [Straw and Wu, 2022].

1.1 Our Contributions

In this study, we investigate gender fairness in personalized pain care recommendations made by using a machine learning algorithm. We utilize the dataset Piette [2022] which is also used in an extensive study on machine-learning driven pain care Piette et al. [2022]. To allow for personalized recommendations, we model the problem at hand using *contextual bandits*, a reinforcement learning problem formulation well-suited to deal with personalized sequential decisions. In this model, the algorithm leverages a range of patient-specific details to tailor its recommendations. In our experiments, we use LinUCB Li et al. [2010], an algorithm that serves as the cornerstone of leading solutions for contextual bandits. Our results show that although adjustments to the parameter utilized by LinUCB result in changes to the quality of pain care recommendations produced by the algorithm, this impact is consistent across both men and women. However, when the algorithm lacks access to certain patient information, notably self-reported pain measurements, the quality of pain care recommendations for women is discernibly worse compared to men.

2 Formalizing Pain Care as a Reinforcement Learning Problem

We model personalized pain care using *contextual bandits* – a reinforcement learning formalization that has been used toward the goal of personalized decision-making in many domains including healthcare. In a contextual bandits problem, at each decision step $t = 1, 2, \dots, T$:

- The algorithm observes a context i.e., some relevant information in the form of a set of feature values.
- Using this contextual information and previous observed feedback, the algorithm selects an option a_t from the set of available options $\mathcal{A}$. After taking an action, the algorithm receives a (possibly randomized) numerical value r_{t,a_t} as feedback. In reinforcement learning terminology, options and feedback are called *actions* and *rewards* respectively. It is important to emphasize that the received reward r_{t,a_t} only corresponds to the action taken a_t and no reward is observed for unchosen actions. Reward r_{t,a_t} informs the algorithm about the goodness of a_t. The higher the received reward, the better the corresponding action.
- The algorithm may choose to improve its action selection strategy with the new observation – {observed context, selected action, received reward}.

With this formulation, the algorithm's goal can be expressed as maximizing the cumulative sum of received rewards $\sum_{t=1}^{T} r_{t,a_t}$. Alternatively, the algorithm's goal can also be expressed as minimizing *regret* formally defined by

$$\mathbb{E}\left[\sum_{t=1}^{T} r_{t,a_t^*}\right] - \mathbb{E}\left[\sum_{t=1}^{T} r_{t,a_t}\right],$$

where $\mathbb{E}$ denotes the expectation and a_t^* is the action with the maximum expected reward at decision step t. Regret can be understood as the total mistake cost.

For personalized pain care, a decision step corresponds to the arrival of a patient. Context can be understood as the relevant patient information. Actions are the available pain recommendations and feedback corresponds to the effects of the given recommendation on the patient.

3 Materials and Methods

First, in Section 3.1, we provide details about the dataset we used in our experiments. Then in Section 3.2, we provide a brief overview of the LinUCB algorithm employed in our study.

3.1 Data

We make use of the dataset provided in Piette [2022]. It contains information about interactions with 164 patients each receiving pain care recommendations across 10 sessions. For each patient interaction, the feature values noted in the dataset are as follows :

- Percentage of days in the current week with steps goal met (0-1)
- Pain intensity change (0-1)
- CBT skill practice this week (0-10)
- Sleep quality this week (0-10)
- Sleep duration this week (0-22)
- Pain interfere1 (0-10)
- Pain interfere2 (0-10)
- Session Number (1-10)

In the above, Pain interfere1 represents a self-reported response to the query – "What number best describes how much pain has interfered with your enjoyment of life today, with 0 meaning pain does not interfere and 10 meaning pain completely interferes?". Moreover, Pain interfere2 is a self-reported response to the query – "What number best describes how much pain has interfered with your general activity today, with 0 meaning pain does not interfere and 10 meaning pain completely interferes?"

The recommendation options (or actions, in the parlance of reinforcement learning) are as follows :

- Option 1 : Interactive voice response (IVR) call. During IVR calls, patients hear a recorded message from their therapist.
- Option 2 : A 15-minute telephone session with the therapist.
- Option 3 : A 45-minute telephone session with the therapist.

It is apparent that option 2 is more human-resource intensive than option 1, while option 3 is the most human-resource intensive among them. To address this disparity, Piette et al. [2022] recommend discounting the expected rewards of option 2 and option 3 with additive factors of -0.02 and -0.06 respectively.

For each patient interaction, the entry recorded in the dataset contains values of all the seven features mentioned above, the recommended action, and the observed reward.

Data Preprocessing: Firstly, we normalized all the feature values apart from the ordinal session numbers into the range $[0, 1]$. Using this logged dataset directly for our evaluation experiments presents the following challenge. The rewards are only recorded for the actions chosen by the logging policy, which are often likely to differ from the actions chosen by the algorithm being evaluated. As a solution, we build a simulator to model patient interactions using this dataset as a foundation. Concretely, we use the patient interactions [Piette, 2022] to build a set of weights for each action which transform the values for the features into the expected reward. Using these weights, rewards are obtained for any set of feature values with some noise being added to the expected reward. We used zero-mean Gaussian noise with a standard deviation equal to the standard deviation observed in the recorded rewards for the corresponding action in the dataset [Piette, 2022].

The dataset [Piette, 2022] divides the patients into three clusters based on subgroups of patients with different levels of pain-related functioning at enrollment. To simulate more patient interactions than given in [Piette, 2022], average feature values for men and women for each session and each cluster are calculated. The feature values are simulated using the average feature values with an additive noise. We used zero-mean Gaussian noise with a standard deviation equal to the standard deviation observed in the recorded feature values. While simulating new patients, we ensure that the distribution of the total population across the three clusters remains consistent with the dataset: Cluster 1 accounts for 24%, Cluster 2 for 42%, and Cluster 3 for 34%. We additionally ensure that the gender distribution within each cluster mirrors that of the original dataset: Cluster 1 comprises 20% women, Cluster 2 comprises 7.1% women, and Cluster 3 comprises 13.8% women.

3.2 Algorithm : LinUCB

For our experiments, we use the algorithm LinUCB introduced by Li et al. [2010]. LinUCB operates by maintaining a set of linear models, each corresponding to a distinct action, and updates these models based on observed features and associated rewards. LinUCB employs an upper confidence bound strategy to balance exploration and exploitation, wherein it chooses actions with higher uncertainty within the bounds of linear model predictions, thus exploring potentially advantageous actions while leveraging existing knowledge. By iteratively updating model parameters and selecting actions according to their estimated values and uncertainty, LinUCB aims to maximize cumulative rewards over time. We chose this algorithm as it is highly influential and forms the basis of other solutions proposed in the literature. Accordingly, we expect that the results obtained using LinUCB will have wide-reaching pertinence and will also be relevant while using other solutions based on LinUCB.

4 Approach and Results

We examine whether differing parameter values influence gender fairness and if the inclusion or exclusion of certain information in context affects gender fairness

4.1 Do differing parameter values affect gender fairness?

The only parameter used by LinUCB is the value α. The role of α in the algorithm design is to control the *exploration* and low values of α lead to less exploration. In the context of reinforcement learning algorithms such as LinUCB, exploration refers to the tendency of algorithms to select actions that appear suboptimal according to the information available then. Exploration is performed with the objective of acquiring information that is to see how profitable the actions are. A reinforcement learning algorithm must perform exploration to effectively navigate unknown environments, especially in its early stages when it lacks prior knowledge of action profitability. Inadequate exploration may cause the algorithm to overlook highly rewarding actions. Conversely, excessive exploration can impede the algorithm's long-term objective of maximizing cumulative reward by sacrificing immediate rewards. Thus, achieving an optimal balance of exploration is crucial for maximizing long-term performance.

In the context of this study, the effect of exploration can potentially be even more out-sized for women as they number 12.5% of the entire population. Accordingly, in this set of experiments, we investigate if changing the value of α affects gender fairness. For each value of α in $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1\}$, we run LinUCB 100 times for 300 and 500 patients with 10 sessions each resulting into 3000 and 5000 interactions respectively. In each of these population sizes, the number of women is limited to 12.5% to imitate the fraction size as in the dataset Piette [2022]. We note that the rewards continue to rise up to 5000 interactions at which point reward scores stabilize, which matches the observation made by Piette et al. [2022].

All the results shown in this article are averaged over 100 runs. Below, we show representative results for $T = 3000$ and $T = 5000$. We use the following three performance metrics:

- Average reward for men (resp. women) = Total reward received by men (resp. women)/number of men (resp. women)
- Fraction of suboptimal action selections for men (resp. women) = Number of times a suboptimal action is selected for men (resp. women)/total number of action selections for men (resp. women)
- Per-step expected regret for men (resp. women)

As you can see in Figures 1, 2, 3, 4, 5 and 6, varying α affects the performance at a broader level as expected. When the exploration parameter α is too small, there was likely insufficient exploration, the algorithm failed to identify the optimal actions which in turn led to low performance. On the other hand, when exploration parameter α is too large,

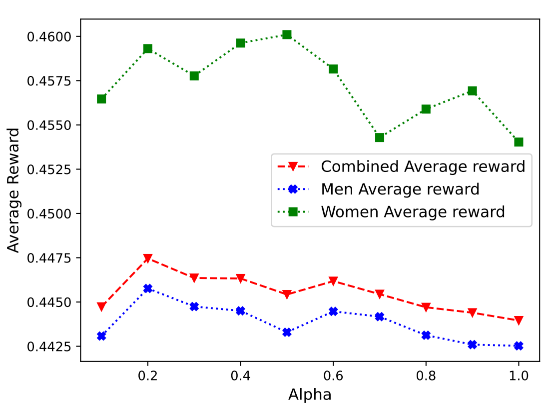

Figure 1: Average Reward for $T = 3000$

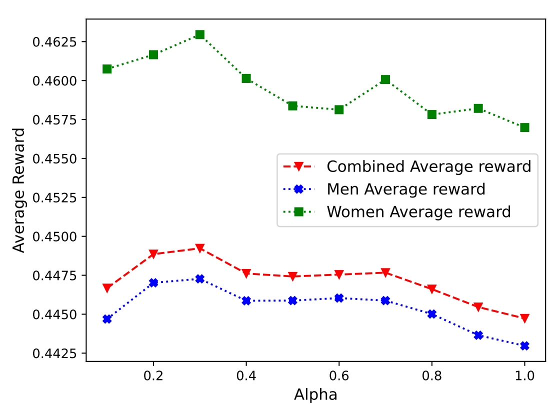

Figure 2: Average Reward for $T = 5000$

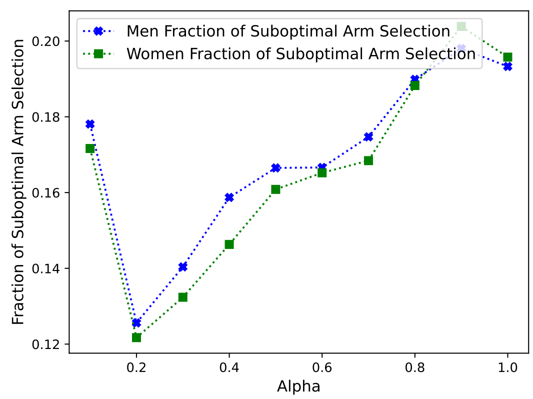

Figure 3: Fraction of Suboptimal Action Selections for $T = 3000$

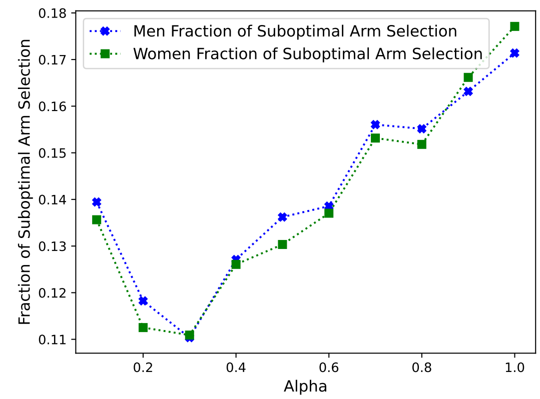

Figure 4: Fraction of Suboptimal Action Selections for $T = 5000$

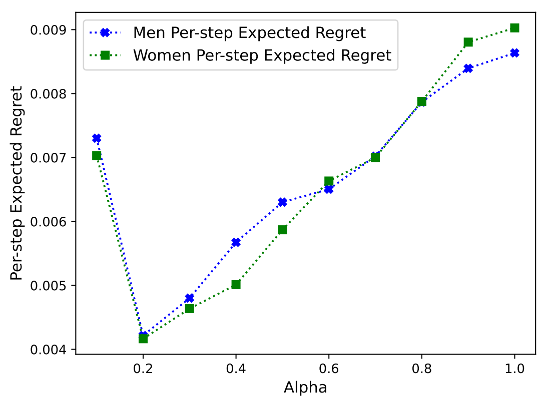

Figure 5: Per-step Regret for $T = 3000$

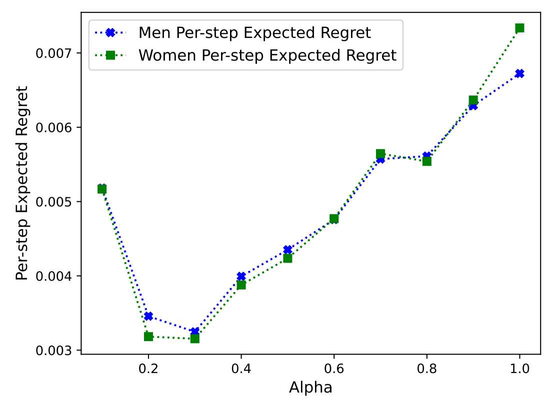

Figure 6: Per-step Regret for $T = 5000$

there was likely excessive exploration, and thus the algorithm wasted some opportunities to select highly rewarding actions which in turn led to low performance.

When it comes to gender fairness, a discernible pattern is not seen in the results. From Figures 1 and 2, it can be seen that the algorithm seems to consistently obtain higher average rewards for women than than for men. However, this does not necessarily mean the algorithm performs better for women than men. Another plausible reason for the results seen in Figures 1 and 2 could be that the optimal actions produce better rewards for women than for men. To circumvent this possibility, results in Figures 3, 4 show the fraction of suboptimal actions selected for men and women respectively. Whereas results in Figures 5 and 6, show per-step regret which can be understood as the rewards foregone due to the algorithm's selection of a suboptimal action. Considering the performance metrics of the fraction of suboptimal actions and per-step regret, the algorithm demonstrates comparable performance for men and women. In particular for the best values of α (0.2 for $T = 3000$ and 0.3 for $T = 5000$), the algorithm demonstrates almost equivalent performance for men and women.

4.2 Does including or excluding certain information in context affect gender fairness?

In the second set of experiments, we see how the features included in the context information can affect gender fairness. We hypothesize that particularity features related to pain will have a significant influence on gender fairness. This is taking into consideration that previous studies reported that genders do not feel pain the same way [Dance, 2019].

We evaluated gender fairness while using LinUCB with the sets of features described in Table 1. For $T = 3000$, we use the the best performing value of $\alpha = 0.2$ as seen in Figures 1, 3 and 5. Analogously, For $T = 3000$, we use the the best performing value of $\alpha = 0.3$ as seen in Figures 2, 4 and 6.

Feature Set	Included Features
Set 1	% of days in current week with steps goal met, Pain intensity change, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Pain interfere1, Pain interfere2, Session Number
Set 2	% of days in current week with steps goal met, Pain intensity change, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Session Number
Set 3	% of days in current week with steps goal met, Pain intensity change, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Pain interfere2, Session Number
Set 4	% of days in current week with steps goal met, Pain intensity change, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Pain interfere1, Session Number
Set 5	Pain intensity change, Sleep quality this week, Sleep duration this week, Pain interfere1, Pain interfere2, Session Number
Set 6	% of days in current week with steps goal met, CBT skill practice this week, Session Number
Set 7	% of days in current week with steps goal met, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Pain interfere1, Pain interfere2, Session Number
Set 8	% of days in current week with steps goal met, Pain intensity change, CBT skill practice this week, Pain interfere1, Pain interfere2, Session Number

Set 9	Pain intensity change, CBT skill practice this week, Sleep quality this week, Sleep duration this week, Pain interfere1, Pain interfere2, Session Number

Table 1: Feature Sets and Included Features

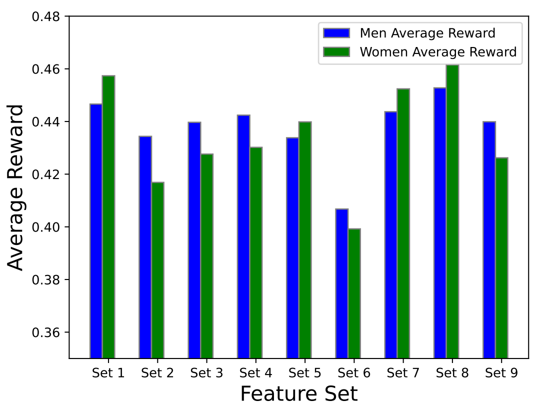

Figure 7: Average Reward for $T = 3000$ and Feature Sets $1 - 9$

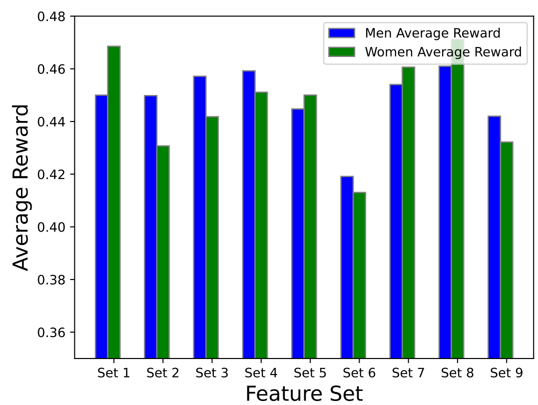

Figure 8: Average Reward for $T = 5000$ and Feature Sets $1 - 9$

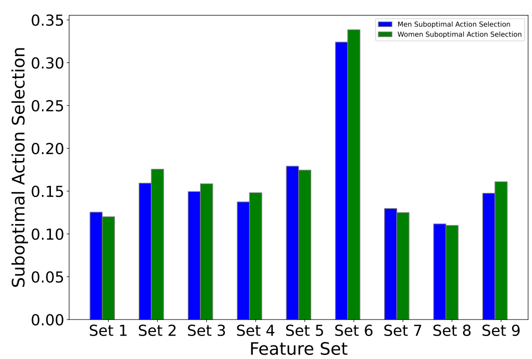

Figure 9: Fraction of Suboptimal Action Selections for $T = 3000$ and Feature Sets $1 - 9$

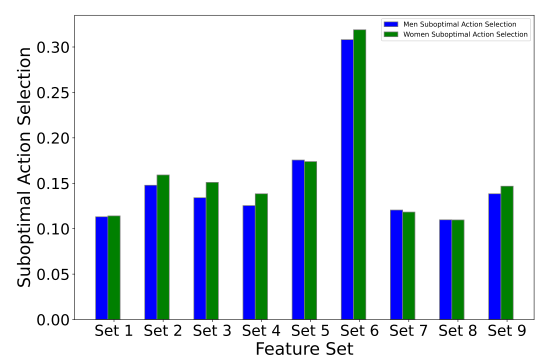

Figure 10: Fraction of Suboptimal Action Selections for $T = 5000$ and Feature Sets $1 - 9$

The results can be seen in Figures 7, 8, 9, 10, 11 and 12. The results demonstrate that for all the performance metrics, the algorithm demonstrates better performance for men as compared to women when features titled “Pain interfere1", “Pain interfere2" or “% of days in the current week with steps goal met" are excluded. This is particularly noteworthy as, for all combinations of features examined, the algorithm consistently exhibits similar performance for men and women. As previously noted, the differences in the quality of pain care recommendations, as measured by the fraction of suboptimal actions and per-step regret, are especially insightful.

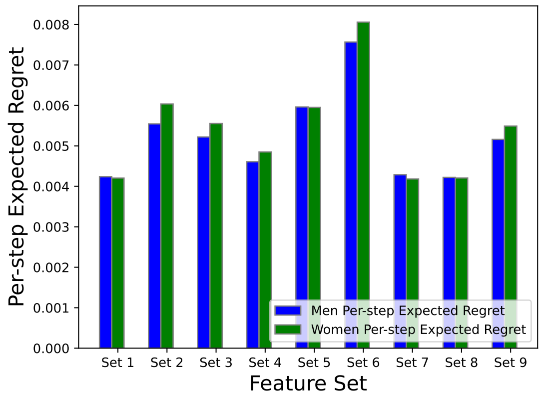

Figure 11: Per-step Regret for $T = 3000$ and Feature Sets $1 - 9$

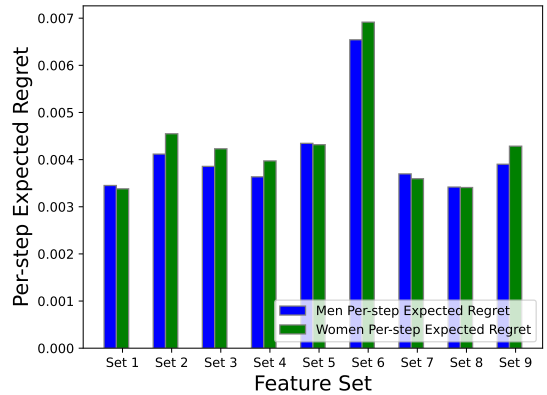

Figure 12: Per-step Regret for $T = 5000$ and Feature Sets $1 - 9$

5 Concluding Remarks

In summary, our study investigates gender fairness in personalized pain care recommendations, highlighting how machine learning algorithms can cause gender disparities in healthcare outcomes. Our findings highlight the importance of incorporating appropriate patient information into reinforcement learning-based pain care recommendations. Omitting specific data, such as self-reported pain measurements, can diminish the quality of recommendations for women relative to men.

References

J. N. Ablin and D. Buskila. Personalized treatment of pain. *Current Rheumatology Reports*, 15(1):298, Dec 2012. ISSN 1534-6307. doi: 10.1007/s11926-012-0298-7. URL `https://doi.org/10.1007/s11926-012-0298-7`.

D. Cirillo, S. Catuara-Solarz, C. Morey, E. Guney, L. Subirats, S. Mellino, A. Gigante, A. Valencia, M. J. Rementeria, A. S. Chadha, and N. Mavridis. Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare. *npj Digital Medicine*, 3(1):81, Jun 2020. ISSN 2398-6352. doi: 10.1038/s41746-020-0288-5. URL `https://doi.org/10.1038/s41746-020-0288-5`.

A. Dance. Why the sexes don't feel pain the same way. *Nature*, 567(7749):448–450, March 2019. doi: 10.1038/d41586-019-00895-. URL `https://ideas.repec.org/a/nat/nature/v567y2019i7749d10.1038_d41586-019-00895-3.html`.

M. Komorowski, L. A. Celi, O. Badawi, A. C. Gordon, and A. A. Faisal. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720, Nov 2018. ISSN 1546-170X. doi: 10.1038/s41591-018-0213-5. URL `https://doi.org/10.1038/s41591-018-0213-5`.

L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, page 661–670, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605587998. doi: 10.1145/1772690.1772758. URL `https://doi.org/10.1145/1772690.1772758`.

M. Matsangidou, A. Liampas, M. Pittara, C. S. Pattichi, and P. Zis. Machine learning in pain medicine: An up-to-date systematic review. *Pain and Therapy*, 10(2):1067–1084, Dec 2021. ISSN 2193-651X. doi: 10.1007/s40122-021-00324-2. URL `https://doi.org/10.1007/s40122-021-00324-2`.

J. Piette. Dataset #1 for piette et al, data in brief: Data for a reinforcement learning intervention to treat chronic pain. Mendeley Data, 2022. URL `https://data.mendeley.com/datasets/p67hy873px/1`.

J. D. Piette, S. Newman, S. L. Krein, N. Marinec, J. Chen, D. A. Williams, S. N. Edmond, M. Driscoll, K. M. LaChappelle, M. Maly, H. M. Kim, K. B. Farris, D. M. Higgins, R. D. Kerns, and A. A. Heapy. Artificial intelligence (ai) to improve chronic pain care: Evidence of ai learning. *Intelligence-Based Medicine*, 6:100064, 2022. ISSN 2666-5212. doi: https://doi.org/10.1016/j.ibmed.2022.100064. URL `https://www.sciencedirect.com/science/article/pii/S2666521222000175`.

L. Roggeveen, A. el Hassouni, J. Ahrendt, T. Guo, L. Fleuren, P. Thoral, A. R. Girbes, M. Hoogendoorn, and P. W. Elbers. Transatlantic transferability of a new reinforcement learning model for optimizing haemodynamic treatment for critically ill patients with sepsis. *Artificial Intelligence in Medicine*, 112:102003, 2021. ISSN 0933-3657. doi: https://doi.org/10.1016/j.artmed.2020.102003. URL https://www.sciencedirect.com/science/article/pii/S0933365720312689.

S. Saria. Individualized sepsis treatment using reinforcement learning. *Nature Medicine*, 24(11):1641–1642, Nov 2018. ISSN 1546-170X. doi: 10.1038/s41591-018-0253-x. URL https://doi.org/10.1038/s41591-018-0253-x.

G. Slack. Reforming pain. *Anesthesiology & Pain Management*, 2022. URL https://stanmed.stanford.edu/individualizing-pain-treatments/.

I. Straw and H. Wu. Investigating for bias in healthcare algorithms: a sex-stratified analysis of supervised machine learning models in liver disease prediction. *BMJ Health & Care Informatics*, 29(1), 2022. doi: 10.1136/bmjhci-2021-100457. URL https://informatics.bmj.com/content/29/1/e100457.

Z. Zimmer, K. Fraser, H. Grol-Prokopczyk, and A. Zajacova. A global study of pain prevalence across 52 countries: examining the role of country-level contextual factors. *PAIN*, 163(9), 2022.