

Personalized, Fair and Adaptive Recommendations for Chronic Pain Care

Pratik Gajane

Abstract

This research proposal addresses the challenges in providing effective chronic pain care through a personalized, fair, and adaptive approach using sequential decision-making. Chronic pain significantly impacts individuals' well-being, necessitating tailored treatments. The proposal outlines three key concerns: personalization to account for patient variability, fairness to address disparities in pain care, and adaptivity to accommodate changing factors influencing treatment efficacy over time. Leveraging sequential decision-making algorithms, this study aims to devise solutions for both stationary and non-stationary settings, encompassing two distinct scenarios: firstly, when the characteristics of changes are known, and secondly, when they are not known. A distinctive feature of our study is the consideration of delayed feedback, reflecting the real-life complexities of chronic pain care, where treatment responses may take some time to manifest. The evaluation plan includes simulation studies assessing utility, personalization, fairness, adaptivity, and the utility-fairness trade-off, contributing to the advancement of AI-driven chronic pain management.

1 Introduction and Related Work

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. In the following, we briefly describe the three major concerns related to chronic pain care that we plan to address using machine learning techniques.

- **Personalization.** 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.
- **Fairness.** Zimmer et al. [2022] highlight inequalities in chronic pain syndromes across various factors such as income and sex. In this proposal, we focus our attention on inequalities in pain care with respect to sexual orientation. It has been shown that sexes do not feel pain the same way [Dance, 2019]. Furthermore, research has also shown that biases with respect to gender and sex are exhibited and perpetuated while using machine learning for healthcare. 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].
- **Adaptivity.** Treating chronic pain also presents additional challenges due to intrinsic or extrinsic factors that change over time such as physiological changes [Gatchel et al., 2014], psychosocial factors [Turk and Gatchel, 2018] and environmental influences Karos et al. [2019]. This dynamic nature of reaction to treatments can render a treatment approach effective at one point but potentially ineffective in the future. Hence there is a critical need to effectively detect the changes in the efficacy of treatments and react to such changes by adapting the recommended treatment accordingly. This non-stationary nature necessitates an adaptive framework for the detection of changes in treatment efficacy and the subsequent adjustment in the recommended treatment, if necessary.

Chronic pain care necessitates treatments over a long period of time taking the feedback from patients into consideration. Thus the formulation of sequential decision making that can learn effectively from user feedback is suitable to model chronic pain care. For example, Komorowski et al. [2018] used sequential decision making solutions to recommend personalized treatments that are on average more reliable than human clinicians. Saria [2018] proposed sequential decision

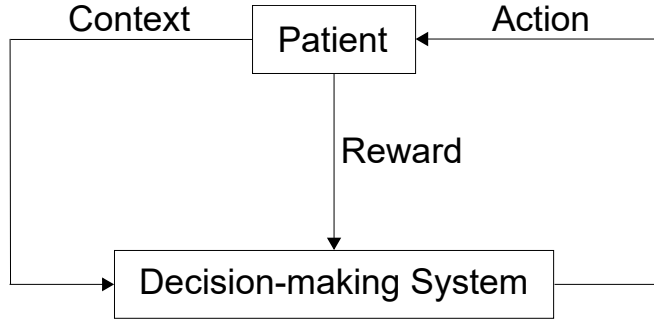


Figure 1: Decision Making for Pain Care

making solutions 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 sequential decision making algorithms to learn effectively from user feedback makes them particularly suitable for making personalized decisions. Moreover, studies have shown that sequential decision making algorithms can make effective decisions even when the feedback changes across time (see for example, our prior works Gajane et al. [2018], Auer et al. [2019] and others such as Besbes et al. [2014], Luo et al. [2018]). As a survey by Zhang and Liu [2021] shows, generic sequential decision making algorithms have also proven effective in addressing fairness considerations. Accordingly, we use the formulation of sequential decision making to model the problem of personalized, fair, and adaptive pain care. In our problem formulation, we consider delayed feedback, which arrives in parts as reactions to treatments can be delayed and spread over time. Moreover, it has been shown that assessments that include delayed feedback contribute to more effective delivery of care in the management of chronic pain [Dworkin et al., 2015, Turk and Okifuji, 2017].

The majority of past research on delayed feedback in sequential decision making assumes that the entire reward of a decision is observed at once, either after some bounded delay [Joulani et al., 2013, Mandel et al., 2015] or after random delays from an unbounded distribution with finite expectation [Gael et al., 2020, Vernade et al., 2017]. In our study, we consider the problem in which the feedback of a decision is spread over an interval with a finite maximum delay value. In our recent work, we studied a basic version of the problem setting which did not include contextual information den Broek et al. [2023]. The same problem was also studied in Romano et al. [2022]. Since the problem setting does not include contextual information, our solutions in den Broek et al. [2023] and the solutions proposed in Romano et al. [2022] are inadequate for the task of providing personalized solutions. In the context of fairness-aware solutions, Huang et al. [2022] and Joseph et al. [2016] propose solutions to achieve certain notions of fairness in personalized sequential decision making. Unlike our proposal, these works assume that the feedback for each decision is immediately available.

2 Problem Setting

We consider the problem setting of *contextual bandits*. This is a variant of the classical multi-armed bandits problem and it has been used toward the goal of personalized decision-making in many domains including healthcare. Accordingly, we consider the formulation of contextual bandits to provide personalized treatment recommendations.

2.1 Stationary Setting: Response to Treatment Does Not Change Across Time

At each step, the decision making system receives contextual information of a patient. This information includes health-relevant attributes and the so-called *protected attributes* such as gender and race. Discrimination based on these protected attributes is prohibited by international legislation. On receiving contextual information, the system chooses a treatment (i.e. an *action*) to administer from the considered options for treatment. The patient response to the administered treatment is given as a numerical value (i.e., a *reward*) to the system. Rewards are typically bounded and a higher reward translates to higher desirability of the associated action. In this context, the more successful treatment is in pain care, the higher will be its reward. Since the results of these treatments may take some time to manifest and may arrive in parts, we assume that rewards are delayed and may arrive via partial rewards that are observed with different delays. The objective here is to learn effective treatments for pain management i.e., learn to select actions that maximize obtained rewards based on the context and received rewards. A representative view of the decision making system is shown in Figure 1.

Since this is a stationary setting, we assume that the rewards of each treatment are drawn from an unknown distribution that does not change with time.

2.2 Non-Stationary Setting: Response to Treatment Changes Across Time

Here, we cannot make the stationarity assumption from the previous subsection since the response to treatment changes across time. The changes could be abrupt or gradual and we will formulate the problem with both these kinds of changes. However, we expect that gradual changes will be more relevant to the problem at hand.

Firstly, we plan to work with the assumption that some characteristics of the changes are known to the algorithm. Then we will further extend the problem setting by removing this assumption.

3 Research Tasks

3.1 Task 1: Fair solutions for stationary setting

Here the goal is to devise fair solutions that can learn efficiently from delayed feedback. One of the possible approaches to satisfy fairness could be to use a *target optimism* term for each action. With each action, the algorithm considers an optimism term indicating how close it will take the algorithm to a target set by the given fairness criteria. To deal with delays, we could use extensions from our approach in an analogous problem in den Broek et al. [2023].

3.2 Task 2: Adaptive fair solutions capable of reacting to changes in response to treatments

Firstly, we plan to propose solutions in the case where some characteristics of the changes are known to the algorithm. For these solutions, we will start with the so-called *sliding-window mechanism* to handle changes. Similar mechanisms have been effectively used to deal with changes in our prior work [Gajane, 2023, Gajane et al., 2019, 2018], so we expect solutions using the sliding-window mechanism to be effective for the problem at hand.

Next, we plan to propose solutions in the case where no characteristics of the changes are known to the algorithm. I have prior experience in designing effective solutions for analogous problems [Auer et al., 2019].

3.3 Task 3: Evaluation of the proposed solution

As a primary evaluation, we propose to evaluate our solutions through simulation studies on previous uses of AI for pain care such as Piette et al. [2022]. The relevant dataset is publicly available under a Creative Commons Attribution 4.0 International licence at Piette [2022]

- **Utility**: The utility of our algorithms can be simply verified by computing the obtained cumulative reward. Higher cumulative rewards translate to a more effective treatment for pain care. Below we describe a more fine-grained verification of our solutions across the three desiderata – personalization, fairness, and adaptivity.
- **Personalization**: We plan to implement solutions from Romano et al. [2022] and our work in den Broek et al. [2023] which will provide non-personalized treatments. We will treat these implementations as a baseline with the aim to study whether our proposed solutions lead to an increase in personalization.

We will shortlist the cases for which recommended treatments by these solutions and the recommended treatments by our proposed solutions differ. We will see the difference in the obtained cumulative rewards for these patients by the baseline and our proposed solutions respectively. For example, if the following histogram shows the treatments suggested by the baseline and our proposed solutions, we will study the cases shown in red for this experiment. The reasoning behind excluding those cases where either of the suggestions is ‘no treatment’ is that they are better counted as mistakes one way or the other rather than missed opportunities for personalization. This experiment will provide us with a more detailed understanding of whether and how our proposed solutions lead to an increase in personalization.

	No treatment	Treatment 1	Treatment 2	Treatment 3
No treatment				
Treatment 1				
Treatment 2				
Treatment 3				

- **Fairness**: In this set of experiments, we divide the patients based on sex, and test if the two groups receive fair treatment considering the following measures which are recommended by medical experts [Rajkomar et al., 2018].
 - Equal outcomes – Equal outcomes refer to the assurance that protected groups have equal benefits in terms of patient outcomes. To see if our proposed algorithm satisfies this measure, we will compare the average reward received by the protected group and the average reward received by the non-protected group.
 - Equal performance – Equal performance refers to the assurance that a model is equally accurate for patients in the protected group and the non-protected group. To see if our proposed algorithm satisfies this measure, we will compare false negative rates (the algorithm recommended no treatment and the patients reported pain in the corresponding feedback) and true negative rates (the algorithm recommended no treatment and the patients reported no pain in the corresponding feedback). We will also compare the scenarios where the patient reported pain despite being recommended a particular treatment.
 - (Proportionately) Equal allocation – Equal allocation ensures that the resources are proportionately allocated to patients in the protected group. Here, we will compute the success rates for each treatment for patients in the protected group and the non-protected group. Then, we will verify if the allocation of each treatment is consistent with the normalized success rate for the patients in the protected group. This might uncover (and motivate us to avoid) situations where a particular treatment has a high success rate for patients in the protected groups and yet it is not proportionally allocated to those patients.
 - Trade-off between fairness and personalization – Here we aim to quantify the trade-off between fairness and personalization. To understand this trade-off, consider the following scenario. Suppose a member of the protected group prefers Treatment 1, however, the fairness measure (e.g., proportionately equal allocation) requires that they should be allocated Treatment 2. For the sake of simplicity, let us assume that the expected probability of success of Treatment 1 and Treatment 2 for the concerned patient is the same. In such cases, personalization and fairness measures like equal allocation can conflict. In this set of experiments, we plan to measure such conflicts. These results might influence the choice of fairness measures in the future.
- **Adaptivity** – Firstly, we will test the performance of our algorithms in situations where the characteristics of the changes are known to the algorithm. For example, we could assume that (an upper bound on) the total variation in the rewards is known to the algorithm. Next, we will test the performance of our algorithms in situations where the characteristics of the changes are not known to the algorithm. Here, we can also check how quickly our algorithms can react to the changes.
- **Utility-fairness trade-off** – For the fairness measures described above, rather than necessitating exact equality between the protected group and the non-protected group, we can allow a small difference parameterized by ϵ . We can use ϵ to regulate the algorithmic behavior and see how it affects the utility. Since a decrease in ϵ translates to a fairer solution, this set of experiments will help us characterize the utility-fairness trade-off.

4 Concluding Remarks

In summary, this research initiative tackles the intricate challenges of delivering effective chronic pain care by adopting a personalized, fair, and adaptive approach through sequential decision-making. This comprehensive strategy addresses patient variability, disparities in pain care, and the evolving nature of treatment efficacy. The planned evaluation, using simulation studies, is designed to make a substantial contribution to the progress of AI-driven chronic pain management. It aims to thoroughly assess utility, personalization, fairness, adaptivity, and the trade-off between utility and fairness.

References

- Jacob N. Ablin and Dan 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>.
- Peter Auer, Pratik Gajane, and Ronald Ortner. Adaptively tracking the best bandit arm with an unknown number of distribution changes. In *Proceedings of the 32nd Conference on Learning Theory (COLT)*, 2019.
- Omar Besbes, Yonatan Gur, and Assaf Zeevi. Stochastic multi-armed-bandit problem with non-stationary rewards. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’14, page 199–207, Cambridge, MA, USA, 2014. MIT Press.

- Davide Cirillo, Silvina Catuara-Solarz, Czuee Morey, Emre Guney, Laia Subirats, Simona Mellino, Annalisa Gigante, Alfonso Valencia, María José Rementeria, Antonella Santucci Chadha, and Nikolaos 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>.
- Amber 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.
- Ronald Van den Broek, Rik Litjens, Tobias Sagis, Nina Verbeeke, and Pratik Gajane. Multi-armed bandits with generalized temporally-partitioned rewards. In *Sixteenth European Workshop on Reinforcement Learning*, 2023. URL <https://openreview.net/forum?id=BcYigwLruz>.
- Robert H Dworkin, Dennis C Turk, Kathleen W Wyrwich, Dorcas Beaton, Charles S Cleeland, John T Farrar, et al. Interpreting the clinical importance of treatment outcomes in chronic pain clinical trials: Imm pact recommendations. *Journal of Pain*, 16(3):219–233, 2015.
- Manegueu Anne Gael, Claire Vernade, Alexandra Carpentier, and Michal Valko. Stochastic bandits with arm-dependent delays. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 3348–3356, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/gael20a.html>.
- Pratik Gajane. Local differential privacy for sequential decision making in a changing environment. In *In AAAI Privacy-Preserving Artificial Intelligence (PPAI)*, 2023.
- Pratik Gajane, Ronald Ortner, and Peter Auer. A sliding-window approach for reinforcement learning in mdps with arbitrarily changing rewards and transitions. In *In the 2nd workshop for Lifelong Learning: A Reinforcement Learning Approach (LLARLA)*, 2018.
- Pratik Gajane, Ronald Ortner, Peter Auer, and Csaba Szepesvári. Autonomous exploration for navigating in non-stationary cmeps. *CoRR*, abs/1910.08446, 2019. URL <http://arxiv.org/abs/1910.08446>.
- Robert J Gatchel, Donald D McGeary, Cindy A McGeary, and Ben Lippe. Interdisciplinary chronic pain management: past, present, and future. *The American psychologist*, 69(2):119–130, 2014. ISSN 0003-066X. doi: 10.1037/a0035514. URL <https://doi.org/10.1037/a0035514>.
- Wen Huang, Kevin Labille, Xintao Wu, Dongwon Lee, and Neil Heffernan. Achieving user-side fairness in contextual bandits. *Human-Centric Intelligent Systems*, 2(3):81–94, Dec 2022. ISSN 2667-1336. doi: 10.1007/s44230-022-00008-w. URL <https://doi.org/10.1007/s44230-022-00008-w>.
- Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. In *Advances in Neural Information Processing Systems*, 2016.
- Pooria Joulani, András György, and Csaba Szepesvári. Online learning under delayed feedback. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML’13*, 2013.
- K. Karos, A. C. D. C. Williams, A. Meulders, and J. W. Vlaeyen. Pain as a threat to the social self: A motivational account. *Pain*, 160(8):1854–1858, 2019.
- Matthieu Komorowski, Leo A. Celi, Omar Badawi, Anthony C. Gordon, and A. Aldo 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>.
- Haipeng Luo, Chen-Yu Wei, Alekh Agarwal, and John Langford. Efficient contextual bandits in non-stationary worlds. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 1739–1776. PMLR, 06–09 Jul 2018. URL <https://proceedings.mlr.press/v75/luo18a.html>.
- Travis Mandel, Yun-En Liu, Emma Brunskill, and Zoran Popović. The queue method: Handling delay, heuristics, prior data, and evaluation in bandits. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Maria Matsangidou, Andreas Liampas, Melpo Pittara, Constantinos S. Pattichi, and Panagiotis 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>.

- John 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>.
- John D. Piette, Sean Newman, Sarah L. Krein, Nicolle Marinec, Jenny Chen, David A. Williams, Sara N. Edmond, Mary Driscoll, Kathryn M. LaChappelle, Marianna Maly, H. Myra Kim, Karen B. Farris, Diana M. Higgins, Robert D. Kerns, and Alicia 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>.
- Alvin Rishi Rajkomar, Greg Corrado, Marshall Chin, Michael Howell, and Mila Hardt. Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine*, 2018.
- Luca Roggeveen, Ali el Hassouni, Jonas Ahrendt, Tingjie Guo, Lucas Fleuren, Patrick Thorat, Armand RJ Girbes, Mark Hoogendoorn, and Paul WG 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>.
- Giulia Romano, Andrea Agostini, Francesco Trovò, Nicola Gatti, and Marcello Restelli. Multi-armed bandit problem with temporally-partitioned rewards: When partial feedback counts. *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*, 2022. doi: 10.24963/ijcai.2022/472.
- Suchi 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>.
- Gordy Slack. Reforming pain. *Anesthesiology Pain Management*, 2022. URL <https://stanmed.stanford.edu/individualizing-pain-treatments/>.
- Isabel Straw and Honghan 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>.
- Dennis C Turk and Robert J Gatchel. *Psychological Treatment of Chronic Pain: Clinical Handbook*. Guilford Publications, 2018.
- Dennis C Turk and Akiko Okifuji. Assessment of patients with chronic pain. In *Psychological approaches to pain management: A practitioner’s handbook*, pages 43–70. Guilford Press, 3rd edition, 2017.
- Claire Vernade, Olivier Cappé, and Vianney Perchet. Stochastic Bandit Models for Delayed Conversions. In *Conference on Uncertainty in Artificial Intelligence*, August 2017. URL <https://hal.science/hal-01545667>.
- Xueru Zhang and Mingyan Liu. *Fairness in Learning-Based Sequential Decision Algorithms: A Survey*, pages 525–555. Springer International Publishing, Cham, 2021. ISBN 978-3-030-60990-0.
- Zachary Zimmer, Kathryn Fraser, Hanna Grol-Prokopczyk, and Anna Zajacova. A global study of pain prevalence across 52 countries: examining the role of country-level contextual factors. *PAIN*, 163(9), 2022.